EuroHyPerCon

The EuroHyPerCon Study: Towards hyperconnectivity for HPC resources

Fotis Karayannis, Innov-Acts, Vassilis Merekoulias, Enomix

JENA Initiative Working Group 1 on HPC/HTC

27.05.2024

EuroHyPerCon fact sheet

- **Title:** Study for **hyper-connectivity** for HPC resources
- **Funding:** EuroHPC JU (LC-02450379)
- **Runtime:** 5 October 2023 – 4 July 2024 (9 months)
  - *May be extended*
- **Partners**
  - Innov-Acts (SME/e-Infrastructures)
  - HLRS (Major HPC centre)
  - Enomix (SME/connectivity + policies)
- **Website:** https://eurohypercon.eu/

EuroHyPerCon Study scope

- **Objective:** EU HPC hyper-connectivity service specification, laying out an implementation roadmap for a secure, federated, and hyper-connected European HPC and data infrastructure
- **Focus:** Requirements analysis & network/services design

Comprehensive Needs and Services Analysis	Forward-Looking Solutions
• Engage with communities • Covering various facets such as traffic, capacity, availability, network architectures, security/privacy, and the evolution of technology	• Aim to accommodate new usages related to scientific instruments and AI, with progressive and flexible solutions to adapt to evolving data traffic needs and changing use cases

Tender proposed approach for EU Hyper-Connectivity

EuroHyPerCon

As outlined in the tender specifications:

- ***Leveraging GÉANT & NRENs' Networks***
  - *Leveraging GÉANT and National Research and Education Networks (NRENs) for HPC hyper-connectivity solutions*
- ***Complementary Connectivity***
  - *Align with ongoing European activities, like the GN5-FPA, to address untargeted HPC-specific needs without redundancy*
- ***Federation Interoperability***
  - *Ensure compatibility and interoperability for future HPC infrastructure federation, considering ties to EU initiatives (e.g., Cloud Federation, DestinE, Human Brain Project, EOSC, European Common Data Spaces)*
- ***Collaborative Study Approach:***
  - *Conduct the study closely with EuroHPC hosting sites, HPC stakeholders, and connectivity players (GÉANT/NRENs) for comprehensive insights and seamless coordination*

