# Belle Computing Activities at CNAF and migration to token

CdG Tier-1

18 February 2022

Dr. Silvio Pardi

# Belle II Infrastructure

19 Countries

55 Sites Registered in DIRAC who replied to the questionnaire

5 Site with Tape systems

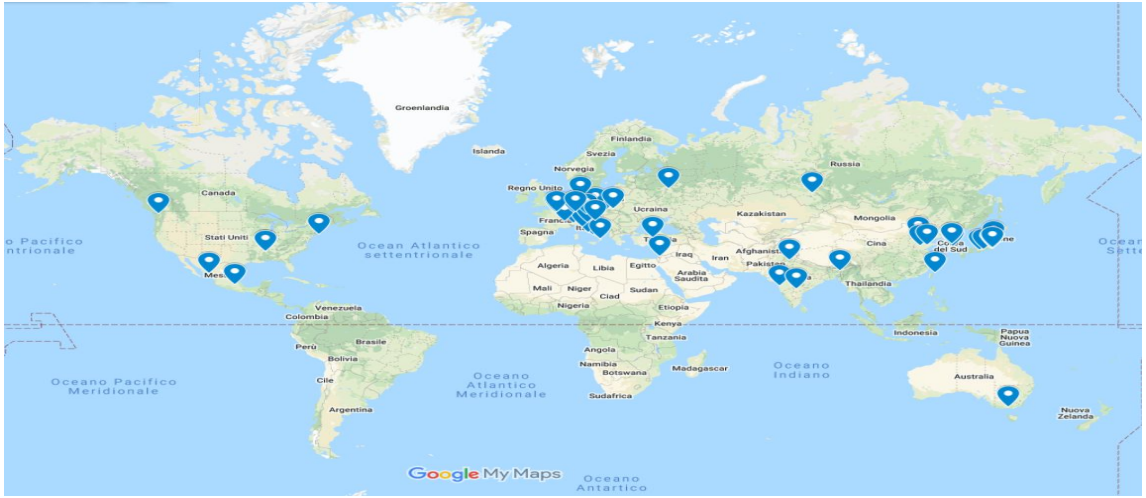


Australia
Austria
Canada
Czech Republic
China
France
Germany
India
Israel
ITALY
Japan
Mexico
Poland
Russia
Slovenia
South Korea
Taiwan
TURKEY
USA

# Site Report 2021

| Resources | CPU Pledged (kHS06) | CPU Pledged jobslots | CPU Opportunistic (kHS06) | CPU Opportunistic jobslots | Total possible CPU (kHS06) | Total possible Jobslots | Storage DISK (TB) | Tape (TB) |
|---|---|---|---|---|---|---|---|---|
| Production | 452 | 31.484 | 310 | 25.377 | 762 | 56.861 | 13.555 | 10.049 |

| Resources at BNL and DESY | CPU Pledged (kHS06) | CPU Pledged jobslots | CPU Opportunistic (kHS06) | CPU Opportunistic jobslots | Total posCPU (kHS06) | Total Jobslots | Storage DISK (TB) | Tape (TB) |
|---|---|---|---|---|---|---|---|---|
| Calib/Recalibration | 36 | 3130 | 0 | 0 | 36 | 3130 | 500 | 600 |

# CPU Resources (as of 1 October 2021)

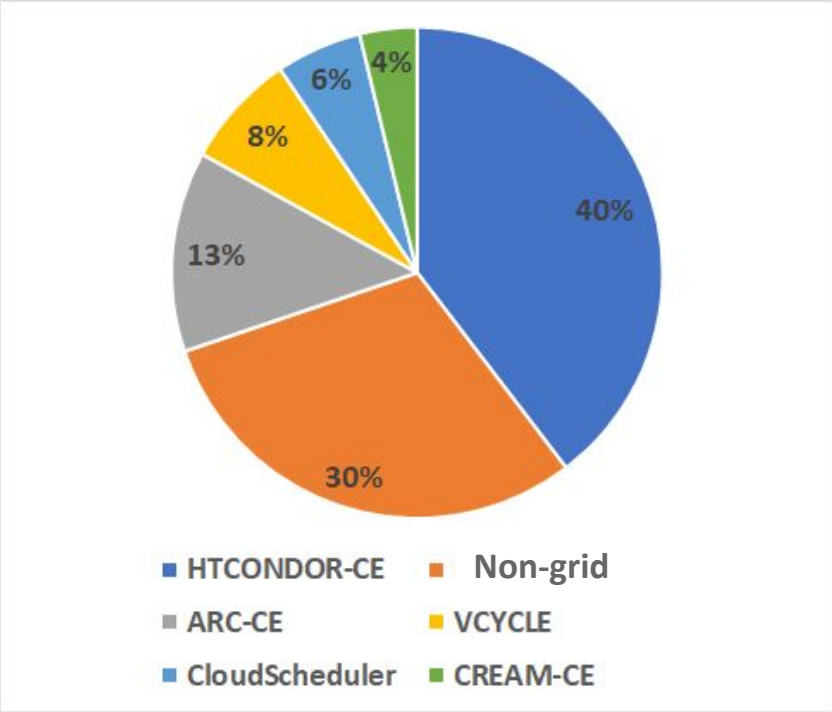


- 55 Different DIRAC Endpoints
  - 24 Sites Providing Pledged CPUs
  - 12 Sites Pledged + Opportunistic
  - 18 Sites Opportunistic Only

