

ANTWERP

Interconnecting EuroHPC Supercomputers for Scientific and Industrial Advancement

The EuroHyPerCon study



UNLEASHING THE
POWER OF EUROPEAN
HPC AND QUANTUM
COMPUTING

Fotis Karayannis, Innov-Acts, EuroHyPerCon study coordinator



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)
- **Partners**
 - Innov-Acts
 - HLRS
 - Enomix
- **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

- Engage with communities
- Covering various facets such as traffic, capacity, availability, network architectures, security/privacy, and the evolution of technology

Forward-Looking Solutions

- 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



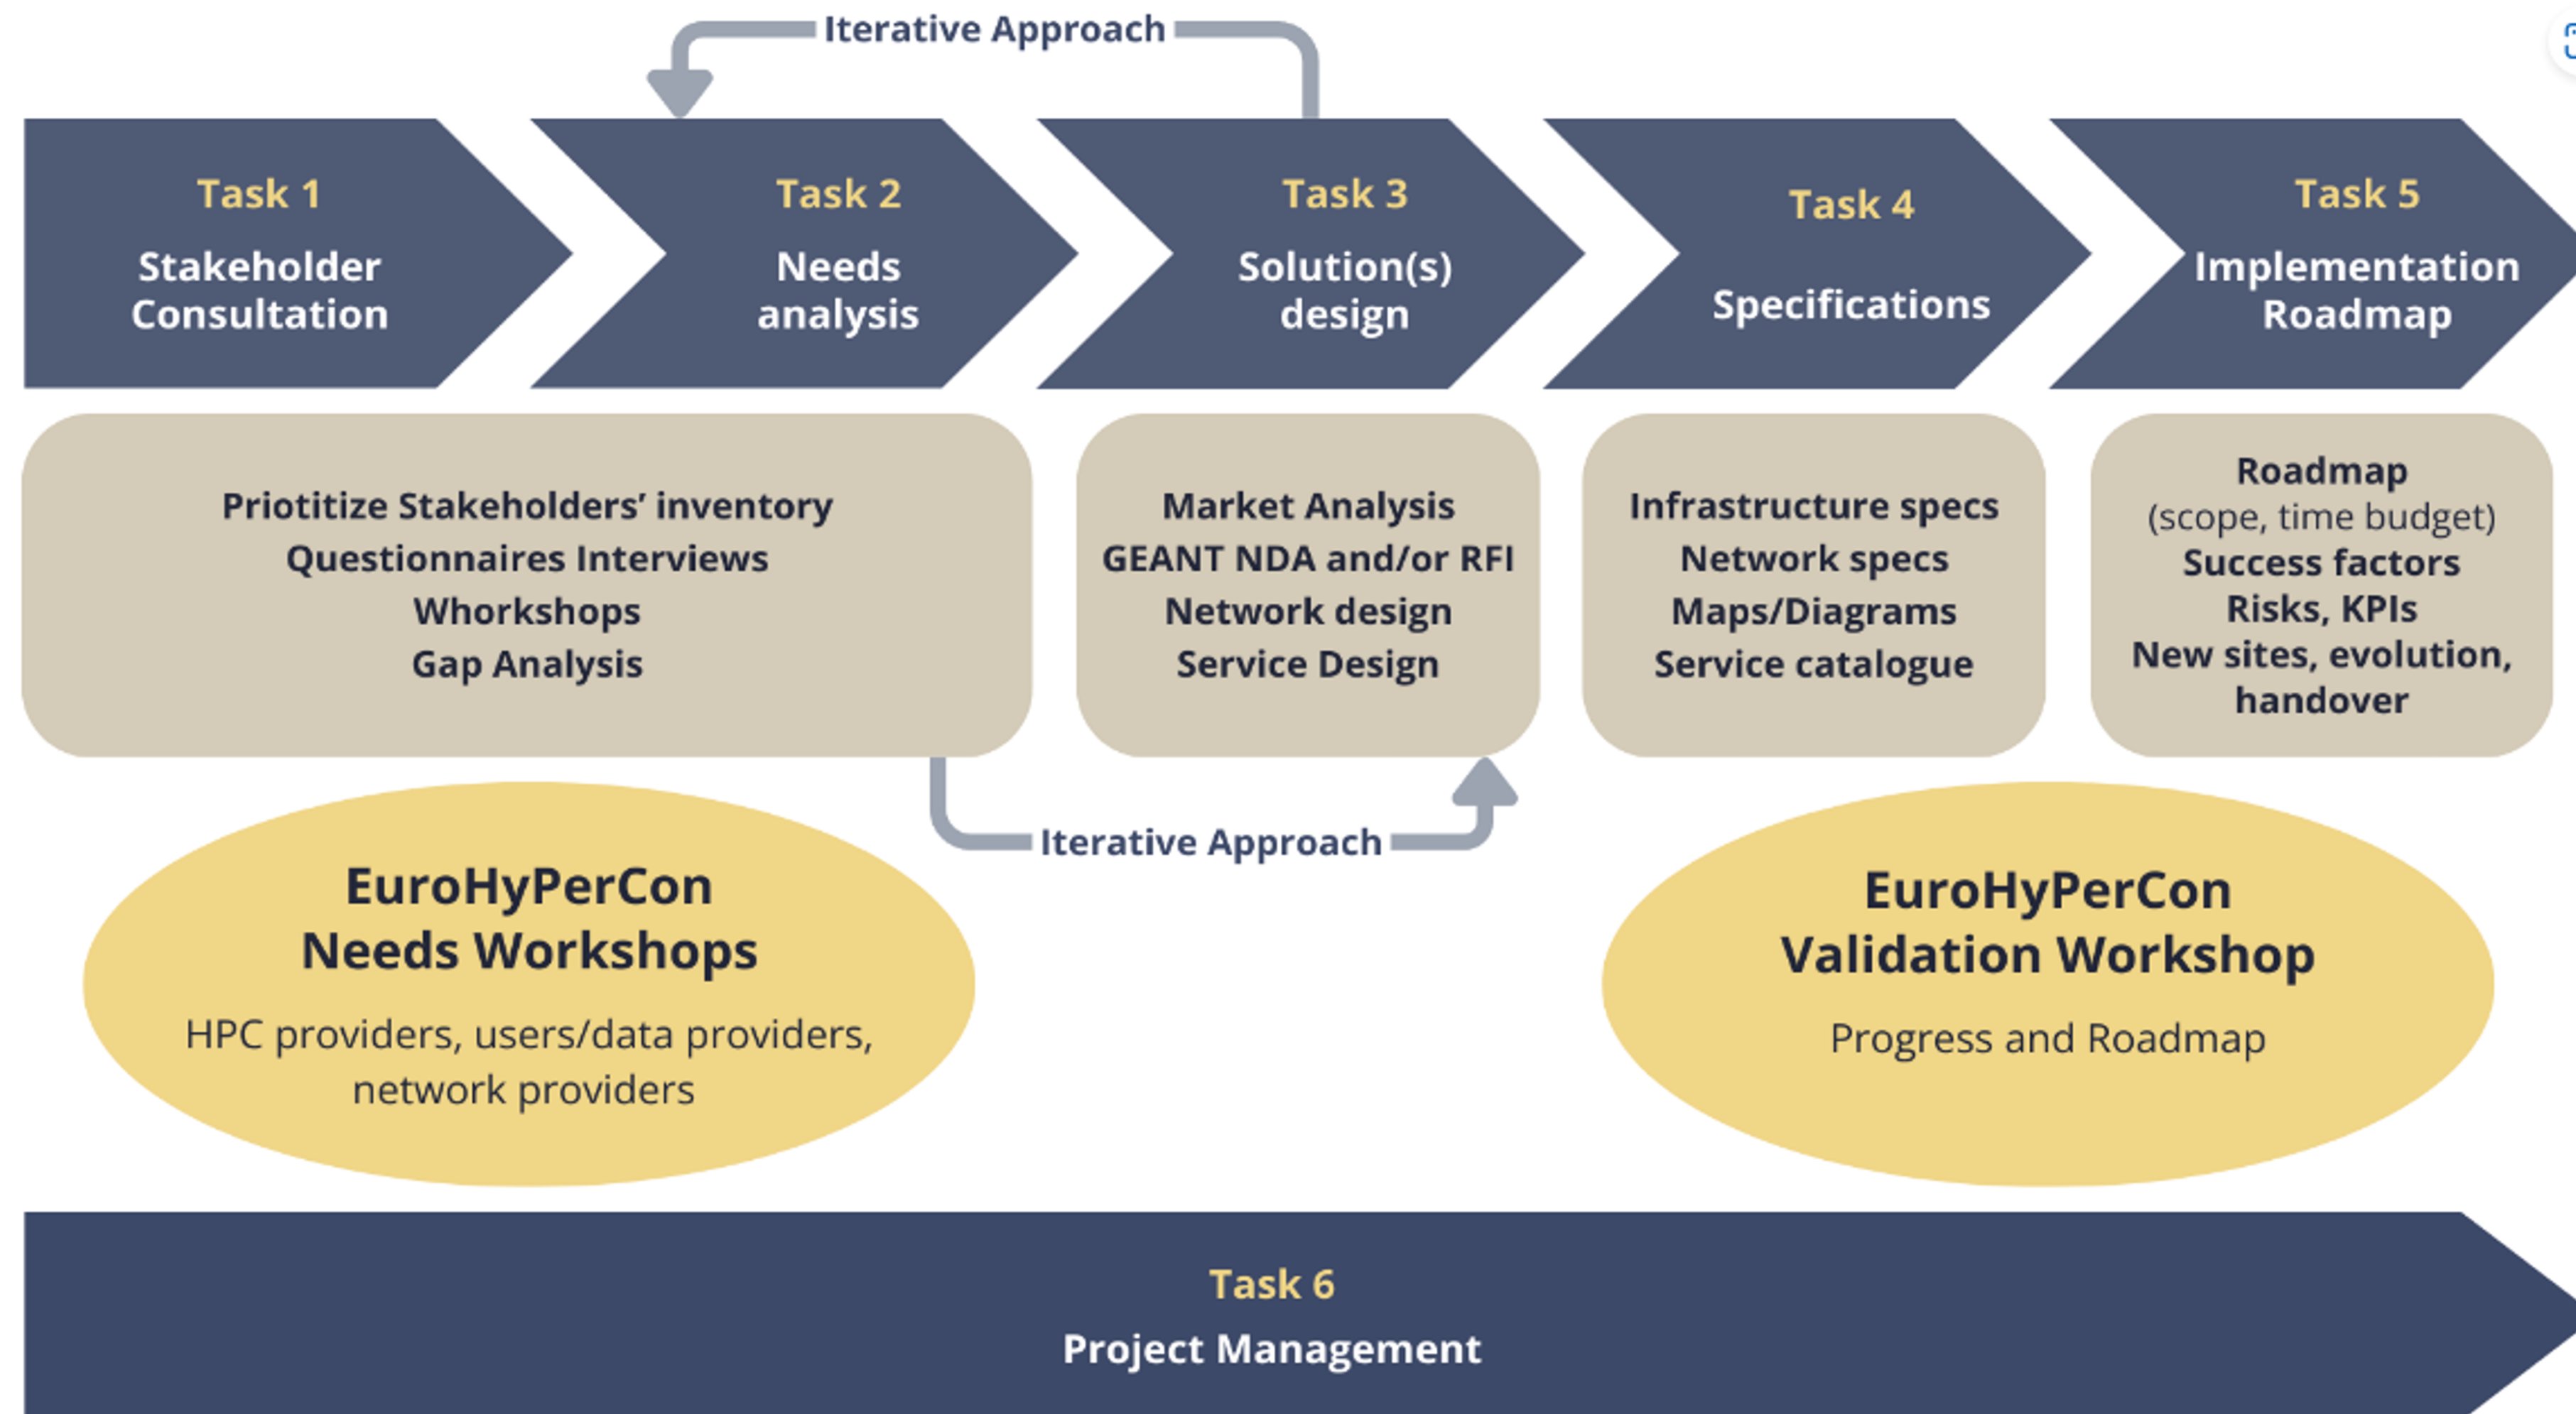
As outlined in the tender specifications:

