

# **COMPUTING FOR ALICE IN THE CZECH REPUBLIC in 2019-2022**

*D. Adamová*

*with the help and support of :*

*J. Chudoba, A. Mikula, J. Hampl,*

*L. Míča, M. Svatoš , P. Šesták , J. Uhlířová, P. Vokáč*

*Budapest 2022*

# Outline

**Status of the WLCG Tier-2 site in Prague**

**Report on the delivery of mandatory resources**

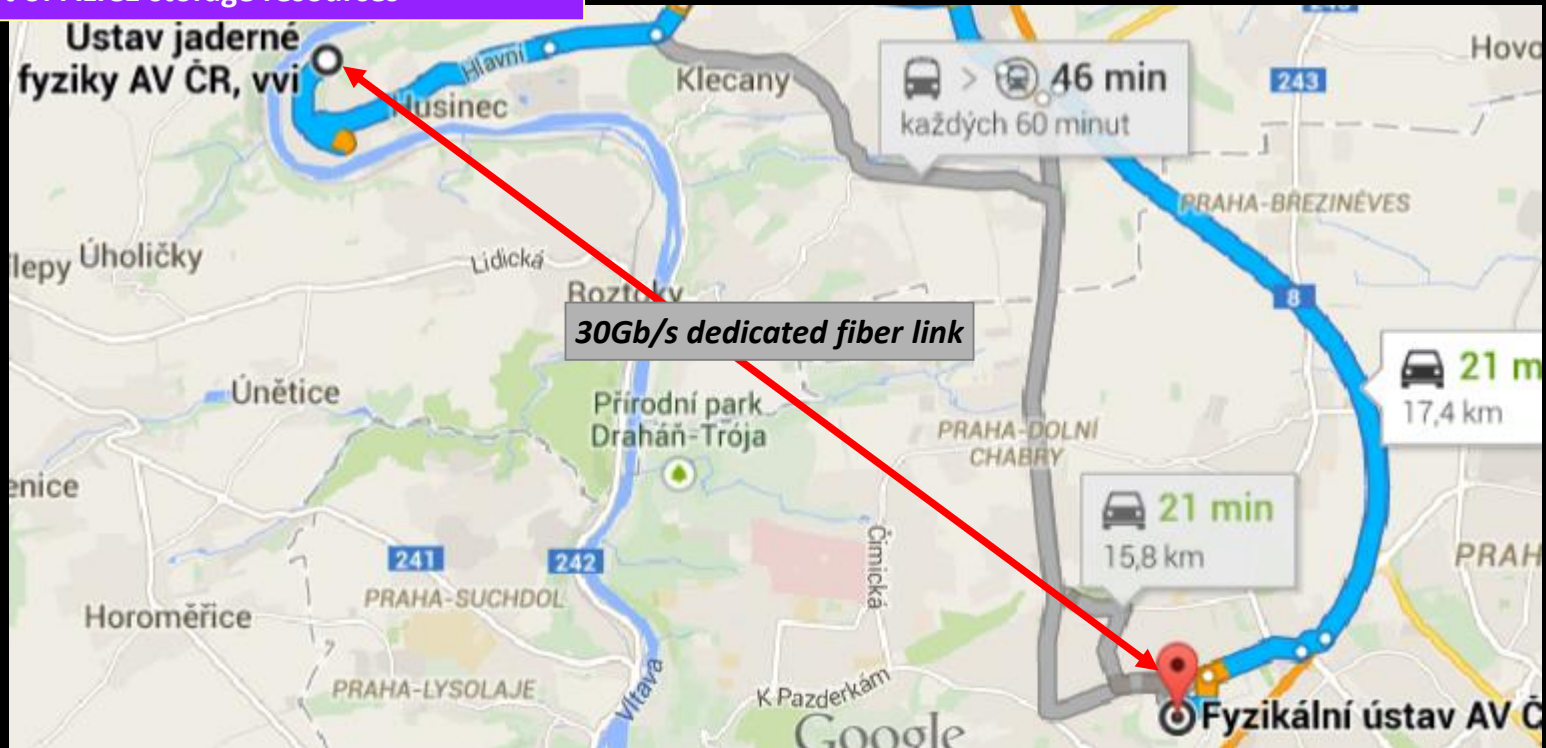
**ALICE operations in 2019-2022**

**A couple of various issues**

**Summary and Outlook**

# The geographical layout

Nuclear Physics Institute AS CR (NPI)  
a large ALICE group, no ATLAS involvement  
A part of ALICE storage resources



30Gb/s dedicated fiber link

Originally, the xrootd storage servers for ALICE were only at NPI. Subsequently, xrootd servers were also put in operation at FZU and also the redirector is there.