| Batch System | # of endpoints |
|---|---|
| HTCondor | 26 |
| SLURM | 8 |
| SGE | 8 |
| PBS | 7 |
| cloud | 5 |
| LSF | 3 |

# Storage Resources (as of 1 October 2021)

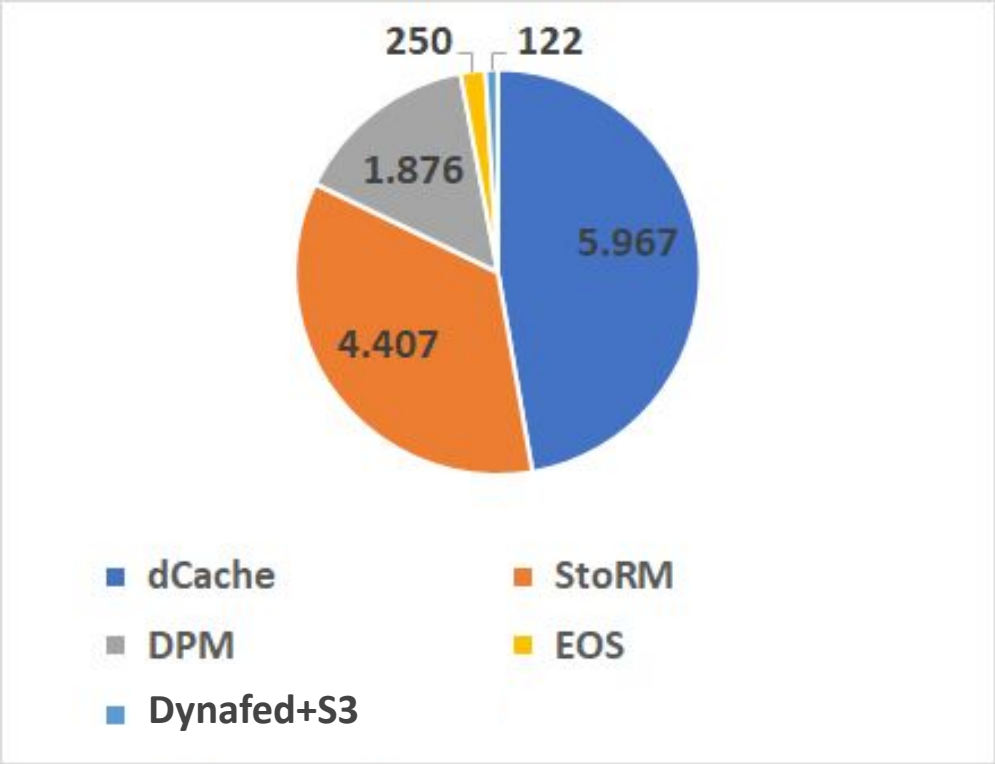


Highlight: 10 DPM Storages managing 1.8 PB
Need a long term strategy for storage migration.

# Siti Italiani (tutto il pledge è fornito)

| | CPU Pledge (kHS06) | CPU Opport. (kHS06) | Storage (TB) | Tape (TB) |
|---|---|---|---|---|
| CNAF | 27 | | 650 | 650 |
| Cosenza | 1 | | | |
| Napoli | 12 | 10 | 390 | |
| Pisa | 8 | 10 | 200 | |
| Torino | 6 | 24 | 350 | |
| Frascati | | 0,5 | 11 | |
| LNL | | 1 | | |
| Roma3 | | 2 | 2 | |
| **TOTALE** | **54 kHS06** | **47,5 kHS06** | **1.603 TB** | **650 TB** |

# Resoconto risorse disponibili per Belle II e richieste

RISORSE DISPONIBILI AL 1 Ottobre 2021

| SITE | CPU | STORAGE | TAPE |
|---|---|---|---|
| CNAF | 27kHS06 | 650TB | 650TB |
| Napoli+Cosenza | 13kHS06 | 390TB | |
| Pisa | 8kHS06 (4 da installare) | 200TB da configurare | |
| Torino | 6kHS06 | 350TB da installare | |
| **TOTALE OGGI** | **54kHS06** | **1.590TB (1.040 Utilizzabile)** | **650TB** |