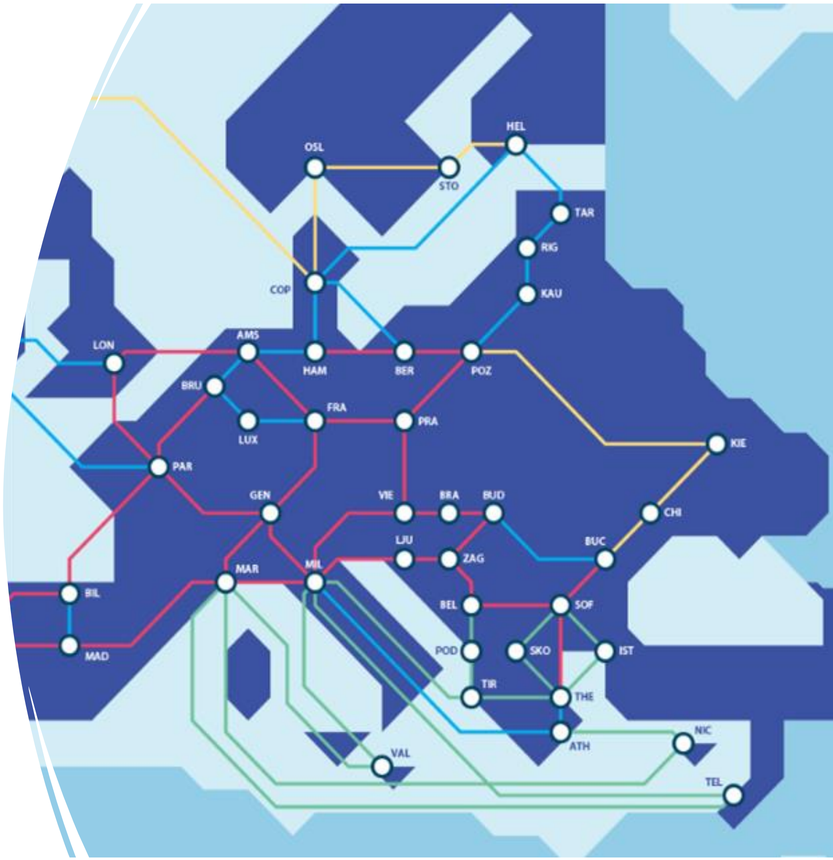

Study methodology

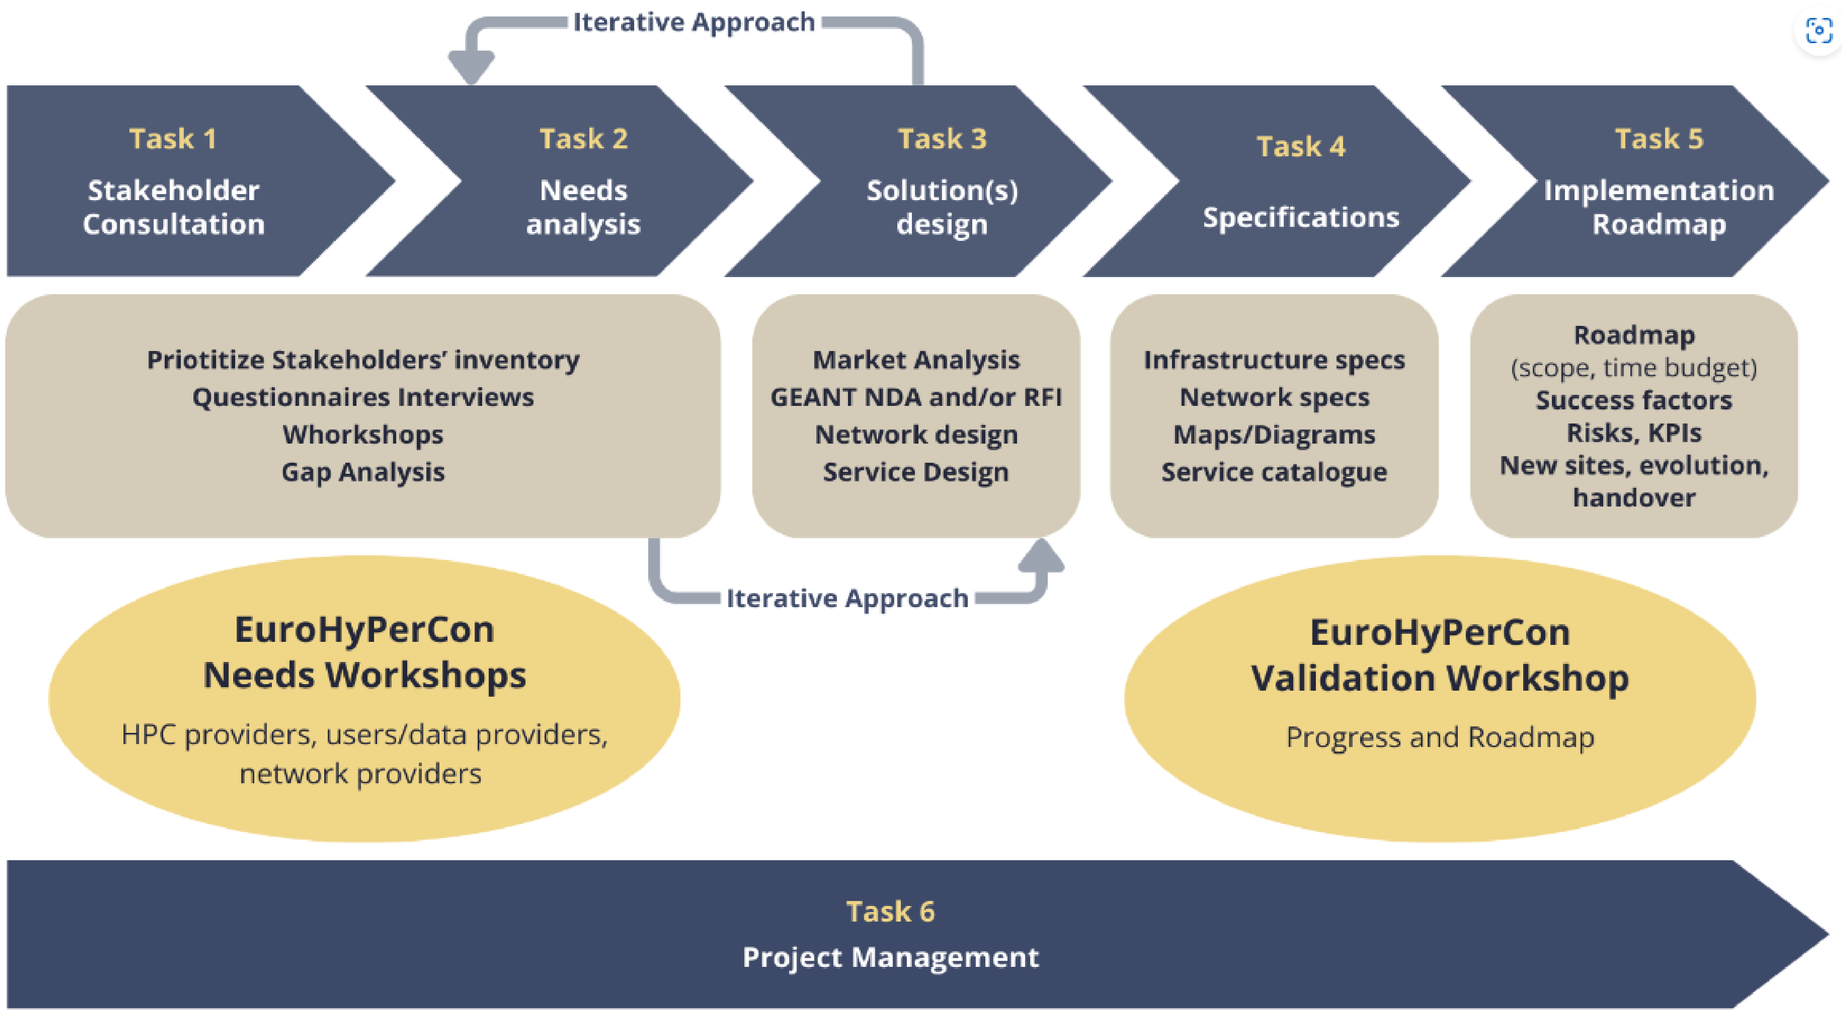

Stakeholders Identification

EuroHyPerCon

HPC Providers	HPC Users	Network Providers	Data Providers and AI Users - Other Stakeholders
• **EuroHPC** Hosting Sites • **Other EU / National** HPC systems	• **Thematic users** of the HPC systems • Big users (e.g., DestinE (ECMWF, EUMETSAT, ESA), CERN, etc.) • Other users	• **GÉANT** • **NRENs** and **regional** research networks • Other connectivity providers	• **Data providers** (e.g., ESFRI & Other RIs, EU Data Spaces) • **AI users** • Online Registration Form

Activities performed

- **Methodology being executed as planned**
  - **Workshops** – There was high interest
    - Stakeholder Identification and User Journeys – 30 October 2023
    - Feedback from HPC users and providers – 22 November 2023
    - Feedback from network providers – 27 November 2023
  - **Focus-Groups/Interviews**
    - Focus group with Exascale & Pre-exascale network providers – 18 December 2023
    - Interview with Destination Earth/ECMWF – 18 December 2023
    - Meeting with EuroHPC JU and GÉANT – 11 January 2024
    - Focus group with Exascale & Pre-Exascale HPC providers – 5 February 2024
    - Interview with Destination Earth/ECMWF-EUMETSAT-ESA – 9 February 2024
    - Interview with Destination Earth/EUMETSAT – 1 March
    - Focus group meeting with AI users – 12 March 2024
  - **Development of EuroHyPerCon stakeholders' database ~ 500 stakeholders** (680 entries)
    - EuroHPC Hosting sites, EU/National HPC Providers **~ 165** (345 systems)
    - HPC Users **(~175)**
    - Data Providers **(~130)**
    - AI stakeholders **(~30)**
    - Geo-location information for some of the stakeholders (GIS-enabled)