Institute of Physics AS CR (FZU)  
Regional computing center  
WLCG Tier-2 site pragueicg2  
All the CPU resources for ALICE and ATLAS  
One part of storage resources for ALICE  
Quite small ALICE group, much larger ATLAS community

# Extended geographical layout



The petascale HPC center IT4I in Ostrava. Currently hosting a cluster Karolina, acquired as a part of the EuroHPC Joint Undertaking. Installed in 2021. Used by pragueLCG2 as an external resource for the ATLAS experiment.

# HEP Computing in Prague: WLCG site prague1cg2 (a.k.a. the farm GOLIAS)

- **A national computing center for processing data from various HEP experiments**
  - A distributed site, the central part located in the Institute of Physics (FZU) in Prague
  - Basic infrastructure already in 2002
  - One of the 13 sites participating in the first LCG campaign in 2003
- **Certified as a Tier2 center of LHC Computing Grid (prague1cg2)**
  - Collaboration with various other Grid projects.
- **April 2008, WLCG MoU signed by Czech Republic (ALICE+ATLAS).**
- **Very good network connectivity:**

Multiple dedicated 10 – 100 Gb/s connections to collaborating institutions, 100 Gb/s connection to LHCONE, to be upgraded to 2\*100 Gb/s soon.
- **Provides computing services for** ATLAS + ALICE, DUNE, AUGER, NOVA, CTA, Fermilab, Astrophysics ...



# Current numbers

- 1 batch system (HTCondor)
- 2 main WLCG VOs: ALICE, ATLAS
- ~ 13,000 job slots on site + 1500 job slots at Charles University in Prague
- ~ 11.07 PB in total on disk storage on site and at NPI ( dCache, XRootD) plus ~1 PB on NFS servers.
- Regular upscale of resources on the basis of various financial supports, mainly the academic grants.
- Monitoring: Prometheus + Grafana, Icinga for alerting.
- Configuration management by Puppet
- Provisioning and SW management by Foreman

About 40 VMs provide core and grid services on 3 KVM hypervisors.




# ALICE disk XRootD Storage Element

## ALICE::Prague::SE

- ~ 3.3 PB of net disk capacity in total
- a distributed storage cluster:
  - Redirector + 2 clients @ FZU,
  - 4 clients @ NPI Rez
- 30 Gb/s *dedicated* connection from FZU to the NPI ALICE storage, to be scaled up to 60 Gb/s

STORAGE BOXES rack at NPI:  
JBOD of lids7, JBOD of lids9, 2xUPS, JBOD  
of lids8 (disk.ujf.cas.cz)

Server xrootd5.farm.particle.cz:  
422 TiB, at FZU



# New hardware 2019 - 2022

## Total contributions for the whole site - CLUSTERS

Number of WNs: 111  
Number of cores: 3552  
Number of slots: 7104  
Total RAM: 24024 GB

## Total contributions for the whole site - STORAGE

Total new NET disk capacity, without the new nfs servers **6802 TB**  
NET disk capacity of the new nfs servers ~ 900 TB  
NET disk capacity of the new XrootD servers 2374 TB

## Total contributions for the whole site - NETWORK

Replacement of Cisco Catalyst 9500 with Nexus 93180 at the main site as the site router->  
this up-scales the local network throughput.

Completed migration at NPI to Cisco NX93180 and the deployment of IPv6.



# New hardware for ALICE 2019 - 2022

- Three new disk servers with the total NET capacity of 2374 TB:


lids7, lids8, lids9 (disk.ujf.cas.cz)

There were new CPU clusters put in operation, in total 7104 slots. ALICE is officially granted 20% of the 'CERN' share which is provided for ALICE + ATLAS.