RISORSE RICHIESTE PER IL 2022

| SITO | ACQUISTI | Origine fondi |
|---|---|---|
| CNAF | 300 TB TAPE | Gia forniti in anticipo |
| CNAF | 170 TB Disk | Dalla giunta. |
| Napoli | 200 TB Disk | A insistere su fondi di progetto IBISCO |

# UTILIZZO CPU

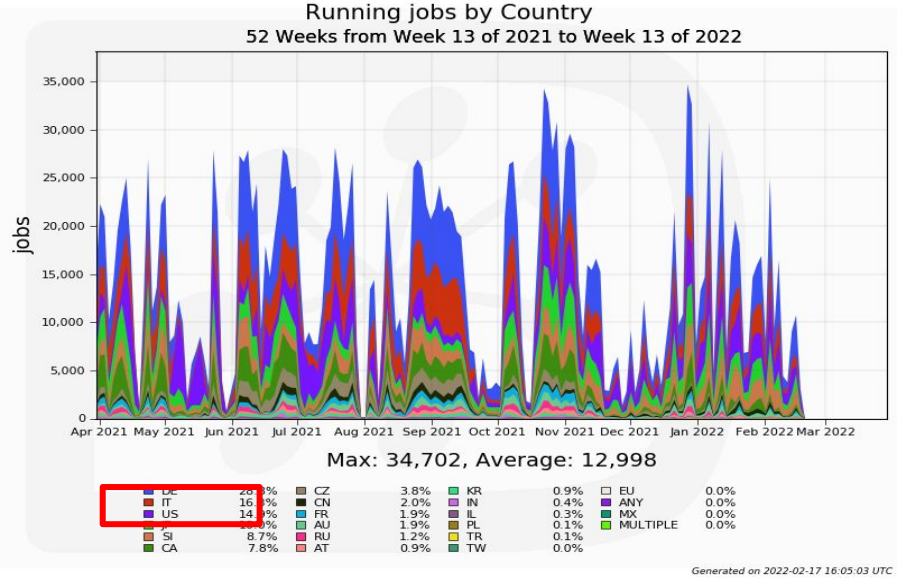


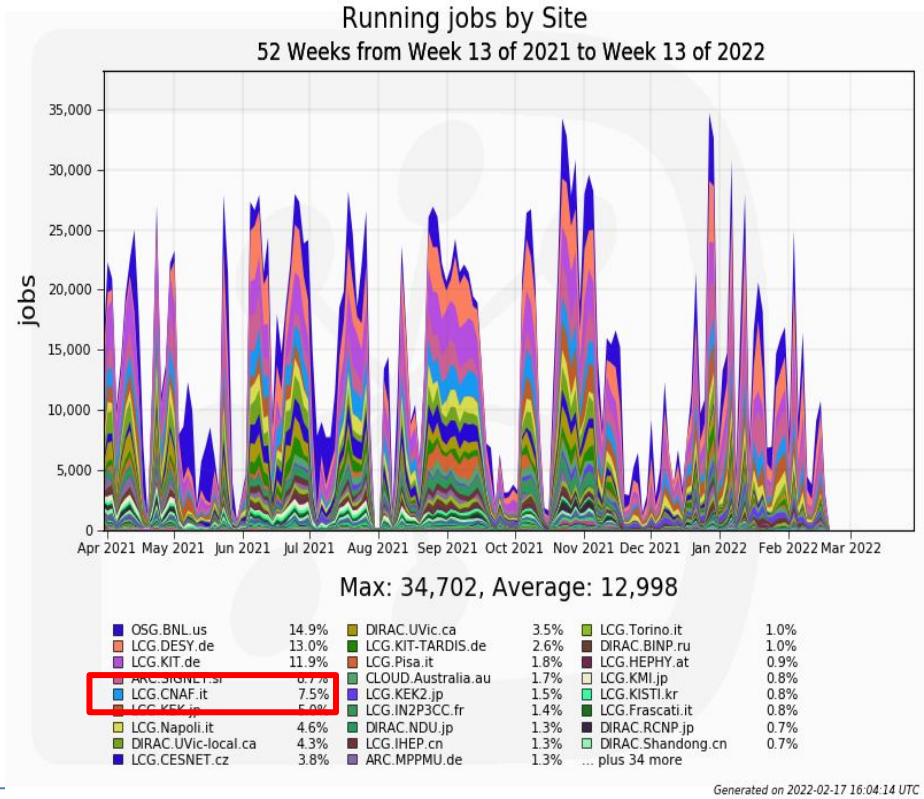