Stakeholders population density - HPC providers/users

EuroHyPerCon

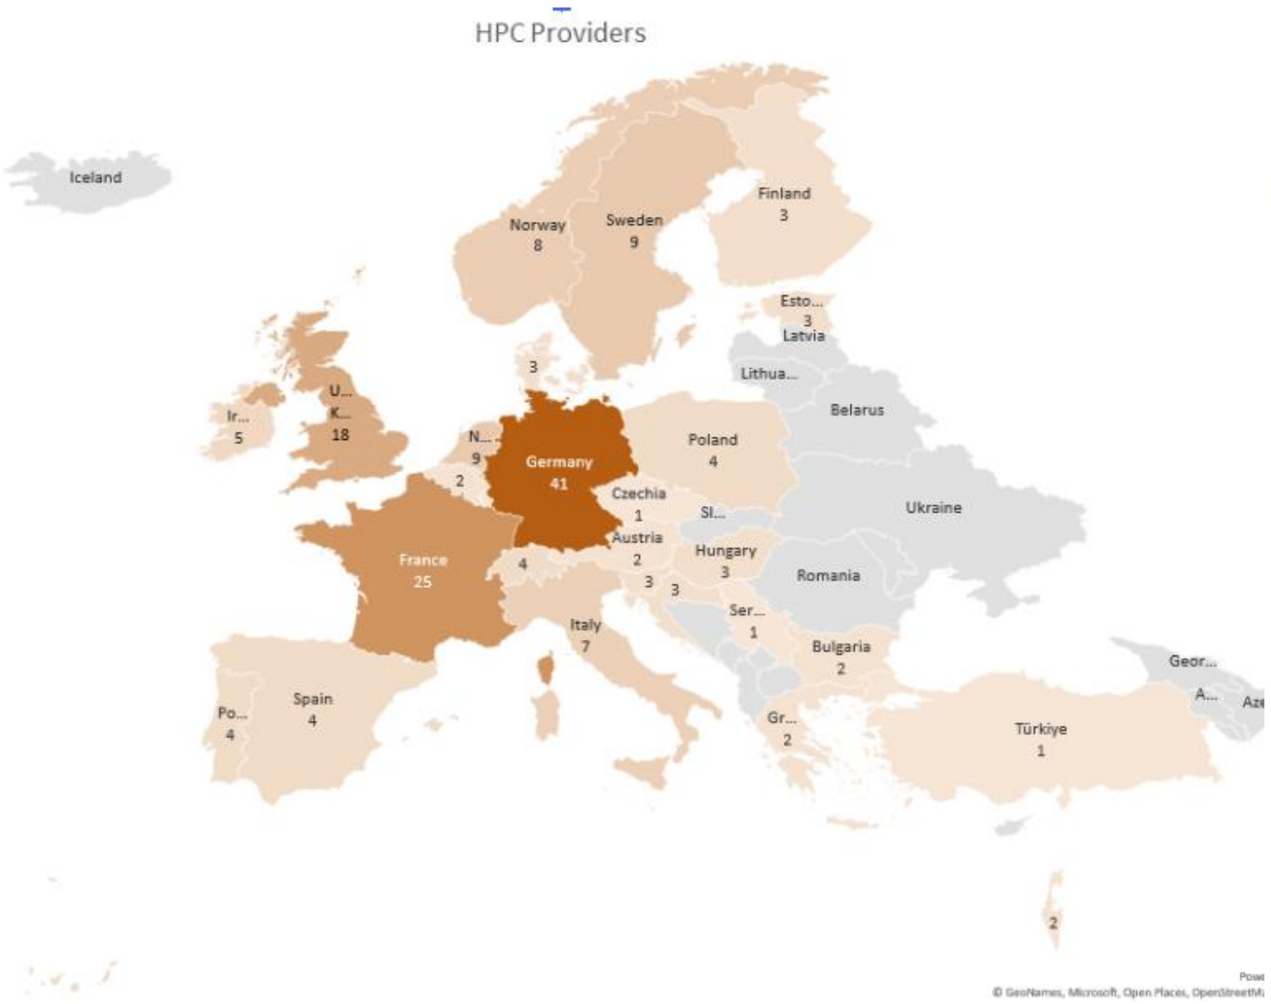

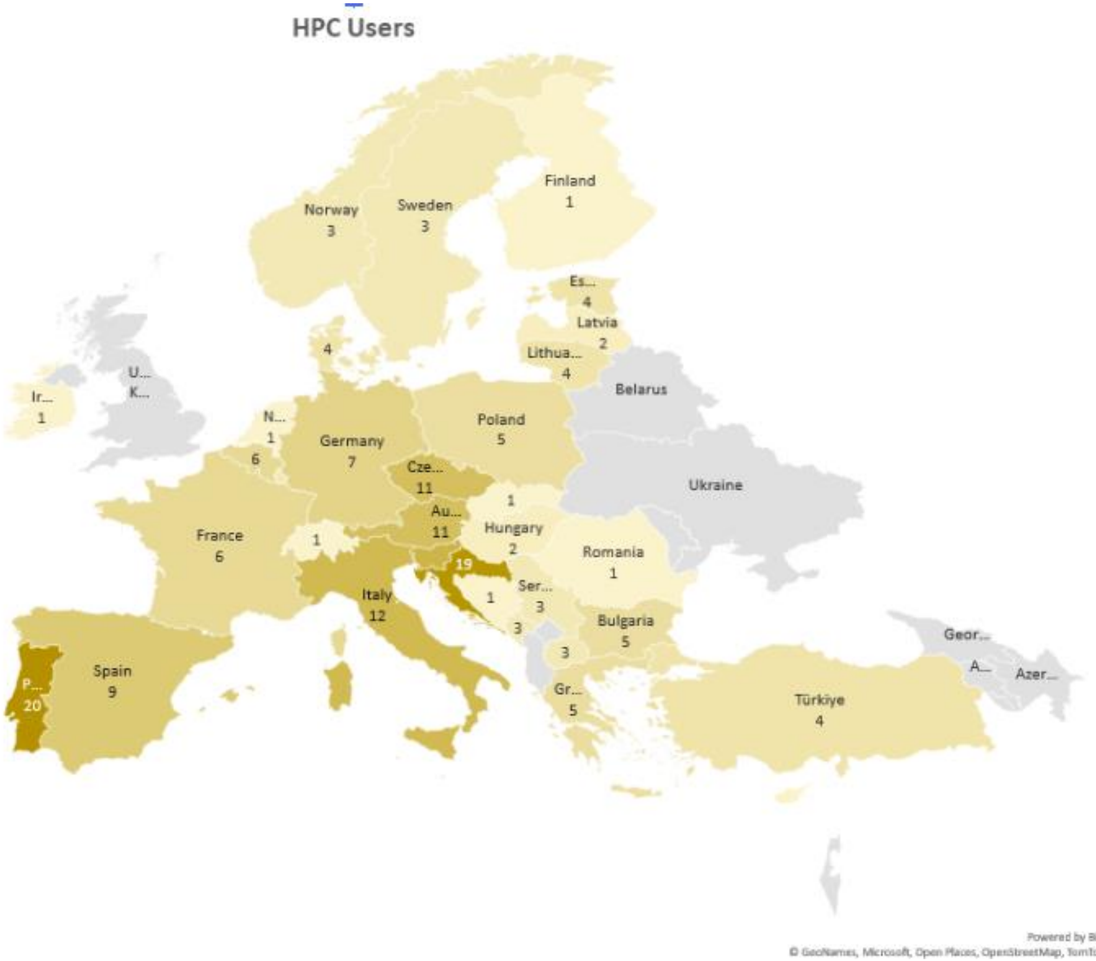

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (1)

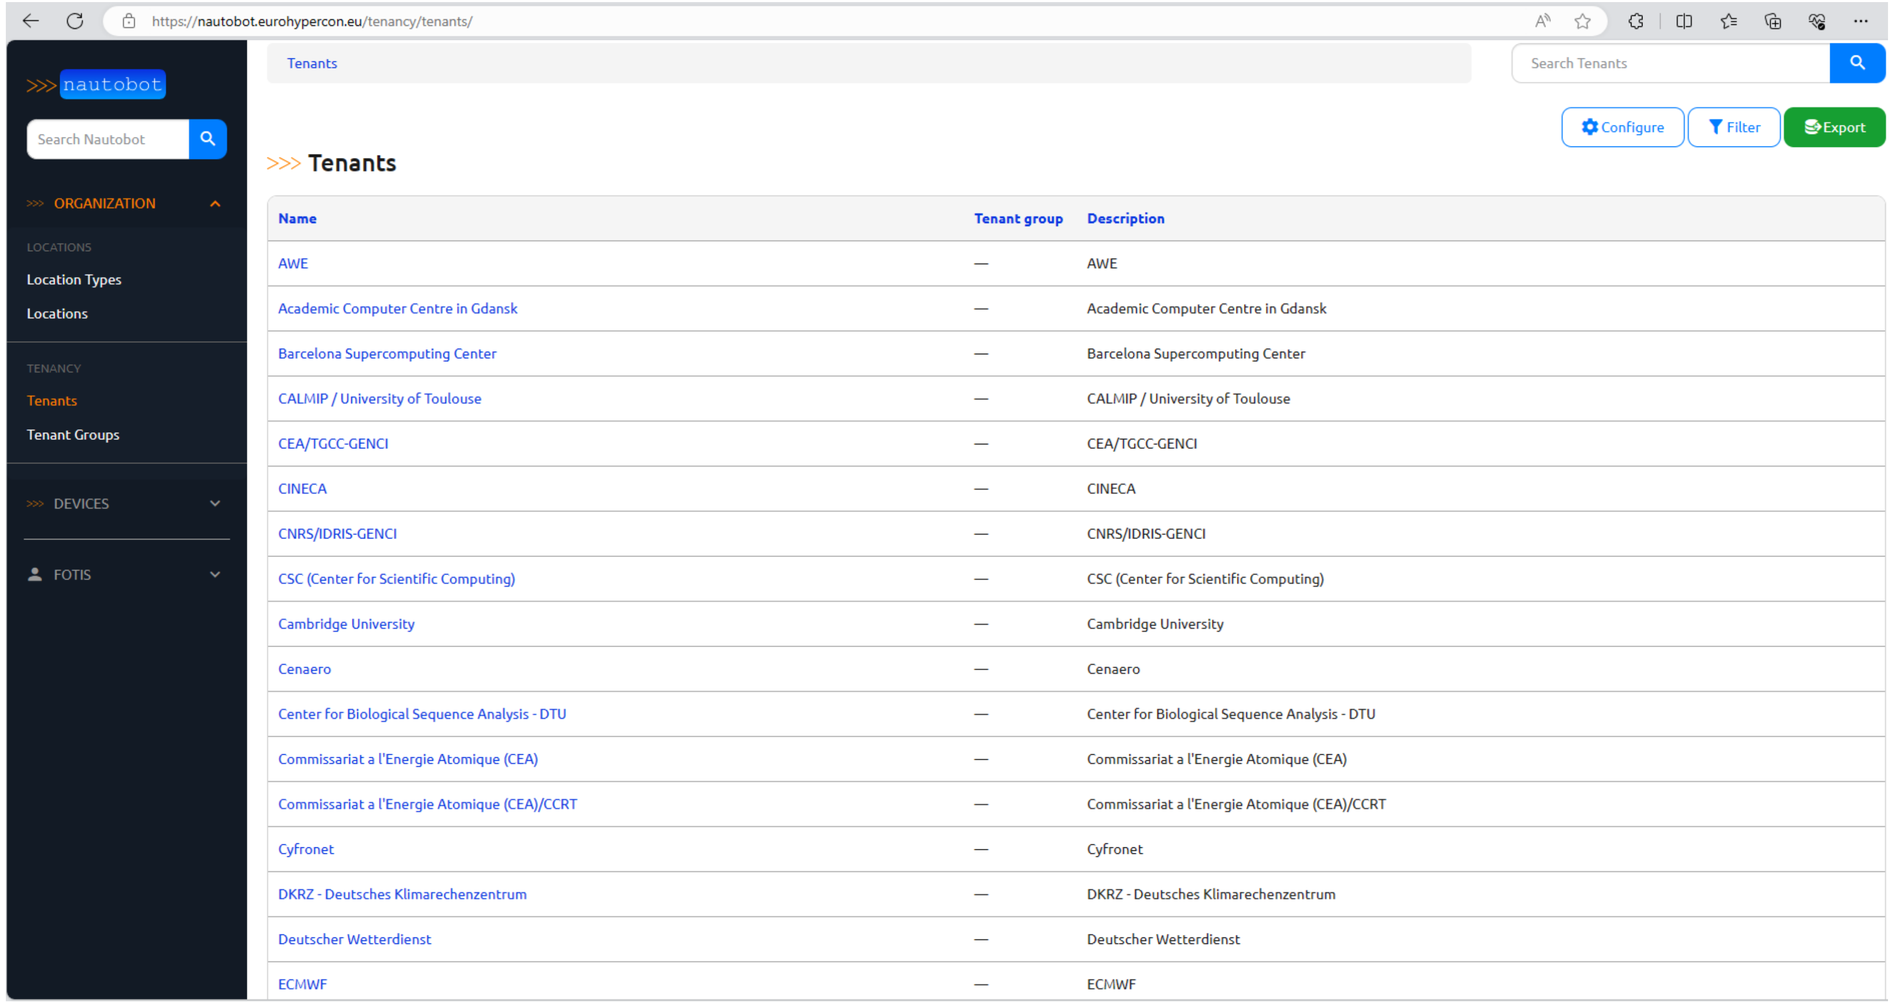

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (2)