- **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



HPC Providers

- **EuroHPC** Hosting Sites
- **Other EU / National** HPC systems



HPC Users

- **Thematic users** of the HPC systems
- Big users (e.g., DestinE (ECMWF, EUMETSAT, ESA), CERN, etc.)
- Other users



Network Providers

- **GÉANT**
- **NRENs** and **regional** research networks
- Other connectivity providers



Data Providers and AI Users - Other Stakeholders

- **Data providers** (e.g., ESFRI & Other RIs, EU Data Spaces)
- **AI users**
- Online Registration Form



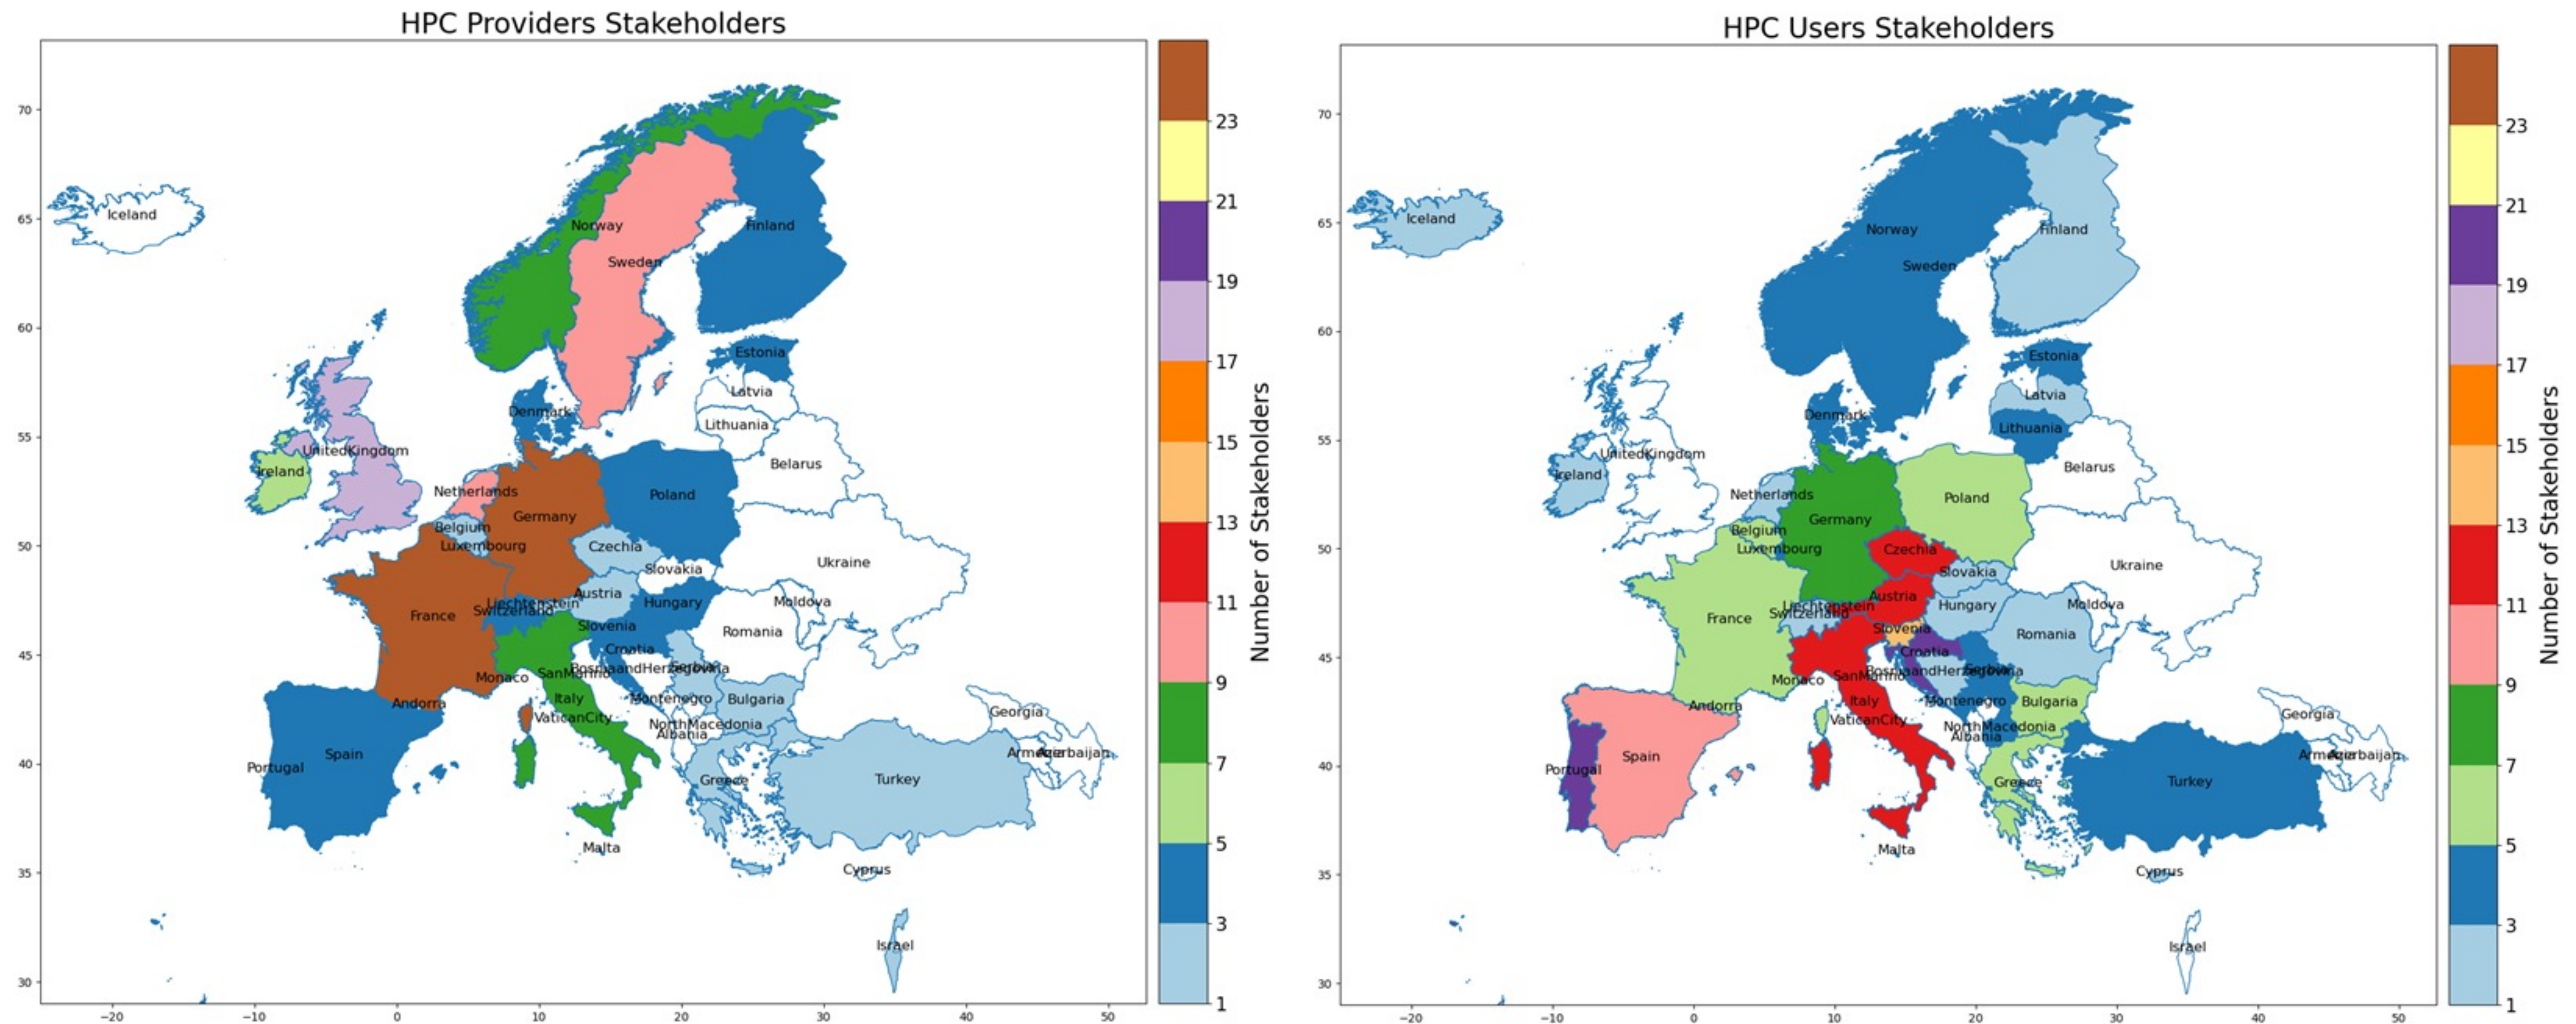
Activities performed



- **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



Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (1)



>>> nautobot

Search Nautobot

>>> ORGANIZATION

LOCATIONS

Location Types

Locations

TENANCY

Tenants

Tenant Groups

>>> DEVICES

FOTIS

Tenants

Search Tenants

ConfigureFilterExport

>>> Tenants

Name	Tenant group	Description
AWE	—	AWE
Academic Computer Centre in Gdansk	—	Academic Computer Centre in Gdansk