At NPI, migration to a new switch Cisco NX93180 and to 30 Gb/s traffic between the main site and NPI

In progress: upgrade of the server room at NPI with new server rack enclosures, 10 kVA UPSes and upgraded power supply

On the picture from top: JBOD LIDS7, head server LIDS7 , router Cisco Nexus, router Cisco C-4500, server LIDS9, UPS, admin and head servers of LIDS, management switch



# Recent changes in the site network

Link 100 Gb/s to LHCONE and 40 Gb/s to the public network

Replacement of Cisco Catalyst 9500 with Nexus 93180 as the site router which provided for a better throughput of the local network

Re-modeling of the internal network structure so that the main internal connectivity is 100 Gb/s

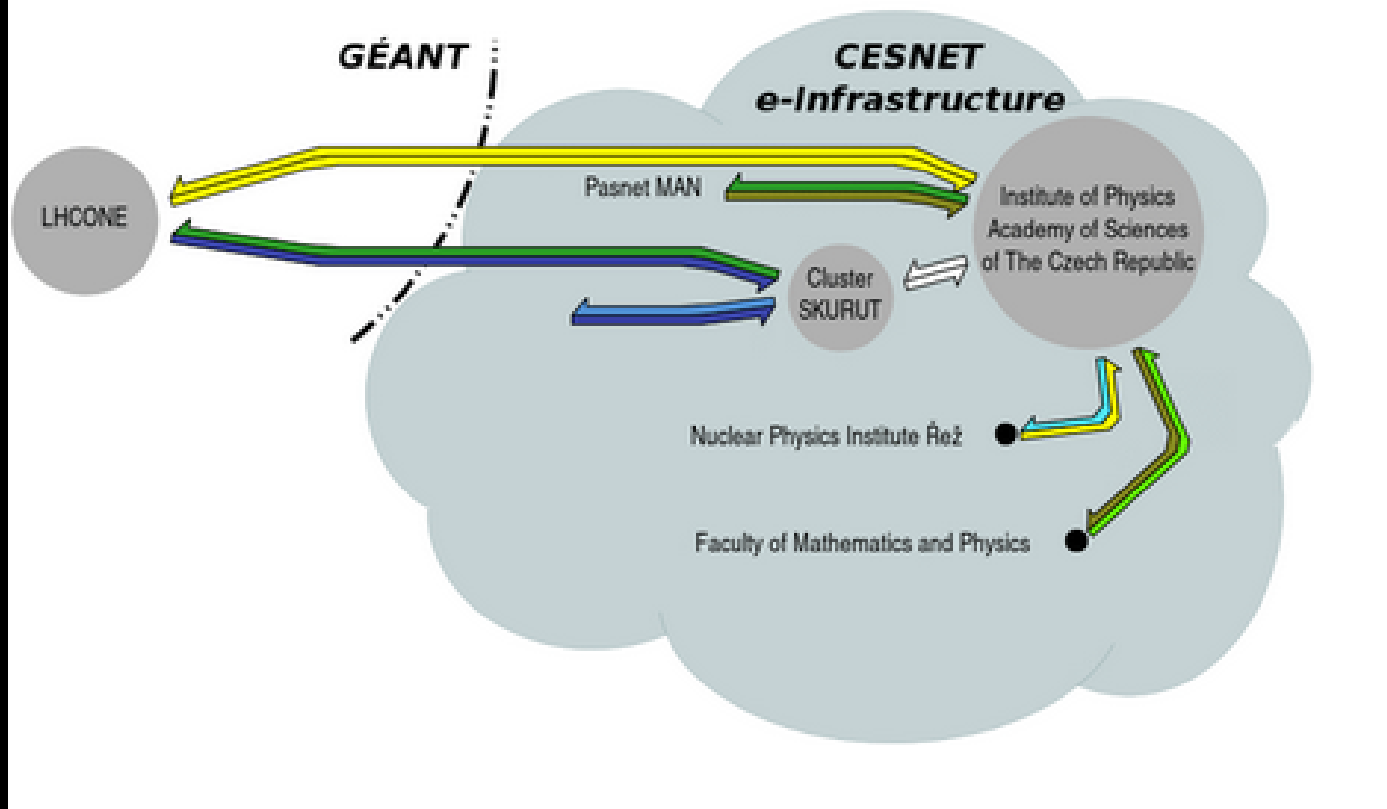
## In preparation

Upgrade of POP at NPI and eventual upscale of the link NPI – FZU to 100Gb/s


Increase of the site connection to LHCONE to 2\*100 Gb/s and to the public network to 100 Gb/s

# Large Infrastructure CESNET: Institute of Physics Academy of Sciences of The Czech Republic - network connections

Public provided by Large Infrastructure CESNET. The presented links are coloured by average load. Click on the line to obtain detailed statistics.