EuroHyPerCon

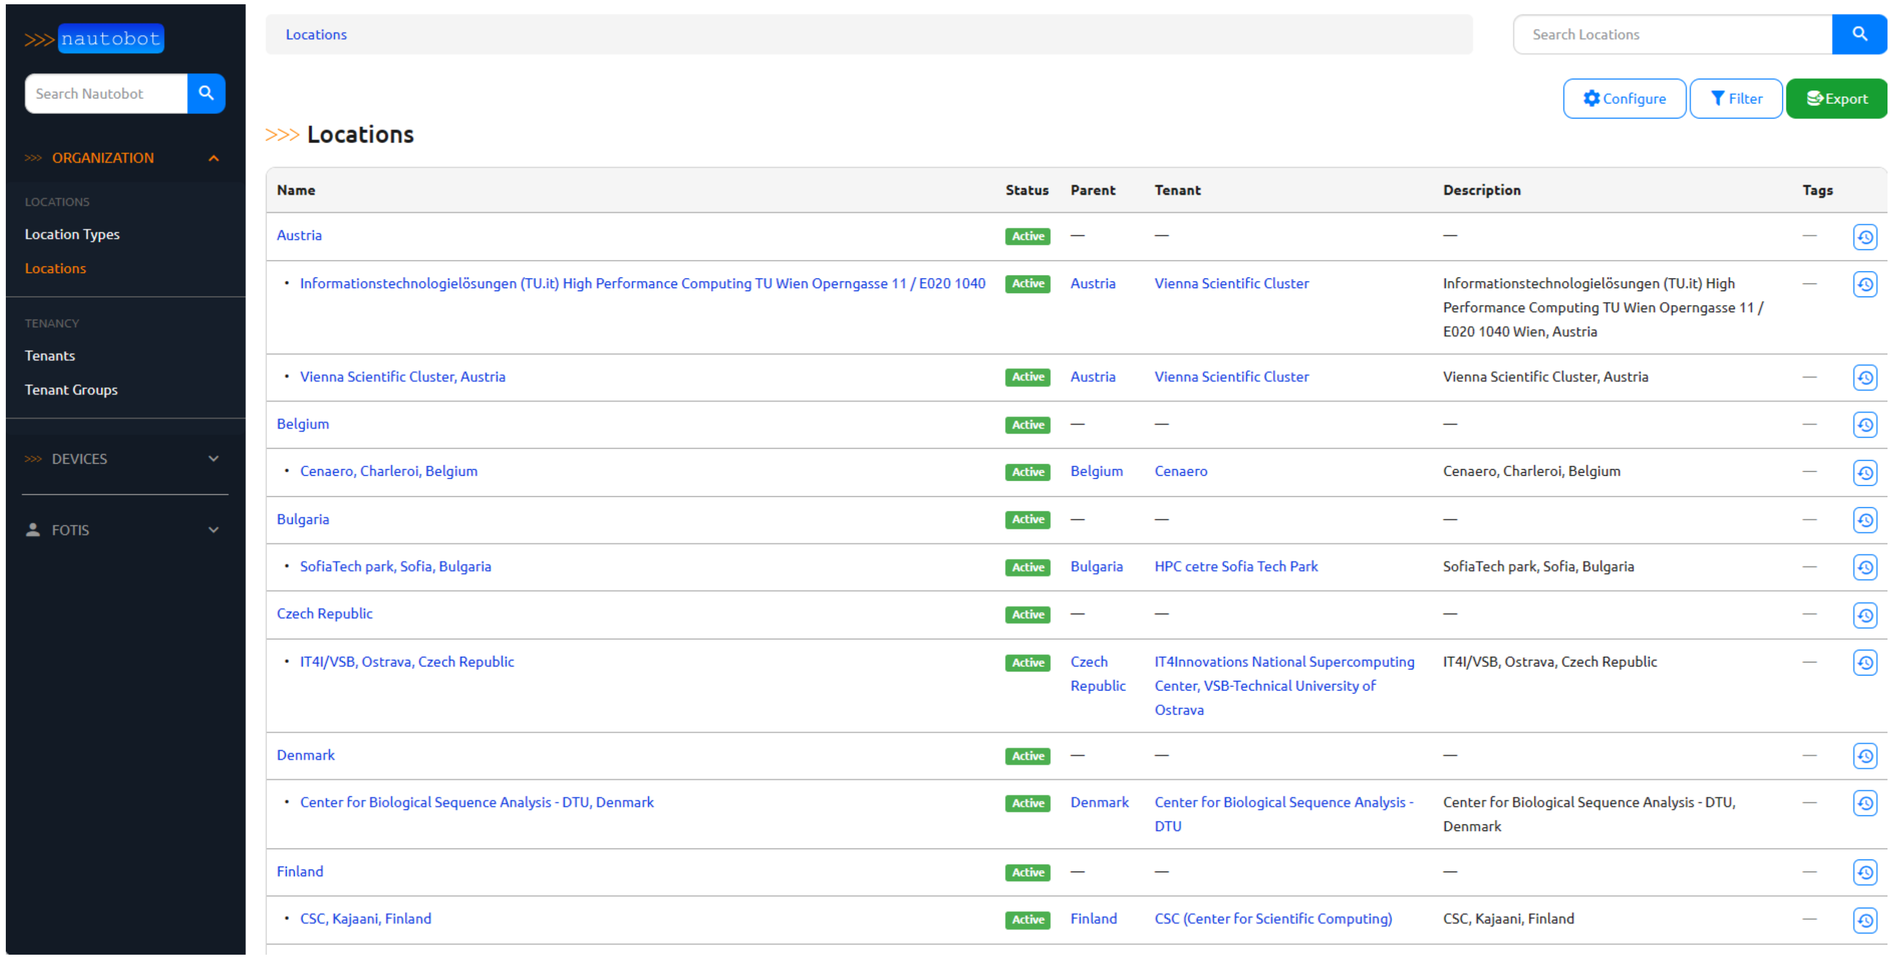

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (3)