Goal per l'italia è 12% nel 2021 e 11% nel 2022

# UTILIZZ

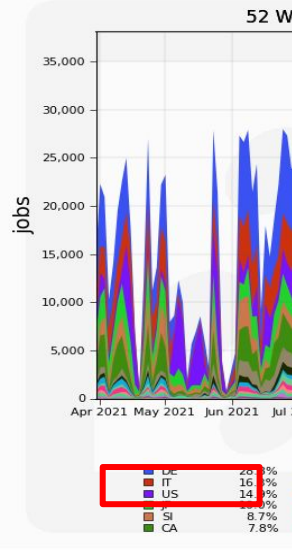


Goal per l'italia

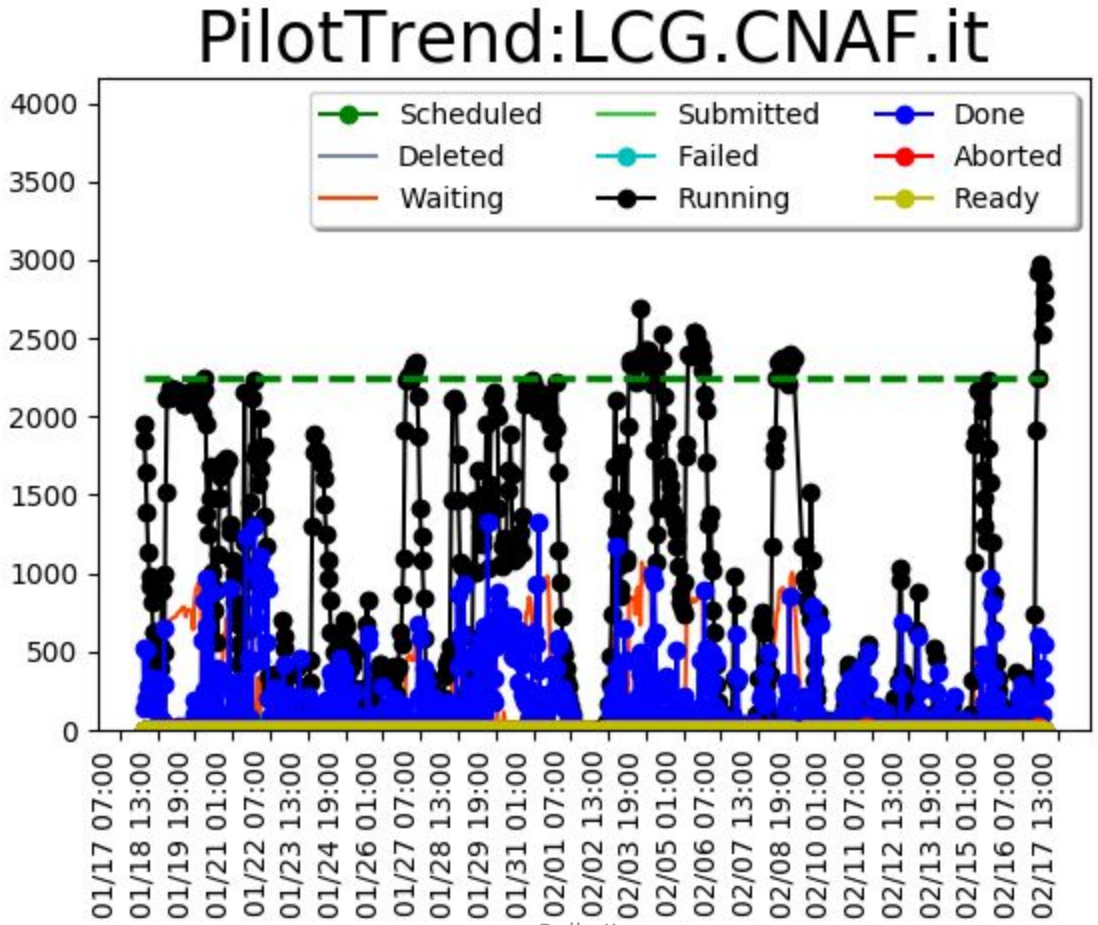


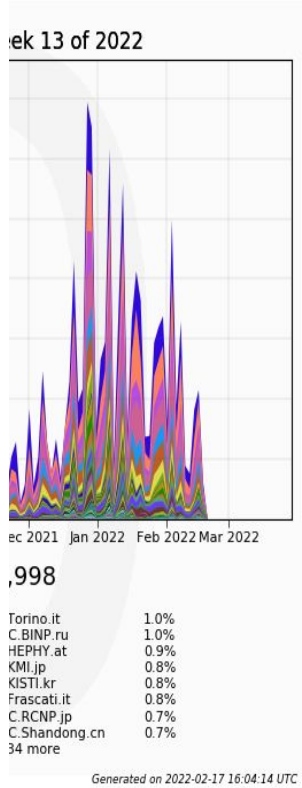

# STORAGE AND TAPE

DISK

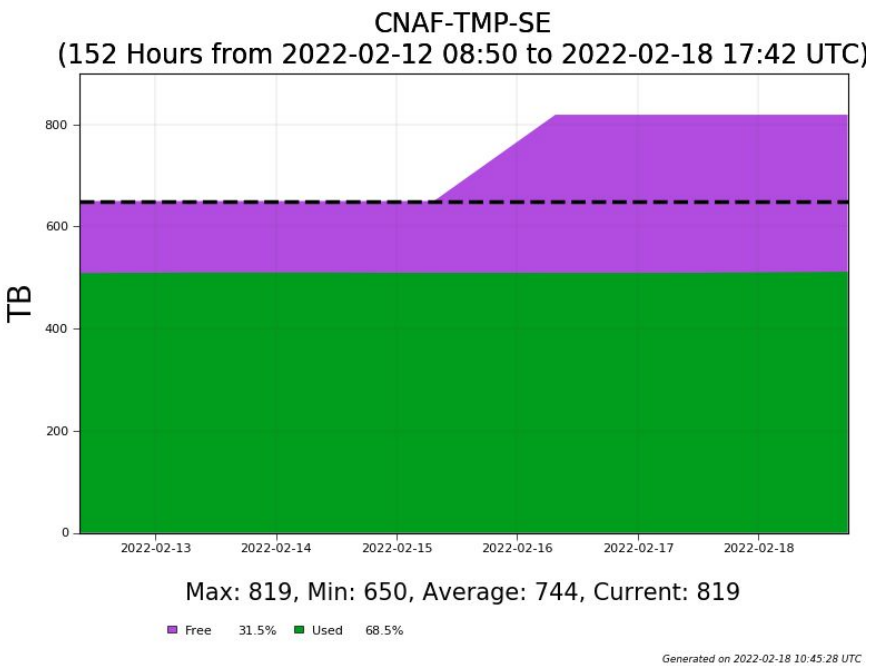


TAPE

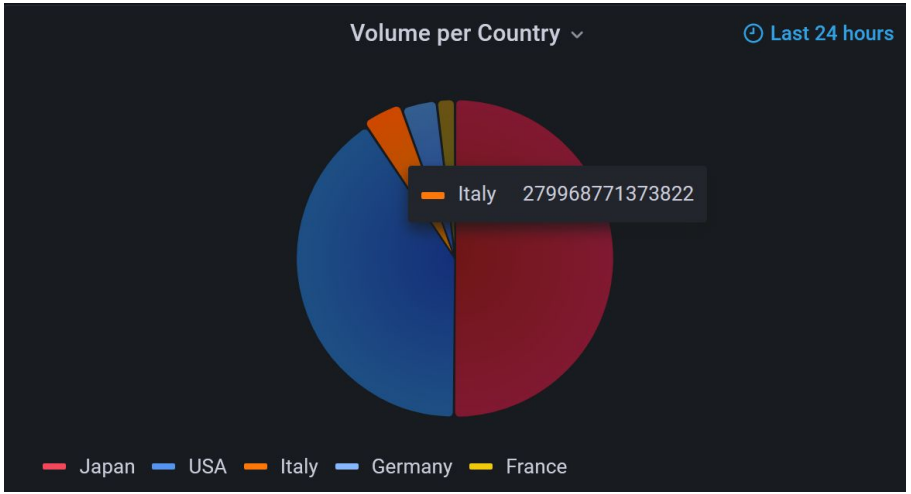

# Belle II Sites Meeting

Topics
- WebDAVS implementation on the storage
- Token Based Authentication for Cluster and Storages

**Meeting 4th February 2022 divided in two session**

-First round  2:00 CET (10:00 JST)
-Second round 7:30 CET (15:30 JST)

Agenda
https://kds.kek.jp/event/41010/
Minute
https://docs.google.com/document/d/1ytk6co4GuClmNhWM8Yucm2q-imdupHn1qrnwAI03ogs/edit#