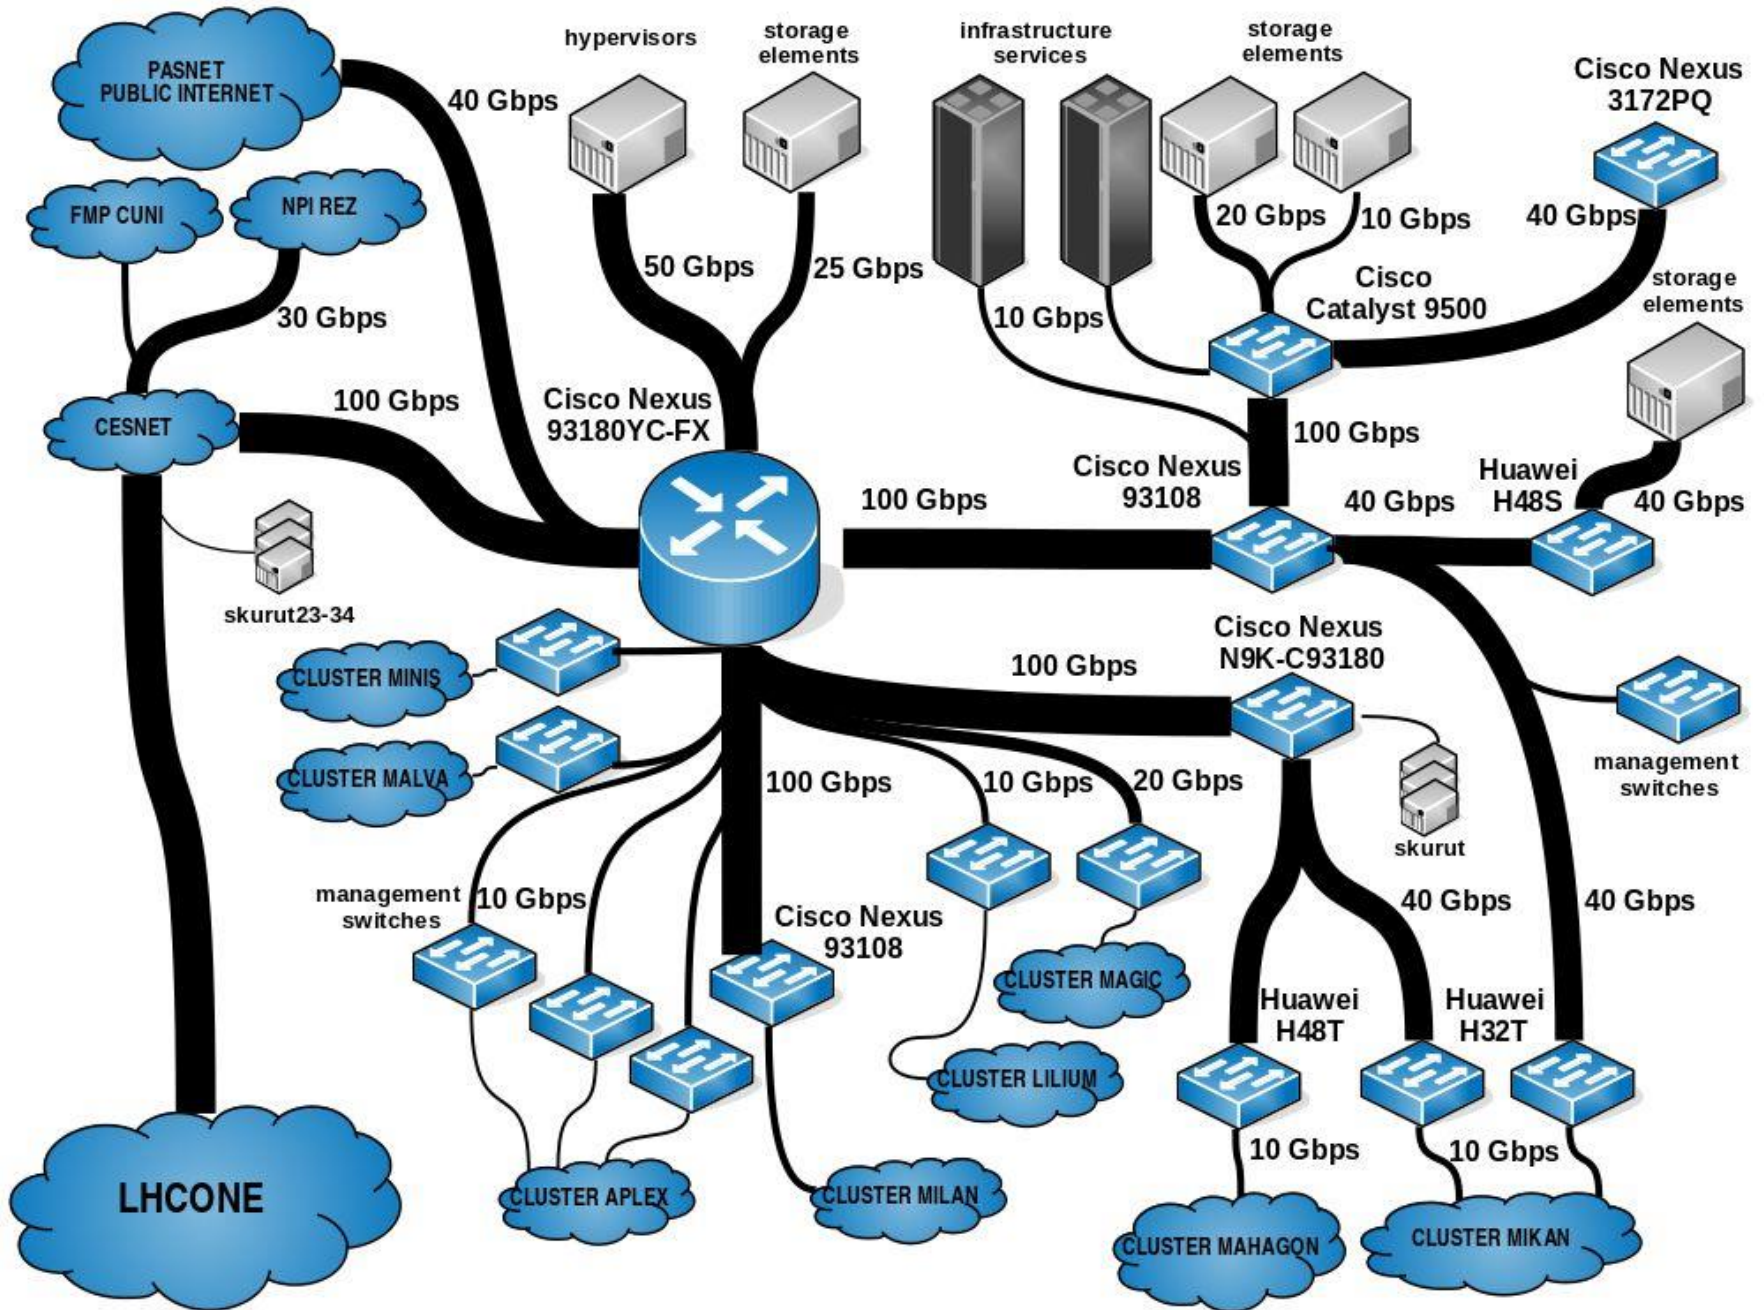


## Data traffic on the link FZU → NPI



**Dedicated private connection provided by CESNET**  
**Traffic during the period 3/2022 – 9/2022**  
**Output traffic max 41.6 Gb/s**





# Delivered vs Required resources 2019 - 2022

<u>2019</u>	ALICE required	Delivered
CPU/year (HS06 hours)	152,928,000	227,570,650
Disk (PB)	1.6	2.0

<u>2020</u>	ALICE required	Delivered
CPU/year (HS06 hours)	157,248,000	294,191,948
Disk (PB)	2.04	1.9

<u>2021</u>	ALICE required	Delivered
CPU/year (HS06 hours)	154,656,000	321,702,339
Disk (PB)	1.9	3.09