EuroHyPerCon

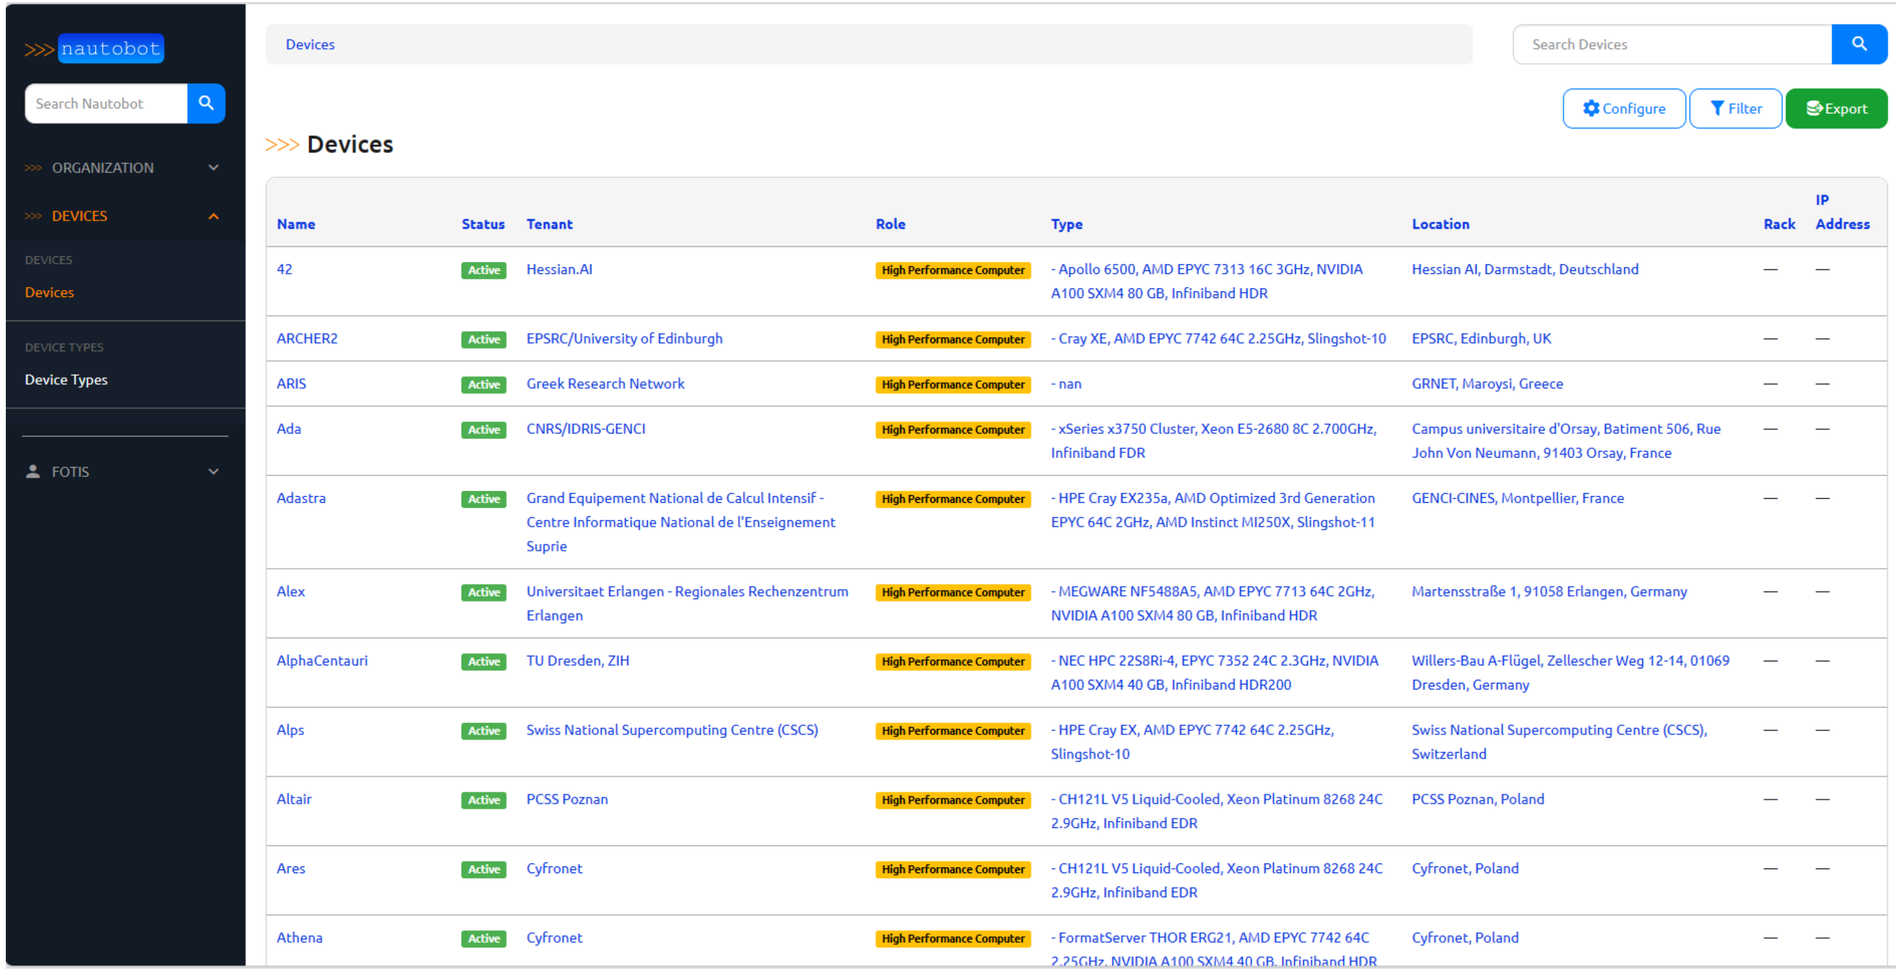

Name	Status	Tenant	Role	Type	Location	Rack	IP Address
42	Active	Hessian.AI	High Performance Computer	- Apollo 6500, AMD EPYC 7313 16C 3GHz, NVIDIA A100 SXM4 80 GB, Infiniband HDR	Hessian AI, Darmstadt, Deutschland	—	—
ARCHER2	Active	EPSRC/University of Edinburgh	High Performance Computer	- Cray XE, AMD EPYC 7742 64C 2.25GHz, Slingshot-10	EPSRC, Edinburgh, UK	—	—
ARIS	Active	Greek Research Network	High Performance Computer	- nan	GRNET, Maroysi, Greece	—	—
Ada	Active	CNRS/IDRIS-GENCI	High Performance Computer	- xSeries x3750 Cluster, Xeon E5-2680 8C 2.700GHz, Infiniband FDR	Campus universitaire d'Orsay, Batiment 506, Rue John Von Neumann, 91403 Orsay, France	—	—
Adastra	Active	Grand Equipement National de Calcul Intensif - Centre Informatique National de l'Enseignement Suprie	High Performance Computer	- HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	GENCI-CINES, Montpellier, France	—	—
Alex	Active	Universitaet Erlangen - Regionales Rechenzentrum Erlangen	High Performance Computer	- MEGWARE NF5488A5, AMD EPYC 7713 64C 2GHz, NVIDIA A100 SXM4 80 GB, Infiniband HDR	Martensstraße 1, 91058 Erlangen, Germany	—	—
AlphaCentauri	Active	TU Dresden, ZIH	High Performance Computer	- NEC HPC 22S8Ri-4, EPYC 7352 24C 2.3GHz, NVIDIA A100 SXM4 40 GB, Infiniband HDR200	Willers-Bau A-Flügel, Zellescher Weg 12-14, 01069 Dresden, Germany	—	—
Alps	Active	Swiss National Supercomputing Centre (CSCS)	High Performance Computer	- HPE Cray EX, AMD EPYC 7742 64C 2.25GHz, Slingshot-10	Swiss National Supercomputing Centre (CSCS), Switzerland	—	—
Altair	Active	PCSS Poznan	High Performance Computer	- CH121L V5 Liquid-Cooled, Xeon Platinum 8268 24C 2.9GHz, Infiniband EDR	PCSS Poznan, Poland	—	—
Ares	Active	Cyfronet	High Performance Computer	- CH121L V5 Liquid-Cooled, Xeon Platinum 8268 24C 2.9GHz, Infiniband EDR	Cyfronet, Poland	—	—
Athena	Active	Cyfronet	High Performance Computer	- FormatServer THOR ERG21, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 SXM4 40 GB, Infiniband HDR	Cyfronet, Poland	—	—

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (4)