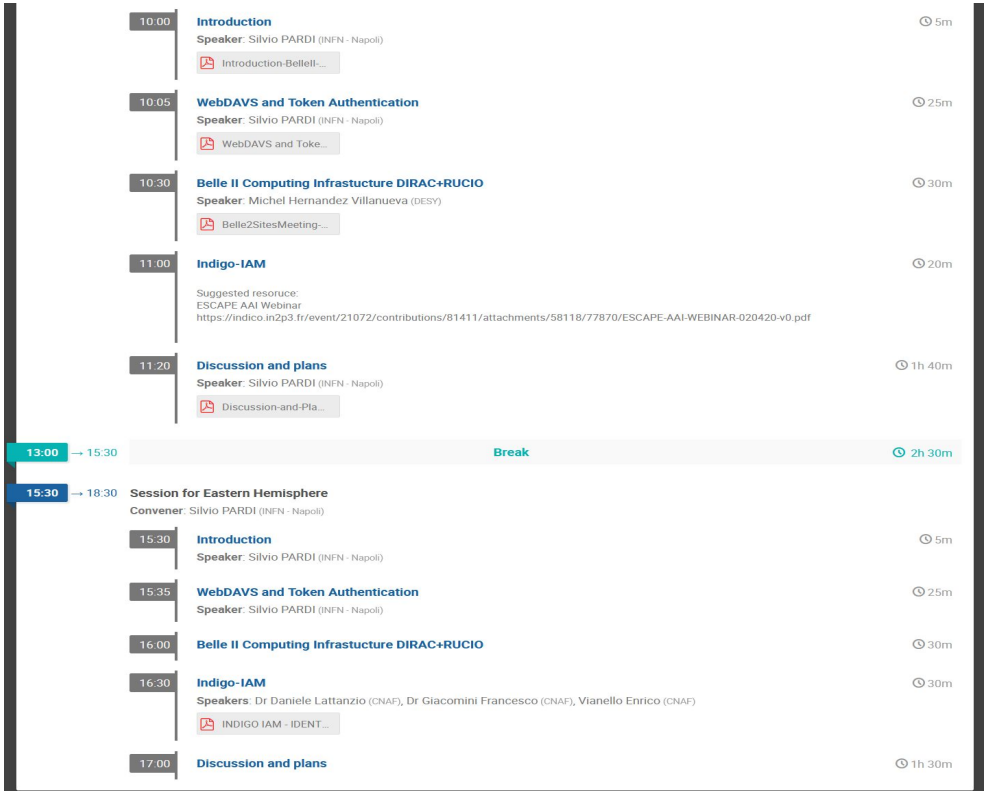

# Plan for WebDAV Deployment

Milestone 1:

Use JSON file (more than 90% done in production)

Milestone 2:

Use WebDAV (50% configured in production activity just started)

Milestone 3:

Third-party-copy Challenge (75% working at least in a modality push or pull)

Timeline: Make everything's working by June 2022

# Activation of WebDAV at CNAF as Access Protocol

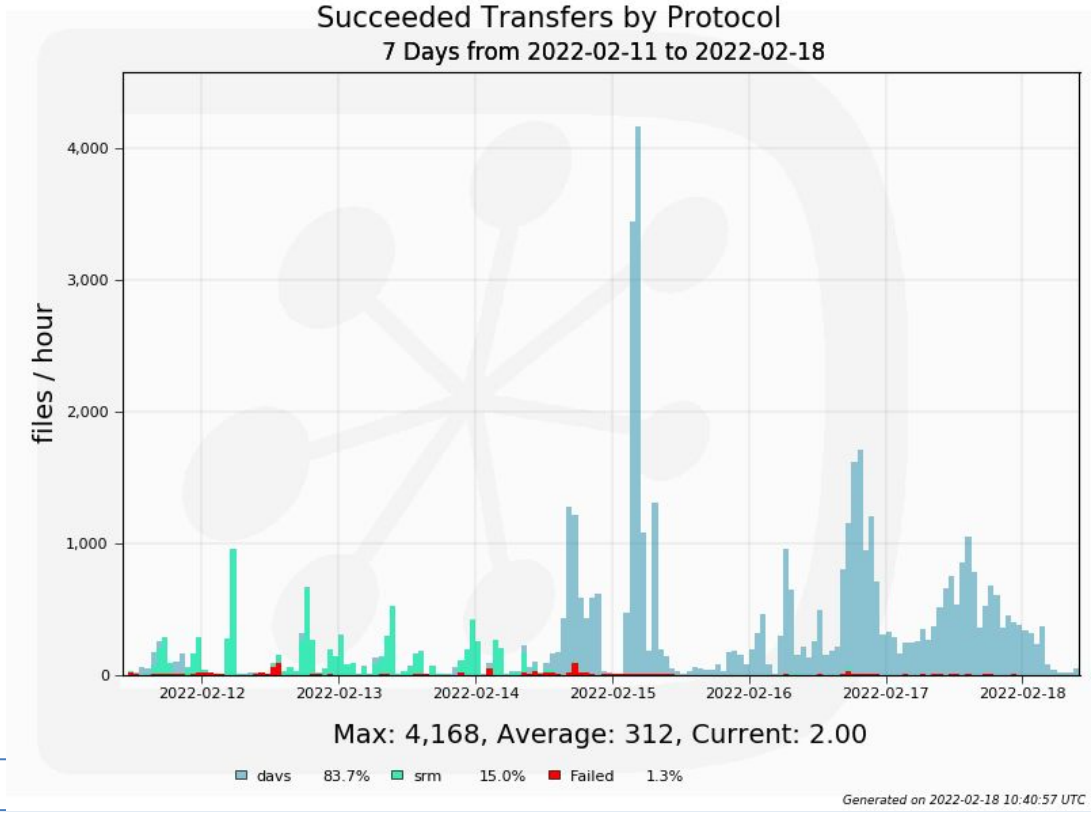

# Third-party-copy

**Third Party Copy test summary**

**FTS SERVER fts.usatlas.bnl.gov**

**Thu Feb 10 17:43:46 CET 2022 - TEST HISTORY<**

**Green if Pull and Push tansfers have been completed successfully, Yellow if at least a Pull or a Push transfer have been completed successfully, Red if Pull and Push tansfers failed**