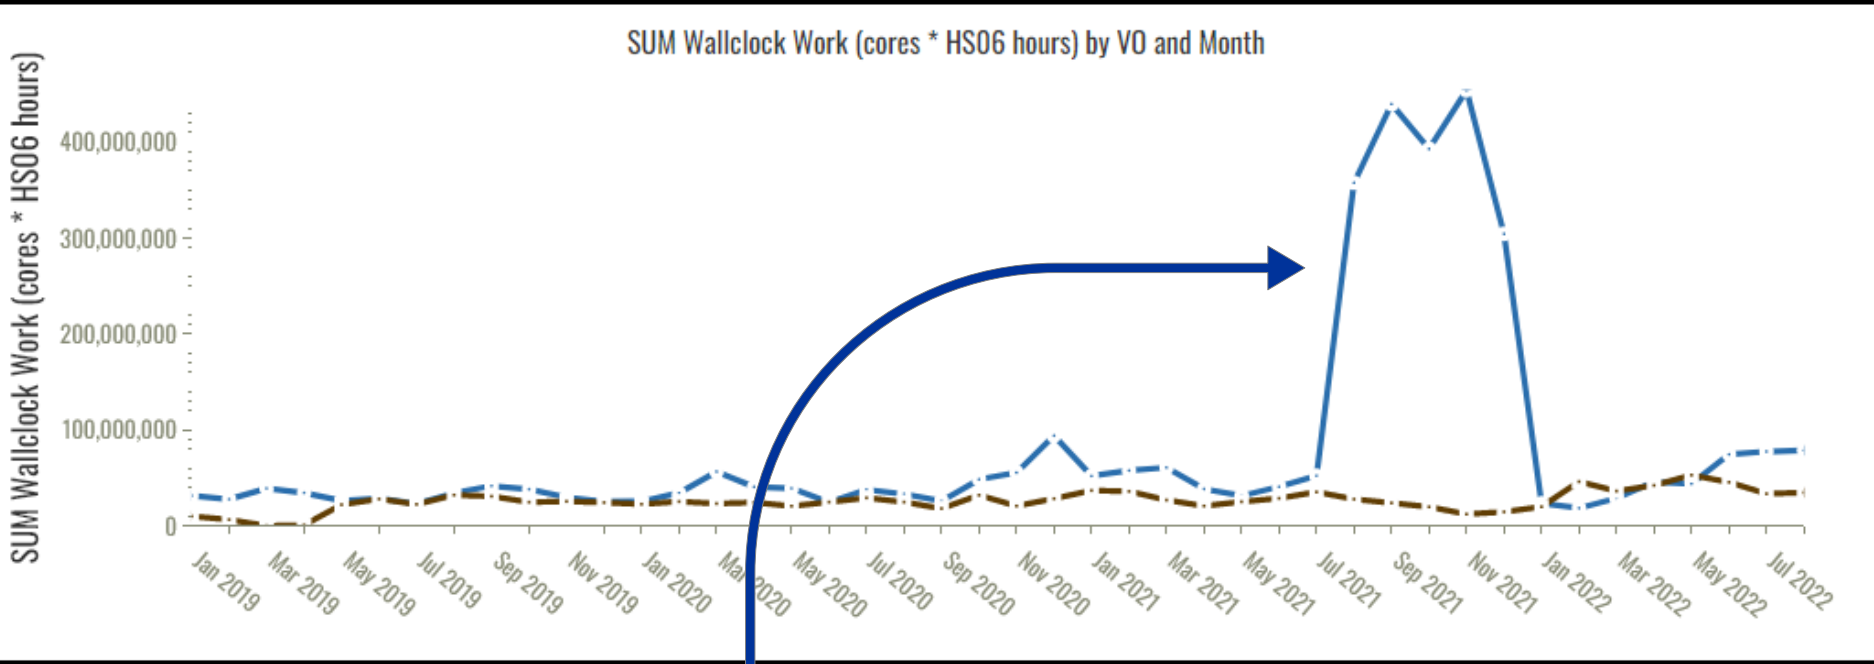
<u>2022</u>	ALICE required	Delivered
CPU/8 months (HS06 hours)	124,128,000	311,013,868
Disk (PB)	2.2	3.33



## Delivered CPU power in 1/2019 – 8/2022

From the EGI Accounting Portal

ALICE required 588,960,000 / delivered 1,140,693,975 HS06 hours




- ALICE delivered in access to what was required mainly due to systematic use of opportunistic resources and due to the site operations almost without outage.
- The peak in the ATLAS contribution is the result of using the new cluster KAROLINA in the HPC center in Ostrava. Available only in 5 months in 2021, ongoing negotiations in 2022.