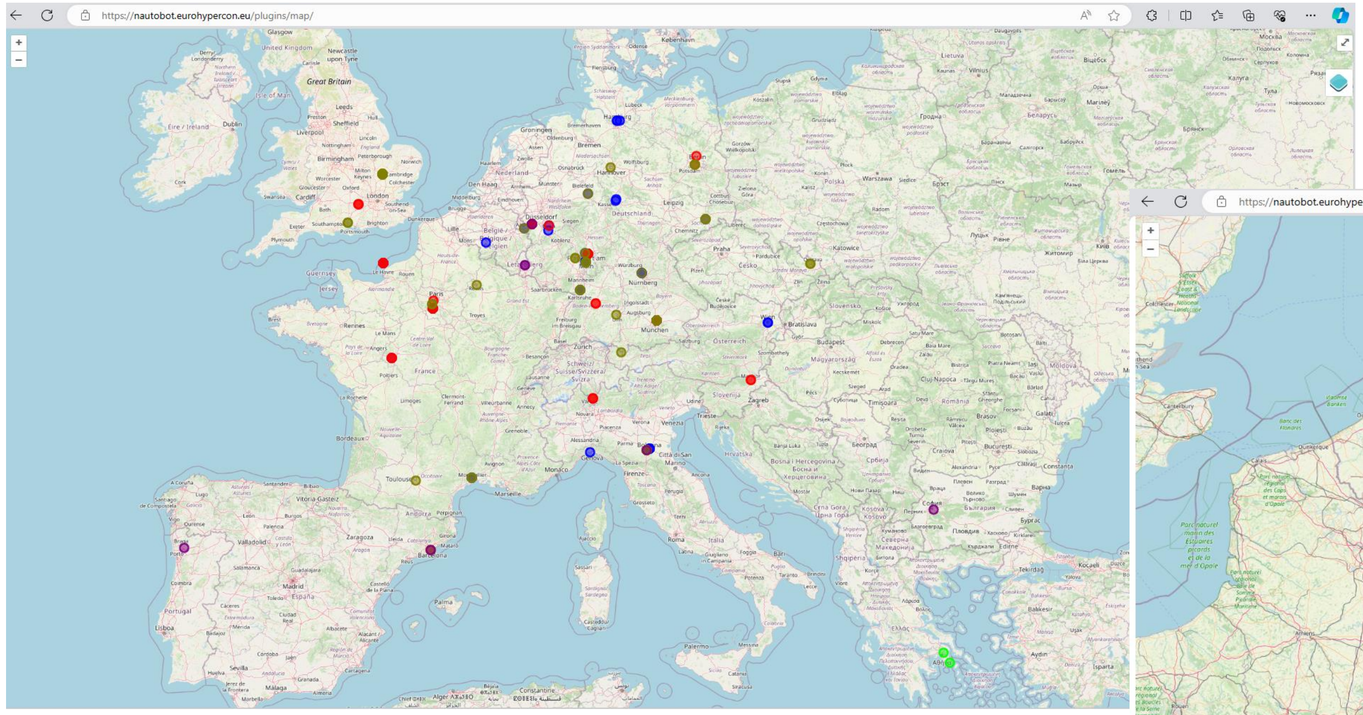

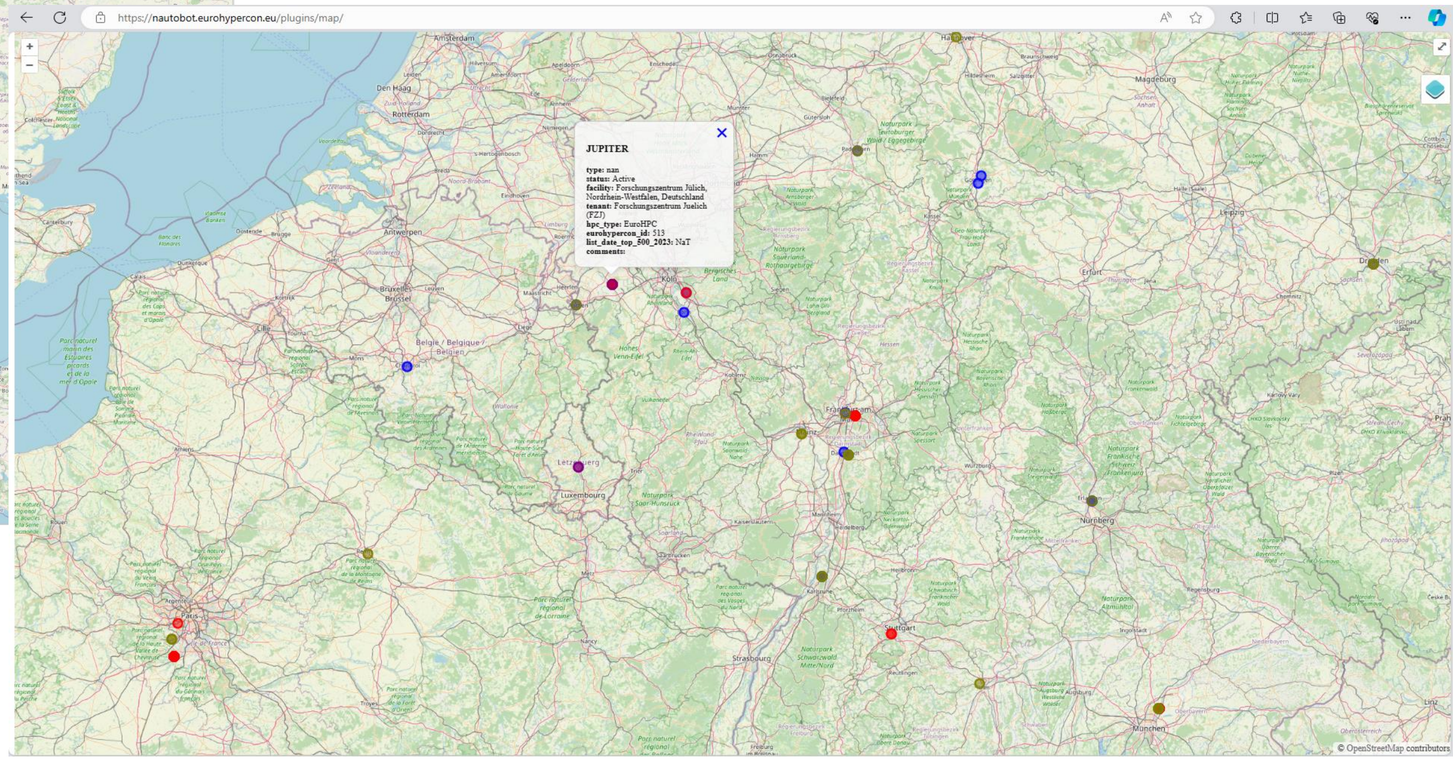

Register as stakeholder: t.ly/e5FE8

EuroHyPerCon questionnaires

- **Questionnaires now closed → 165 full responses (>400 partial, not all useful)**
  - **HPC Users: 111**
  - **HPC Providers: 32**
  - **Network Providers: 22**
- Initial deadline in January (first phase) – More inputs received, e.g. AI stakeholders (iterative approach)
- Questionnaires remained open until the end of March – Now closed!

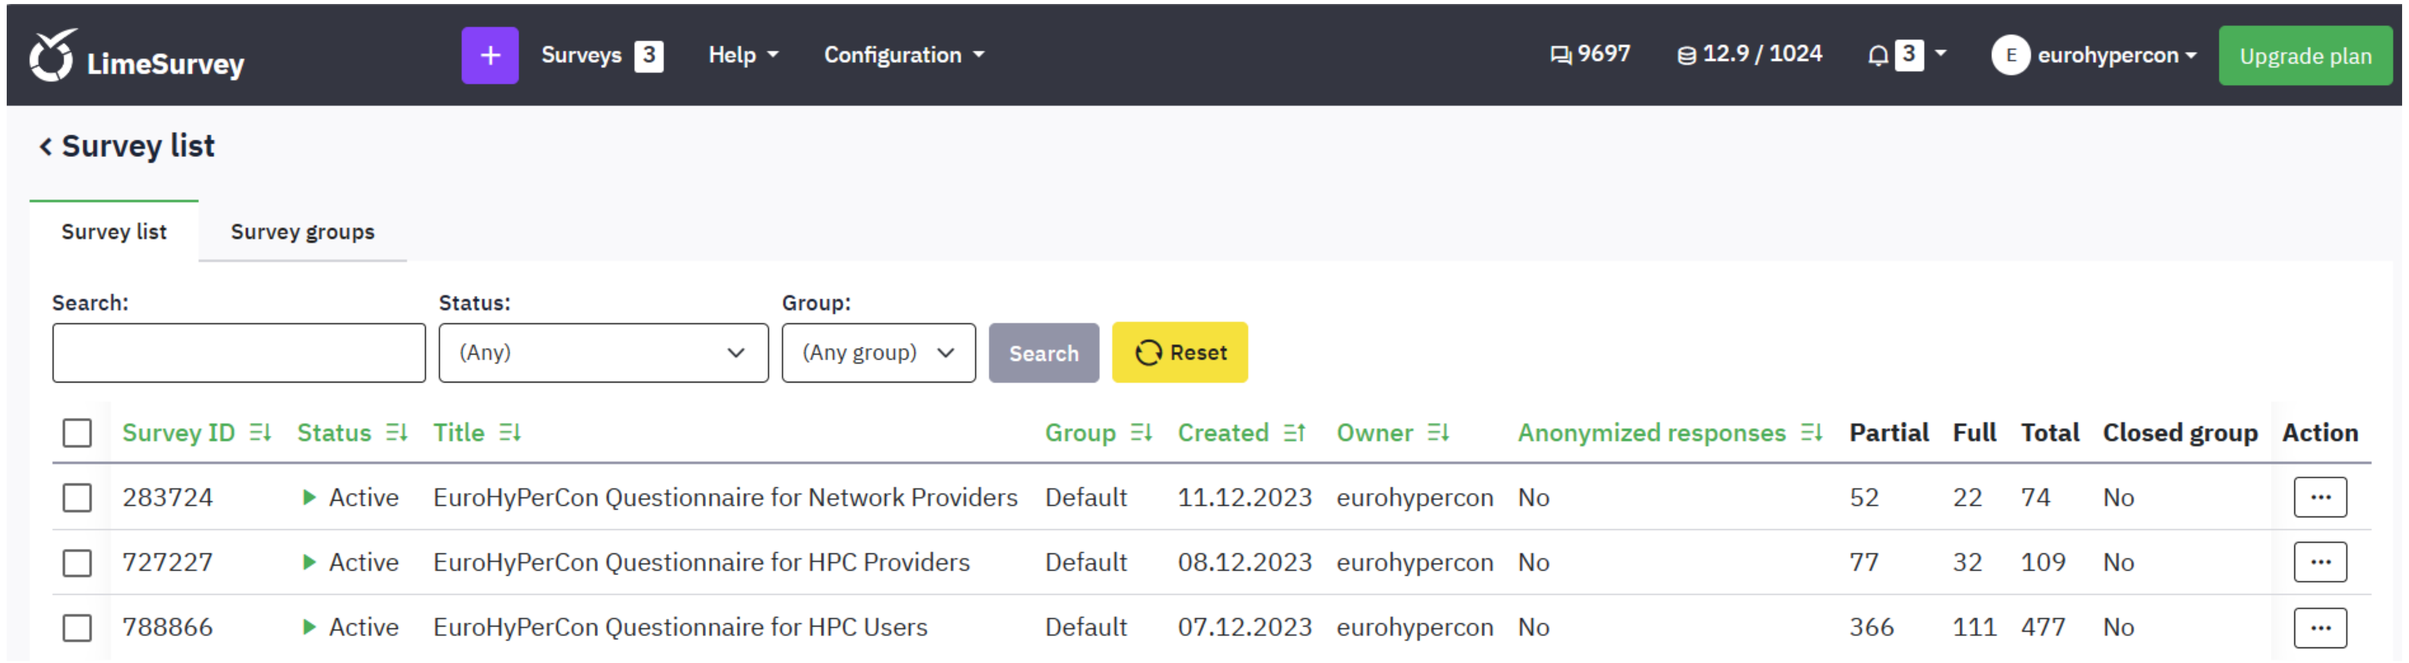

Questionnaires' analysis – Work in progress

- **Feedback from users' questionnaires show national and pan-European needs**
  - Considerable cross-border needs only for operational EuroHPC hosting sites
  - Some active countries have more responses
  - Some countries could have benefited from more answers