Barcelona Supercomputing Center	—	Barcelona Supercomputing Center
CALMIP / University of Toulouse	—	CALMIP / University of Toulouse
CEA/TGCC-GENCI	—	CEA/TGCC-GENCI
CINECA	—	CINECA
CNRS/IDRIS-GENCI	—	CNRS/IDRIS-GENCI
CSC (Center for Scientific Computing)	—	CSC (Center for Scientific Computing)
Cambridge University	—	Cambridge University
Cenaero	—	Cenaero
Center for Biological Sequence Analysis - DTU	—	Center for Biological Sequence Analysis - DTU
Commissariat a l'Energie Atomique (CEA)	—	Commissariat a l'Energie Atomique (CEA)
Commissariat a l'Energie Atomique (CEA)/CCRT	—	Commissariat a l'Energie Atomique (CEA)/CCRT
Cyfronet	—	Cyfronet
DKRZ - Deutsches Klimarechenzentrum	—	DKRZ - Deutsches Klimarechenzentrum
Deutscher Wetterdienst	—	Deutscher Wetterdienst
ECMWF	—	ECMWF

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (2)



>>> nautobot

Search Nautobot

>>> ORGANIZATION

LOCATIONS

Location Types

Locations

TENANCY

Tenants

Tenant Groups

>>> DEVICES

FOTIS

Locations