# Running jobs profile: 1/2019 – 8/2022


## Active jobs in Prague





# CPU delivery share: 1/2021 – 8/2022

## Average running jobs




Czech republic share during 1/2021 – 8/2022 was ~ 5.5%`



# Running jobs profile – local Grafana

Condor: slots usage




Running jobs profile for the period 1/2021 – 6/2022, by the local monitoring at the central site in Prague. Main CPU consumers are ALICE and ATLAS.



# Traffic on ALICE::Prague::SE: 1/2019 – 8/2022



Network traffic on ALICE::Prague::SE

Traffic on Prague SE : 1/2019 – 8/2022  
Total data download: 146.9 PB  
> 4.3 GByte/s in peaks



- lids.ujf.cas.cz
- lids4.ujf.cas.cz
- lids5.ujf.cas.cz
- lids6.disk.ujf.cas.cz
- lids6.ujf.cas.cz
- lids7.disk.ujf.cas.cz
- lids7.ujf.cas.cz
- lids8.disk.ujf.cas.cz
- lids8.ujf.cas.cz
- lids9.disk.ujf.cas.cz
- xrdhead.farm.particle.cz
- xrdtest.farm.particle.cz
- xrootd4.farm.particle.cz
- xrootd5.farm.particle.cz
- xrootd6.farm.particle.cz

# Use of the resources of the HPC center in Ostrava, by ATLAS



## Contributions from the external resources to the praguecg2 ATLAS operations

In 2021: HPC: 84%, Grid:12%, BOINC:4%.

In 2020: HPC: 47%, Grid:52%, BONIC:1%.

# Summary and Outlook

In the period 2019 – 2022, the Czech Tier2 site delivered to ALICE resources in access to what was required, with the exception of disk space in 2020.

The CPU share officially allocated for ALICE would not do the job if there were not the opportunistic resources available. We do our CPU pledges carefully and keep them lower just in case the opportunistic resources were not there. In addition, an old cluster providing ~1500 slots will be de-commissioned by the end of 2022. Because we expect the energy prices go up by factor 5 starting in January 2023, probably more older clusters will be switched off.

There were no substantial problems with delivering the disk capacity. Currently we offer 3.3 PB, but pledged for the 2023 less. The problem is that the server xrootd4.farm.particle.cz which has been in operation almost 9 years will most likely be de-commissioned by the end of 2022.

There is an issue of using the resources of the HPC center in Ostrava. Until now ALICE was not able to use them, but the new cluster KAROLINA is suitable. It was available for prague1cg2 in 2021, when after it was put in operation there were plenty of opportunistic resources. Since then we could not use the cluster until some issues would be resolved.

In 2023, we currently do not have any promises of money for new hardware. We only can expand the disk capacity at NPI adding disks to one storage box which is not yet fully occupied. But of course we will do our best to keep up the reliability and performance level of the services and deliver the maximum we can.



*This work was supported by a project OP RDE CERN Computing (CZ.02.1.01/0.0/0.0/1 6013/0001404) from EU funds and MEYS.*



EUROPEAN UNION  
European Structural and Investment Funds  
Operational Programme Research,  
Development and Education



MINISTRY OF EDUCATION,  
YOUTH AND SPORTS



***Backups***



**Server xrootd5 (580 TB)**



# Czech Republic resources delivery

***MANDATORY ACCORDING TO THE ALICE CONSTITUTION: 1.96% of the total required T2 resources***

## Resources delivery in 2016:

### CPU (kHepSpec06)

ALICE requirement:	7,2
REBUS info: pledged	5,0
REBUS info: delivered	8,85

## Resources delivery in 2017:

### CPU (kHepSpec06)

ALICE requirement:	12,22
REBUS info: pledged	5,0
REBUS info: delivered	9,4

## Resources delivery in 2016:

### Disk (PB):

ALICE requirement:	0,92
REBUS info: pledged	1,4
REBUS info: delivered	1,79

## Resources delivery in 2017:

### Disk (PB):

ALICE requirement:	1,12
REBUS info: pledged	1,4
REBUS info: delivered	1,79

**For the ALICE Czech group easier to scale up the storage capacity than to pile up CPUs.**

***ALICE requests in 2016 of computing and storage resources were satisfied, though the CPU pledges were low.:***

- by stable operations with almost no outage***
- by using the servers out of warranty***
- by extensive use of other projects resources***