**DESTINATION**

| | BNL-TMP-SE | CNAF-TMP-SE | DESY-TMP-SE | IN2P3CC-TMP-SE | KIT-TMP-SE | KEK-DISK-TMP-SE | Napoli-TMP-SE | SIGNET-TMP-SE | UVic-TMP-SE1 | UVic-TMP-SE2 | Australia-SE | HEPHY-TMP-SE | IHEP-TMP-SE | CESNET-TMP-SE | LAL-TMP-SE | IPHC-TMP-SE | MPPMU-TMP-SE | TAU-TMP-SE | Pisa-TMP-SE | Torino-TMP-SE | ROMA3-TMP-SE | Frascati-TMP-SE | CYFRONET-TMP-SE | NTUCC-TMP-SE | ULAKBIM-TMP-SE |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| BNL-TMP-SE | | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | Pull FINISHED | ERROR | Pull FINISHED | ERROR | ERROR | Pull FINISHED | ERROR | Pull FINISHED | Pull FINISHED | Pull FINISHED | 1 test ERROR | Pull FINISHED | ERROR | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | Pull FINISHED |
| CNAF-TMP-SE | FINISHED | | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | ERROR | ERROR | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | FINISHED | FINISHED |
| DESY-TMP-SE | FINISHED | FINISHED | | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | Push FINISHED | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | 1 test ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| IN2P3CC-TMP-SE | FINISHED | FINISHED | FINISHED | | FINISHED | ERROR | FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | Push FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| KIT-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | Push FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | 1 test ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| KEK-DISK-TMP-SE | ERROR | Push FINISHED | ERROR | ERROR | ERROR | | ERROR | ERROR | Push FINISHED | ERROR | ERROR | ERROR | ERROR | Push FINISHED | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR |
| Napoli-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| SIGNET-TMP-SE | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | 1 test ERROR | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED |
| UVic-TMP-SE1 | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | | Push FINISHED | ERROR | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | UNSHED | FINISHED | ERROR | FINISHED |
| UVic-TMP-SE2 | Push FINISHED | Pull FINISHED | ERROR | FINISHED | Pull FINISHED | ERROR | Pull FINISHED | ERROR | Pull FINISHED | | ERROR | Pull FINISHED | ERROR | Push FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | ERROR |
| Australia-SE | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | Pull FINISHED | ERROR | Pull FINISHED | ERROR | | Pull FINISHED | ERROR | Pull FINISHED | Pull FINISHED | Pull FINISHED | 1 test ERROR | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | Pull FINISHED |
| HEPHY-TMP-SE | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | ERROR | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED |
| IHEP-TMP-SE | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | Pull FINISHED | ERROR | ERROR | ERROR | Pull FINISHED | ERROR | ERROR | Pull FINISHED | | Pull FINISHED | ERROR | ERROR | ERROR | Pull FINISHED | ERROR | Pull FINISHED | Pull FINISHED | ERROR | ERROR | ERROR | ERROR |
| CESNET-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | ERROR | Pull FINISHED | Push FINISHED | | FINISHED | FINISHED | Pull FINISHED | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| LAL-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | | FINISHED | ERROR | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | FINISHED |
| IPHC-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | | ERROR | FINISHED | Pull FINISHED | FINISHED | UNSHED | 1 test ERROR | Push FINISHED | ERROR | FINISHED |
| MPPMU-TMP-SE | 1 test ERROR | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | ERROR | 1 test ERROR | 1 test ERROR | 1 test ERROR | ERROR | 1 test ERROR | Push FINISHED | Push FINISHED | Push FINISHED | | 1 test ERROR | ERROR | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | 1 test ERROR |
| TAU-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | Pull FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | 1 test ERROR | | ERROR | Pull FINISHED | UNSHED | 1 test ERROR | FINISHED | ERROR | FINISHED |
| Pisa-TMP-SE | 1 test ERROR | ERROR | 1 test ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | 1 test ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | ERROR | | ERROR | ERROR | ERROR | 1 test ERROR | ERROR | ERROR |
| Torino-TMP-SE | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | ERROR | | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED |
| ROMA3-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | ERROR | Pull FINISHED | | FINISHED | FINISHED | ERROR | FINISHED |
| Frascati-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | 1 test ERROR | Pull FINISHED | UNSHED | | FINISHED | ERROR | FINISHED |
| CYFRONET-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | Pull FINISHED | FINISHED | Pull FINISHED | Pull FINISHED | UNSHED | FINISHED | | ERROR | FINISHED |
| NTUCC-TMP-SE | Push FINISHED | Push FINISHED | FINISHED | Push FINISHED | Push FINISHED | ERROR | FINISHED | ERROR | ERROR | Push FINISHED | Push FINISHED | Push FINISHED | ERROR | Push FINISHED | FINISHED | FINISHED | ERROR | Push FINISHED | ERROR | Push FINISHED | Push FINISHED | Pull FINISHED | FINISHED | | FINISHED |
| ULAKBIM-TMP-SE | FINISHED | FINISHED | FINISHED | FINISHED | FINISHED | ERROR | FINISHED | ERROR | FINISHED | Push FINISHED | Push FINISHED | Pull FINISHED | ERROR | FINISHED | FINISHED | FINISHED | 1 test ERROR | FINISHED | ERROR | Pull FINISHED | UNSHED | FINISHED | FINISHED | ERROR | |