- Additional inputs from all other means (workshops, interviews, focus groups, etc.)
  - Inputs are analysed and cross-checked – Final phase of analysis

Users Questionnaires – Some statistics (1/2)

EuroHyPerCon

Type of HPC usage - Now

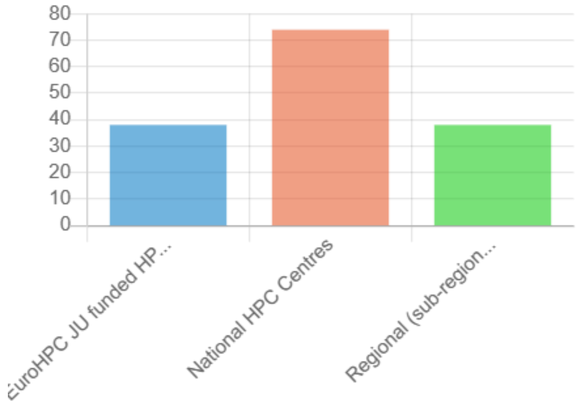

HPC usage timeframe

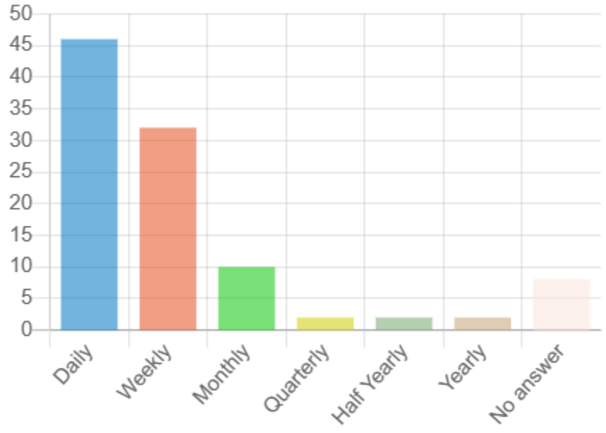

Security requirements?

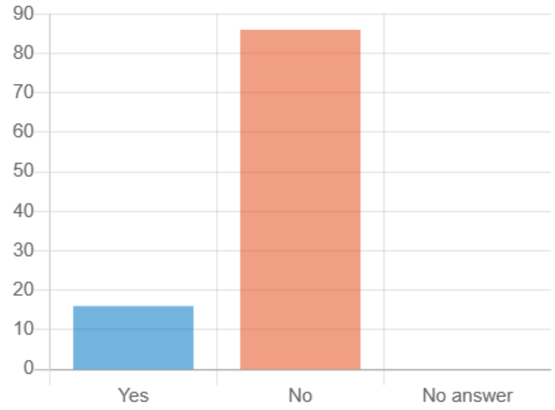

Type of HPC usage - 2030

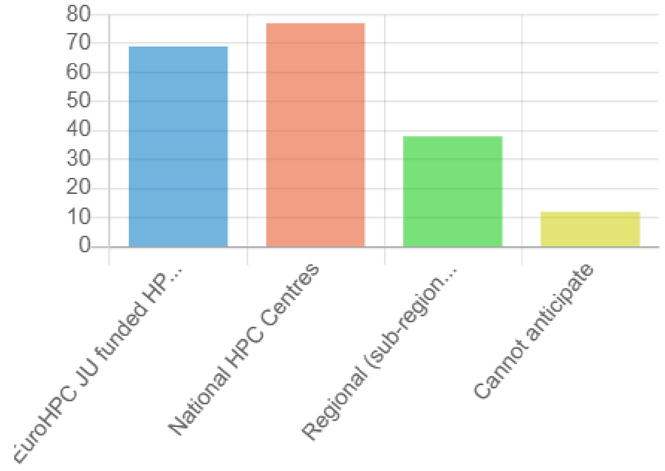

Amount of resources used

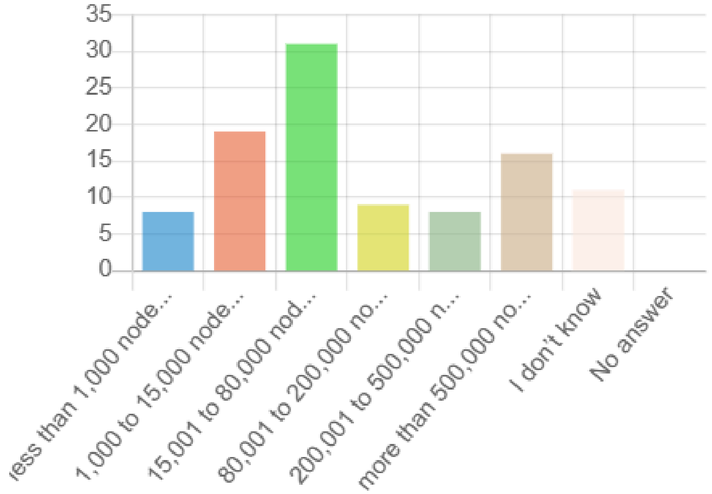

Most data transferred…

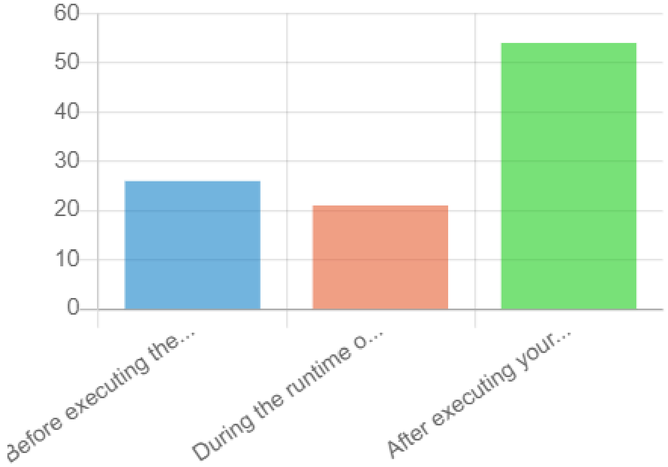

Users Questionnaires – Some statistics (2/2)

Data size

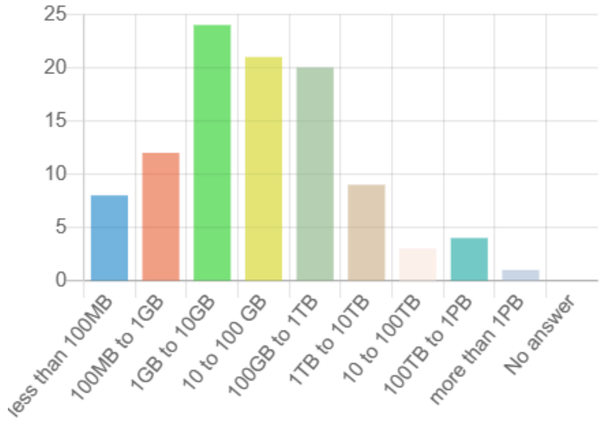

Data transfer times

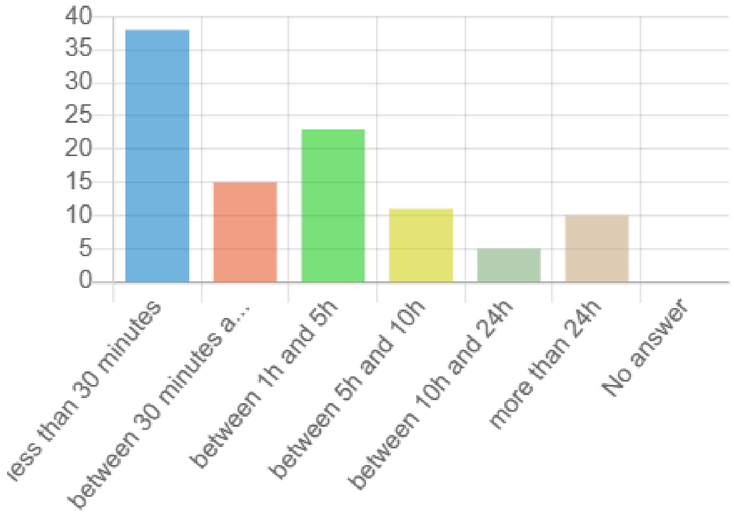

Satisfied with data transfer times?

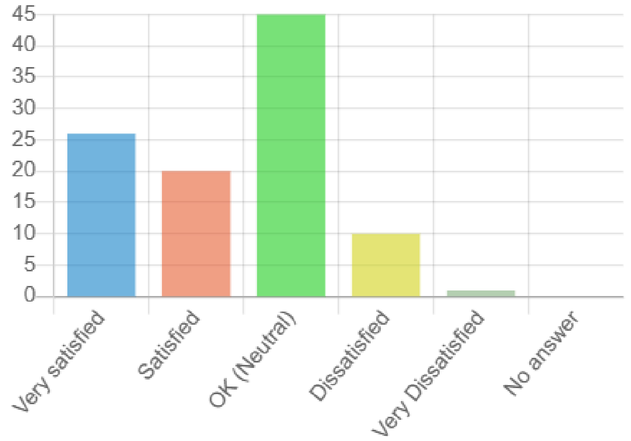

Data size growth (%) 2025-2028-2030+

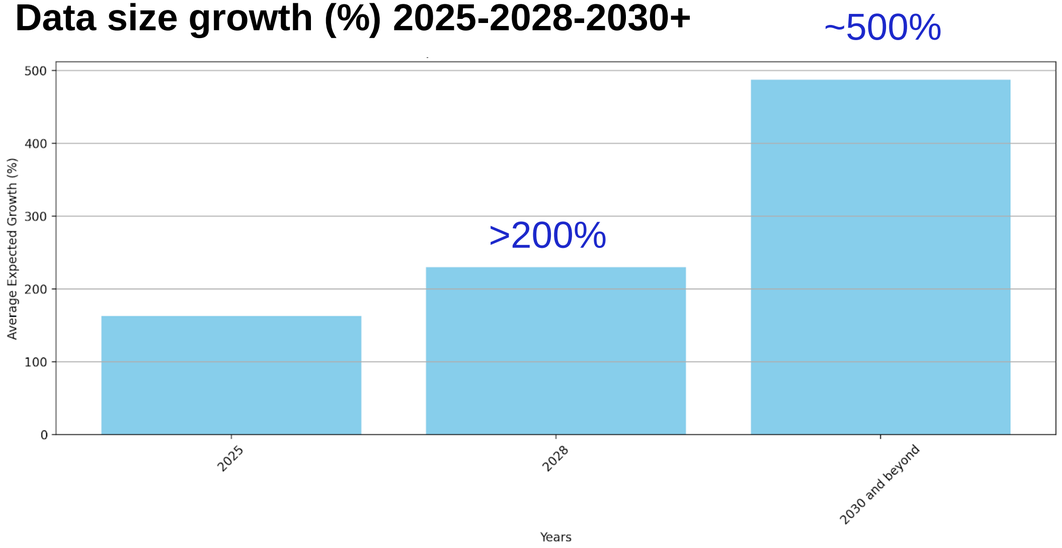

Some of the challenges: Local network, policies limiting bandwidth usage, routing rules, storage capacity limitations at HPC centre, security/firewalls/ssh connection failures

Summary of preliminary findings

EuroHyPerCon

Multiple inputs from several workshops, focus groups, interviews and questionnaires:

- Users are satisfied by services provided by **GÉANT & NRENs**
- The majority of NRENs & GÉANT ready to **upgrade** access and backbone links reaching n x 400 Gbps and then Tbps levels;
  - GÉANT: Soon **400Gbps for backbone/user access**
- Main issues on accessing & uploading/downloading data to/from HPC Providers
  - **Security** related aspects: SSH access may affect network performance
  - Different levels of security/practices across sites → **harmonization** needed
- **Majority of users request national HPC resources / some pan-European**
  - Pan-European requirements can be mostly satisfied by GÉANT
  - **DestinE:** champion user/data provider: Data infrastructure deploying commercial solutions→ A bespoke solution may be required

EuroHyPerCon

Way forward

- **End-to-end** solution for EuroHPC, spanning panEuropean+national segments (to reach HPC sites):
  - **Leverage GÉANT / NRENs** that could meet the vast requirements, currently reaching all countries and also having global reach
    - Adaptation to HPC needs, upgrades when needed, evolution over time
    - **Plain IP service: Class-based access ports** (e.g. Class A 400Gbps to 1Tbps, Class B 200Gbps to 400Gbps, etc.)
    - Over the top services (NOC/user support/helpdesk, transport security, etc.)
  - **Bespoke solutions for big users/data providers** (e.g. DestinE)
  - **Connect external Cloud providers:** Commercial (Amazons)/user-deployed (Nextclouds)
    - **Peerings** with major cloud providers needed in relevant locations with ample capacity to facilitate forward-looking high-throughput exchange of data.
- Aim for **"as a service"** solution; outsource to network providers;
  - EuroHPC JU will need to only oversee; no need for an internal network management team;
- **Integrated connectivity service provision** - From HPC Providers to an EuroHPC Ecosystem!
  - Homogenized services/practices across HPC sites: Common access/methods for upload/download data, VPN/encryption services; This will enhance user experience! → **Input to Federation call project**

Hypothesis validation:
Techno-economic analysis of different solutions

EuroHyPerCon

- **Approach A - GÉANT / NRENs**
  - Upgrade access speeds of ~80-100 HPC systems / ~80-100 Data Providers
  - Need cost info about access upgrades from a good sample of NRENs (East/West, North/South, big/medium/small) and backbone upgrades from GÉANT
    - CAPEX (Upgrade access cards/equipment/links and backbone cards/equipment/links) and OPEX (New personnel needed, extra energy for the new equipment, etc.)
      - For both the Optical (layer 1) and IP/MPLS layers (Layers 2/3)
  - Cost evolution: 2025, 2028, 2030+

  → Proposed approach: **Incremental cost** based on GÉANT and sample of NRENs (compared to 2024 status)

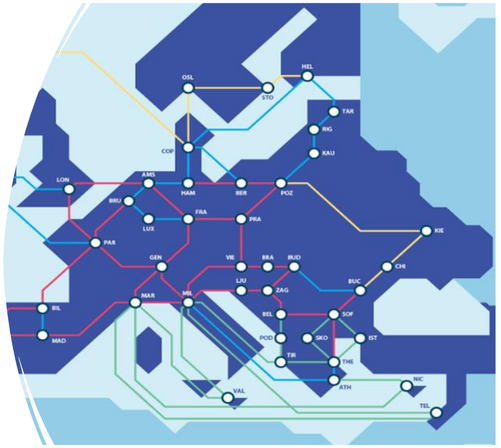

- **Approach B - IP transit service from commercial providers**
  - IP transit interconnecting a set of end points:
    - ~165 HPC systems / ~130 Data Providers / ~60 Internet Exchange Points
  - Based on recent procurement prices (industry benchmarks) – Telegeography.c
  - Inspired by the current DestinE model (which is based on a commercial provider)
  - Cost evolution: 2025, 2028, 2030+
  - → Commercial Service based on IP transit + Local Loops (e.g. DWDM / Metro Ethernet)

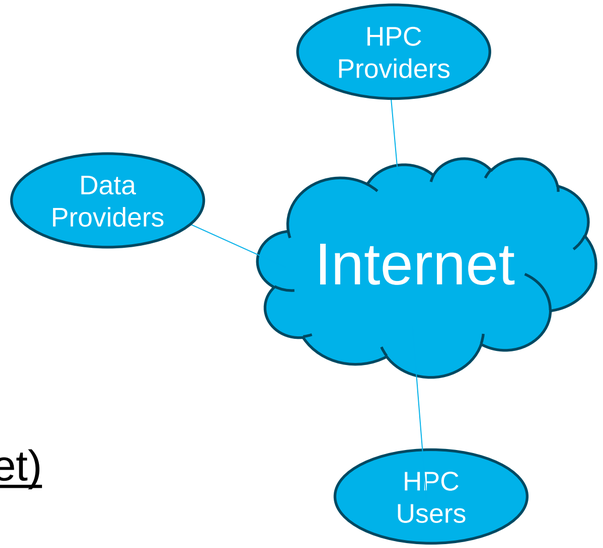

- **Technical analysis for both solutions to run in parallel**

Conclusions and Next steps

EuroHyPerCon

- Stakeholders' identification performed
  - Surveys closed
- Needs analysis submitted for comments from the EuroHPC JU
- Proceeding with Gap analysis (from current solutions, i.e. GEANT/NRENs)
- Then alternative solutions design
  - Techno-economic analysis ongoing
- Proceed with technical assessment: fitness for purpose, performance, support for innovation, security, impact..
- Validation workshops
  - One online with NRENs on Wednesday

Questions?

EuroHyPerCon

More info

- *https://eurohypercon.eu (surveys, stakeholder registration form, summary of workshops)*
- *info at eurohypercon . eu*
- *surveys at eurohypercon . eu*

Thank you!