# Token Based AAI implementation in WLCG [2]

WLCG Road Map up to 2024

https://twiki.cern.ch/twiki/bin/view/LCG/WLCGAuthorizationWG
https://docs.google.com/document/d/11fcZU8fEsfjDiSkjh95nVr4tNXLPCA_xwr2SwriBpiw/edit#
https://indico.cern.ch/event/876810/

| | | | | |
|---|---|---|---|---|
| M.2 | June 3/4 2021 | WLCG hosts "CE and pilot factory hackathon" | None | Pilot framework providers |
| M.3 | July 2021 | Production IAM Instance(s) Available for at least 1 LHC experiment, likely CMS and possibly ATLAS | None | WLCG Ops, IAM, CERN IT |
| M.4 | Oct 2021 | Pilot job submissions may be performed with tokens. | M.3 | Experiments, pilot framework providers, OSG/EGI, sites, Monitoring |
| M.5 | Dec 2021 | VOMS-Admin shutoff for CMS. IAM is sole authz provider for those (including for VOMS server) | M.3 | WLCG Ops, CERN IT |
| M.6 | Feb 2022 | OSG ends support for the Grid Community Toolkit | M.1, M.4 | OSG |

| | | | | |
|---|---|---|---|---|
| M.7 | Mar 2022 | All storage services provide support for tokens | M.1 | WLCG Ops, Storage providers |
| | ? | All VOs shut off VOMS-Admin | | |
| | Sept 2022 | End of HTCondor support for GSI Auth (link) | | |
| M.8 | Oct 2022 | Rucio transfers performed with token auth in production | M.7 | Rucio, Experiments |
| M.9 | Mar 2023 | Experiments stageout & data reads performed via tokens. | M.7 | Experiments |
| M.10 | Mar 2024 | X.509 client auth becomes optional. | M.9, M.8, M.4 | Experiments |

# Token authentication for Computing resources

- HTC-Condor-CE version 9.X support Token based Authentication
- ARC-CE
  https://docs.google.com/presentation/d/11lczxLFmS40u4rZ1Zw8eJZZYkATEugUlIcZ3FWsPP9g/edit#slide=id.gc6f919934_0_5
- DIRAC not yet

# Token Based Authentication

At Now Belle II does not have any service with token based authentication:

## DIRAC

DIRAC is not ready

- Although there have already been developments since a while: OIDC/OAuth and SSO (EGI Check-in)
  - Lytovchenko, May 2019, DUW https://indico.cern.ch/event/756635/contributions/3383243/
- There will be some experimental support for tokens in DIRAC V8.0
  - not necessarily everything can be done with tokens
- We are still with v7r1
  - going to v7r2 in the coming months, and then to v7r3, with many jumps -- DIRACOS2, python3, ...

## Rucio

- Rucio is ready, though the current workflow may be refined
- https://indico.cern.ch/event/1037922/contributions/4528620/

## AMGA

# Towards a Token Based Authentication Challenge for Belle II

Proposal for first step tests

1. Setup a Token-based AAI service Indigo-IAM based
2. Send a job to a CE with token authentication (Low level test)
3. Copy a file to an SE with token authentication (Low level test)

# Token Based Authentication WG

BNL: HTCondor-CE and dChache

CNAF: indigo-IAM and STORM

DESY: HTCondor-CE dCache

KIT:  HTCondor-CE, dCache and xrootd

SIGNET: ARC-CE and dCache

IAM service in preparation

Next Step: Start some tests by early March

Milestone: Test done by June 2022

BACKUP

# Data Management with Rucio

- Rucio is an advanced Distributed Data Management System initially developed for the ATLAS experiment.
- Last year, our Distributed Data Management system was successfully integrated with Rucio.
- Rucio takes care of interacting with the File Transfers Service (FTS) to move around files, and takes care of the physical deletion of the file replicas.
- It is able to use multiple protocols to interact with the Storage Element, in particular WebDAV.