Search Locations

Configure

Filter

Export

>>> Locations

Name	Status	Parent	Tenant	Description	Tags
Austria	Active	—	—	—	—
• Informationstechnologielösungen (TU.it) High Performance Computing TU Wien Operngasse 11 / E020 1040	Active	Austria	Vienna Scientific Cluster	Informationstechnologielösungen (TU.it) High Performance Computing TU Wien Operngasse 11 / E020 1040 Wien, Austria	—
• Vienna Scientific Cluster, Austria	Active	Austria	Vienna Scientific Cluster	Vienna Scientific Cluster, Austria	—
Belgium	Active	—	—	—	—
• Cenaero, Charleroi, Belgium	Active	Belgium	Cenaero	Cenaero, Charleroi, Belgium	—
Bulgaria	Active	—	—	—	—
• SofiaTech park, Sofia, Bulgaria	Active	Bulgaria	HPC cetre Sofia Tech Park	SofiaTech park, Sofia, Bulgaria	—
Czech Republic	Active	—	—	—	—
• IT4I/VSB, Ostrava, Czech Republic	Active	Czech Republic	IT4Innovations National Supercomputing Center, VSB-Technical University of Ostrava	IT4I/VSB, Ostrava, Czech Republic	—
Denmark	Active	—	—	—	—
• Center for Biological Sequence Analysis - DTU, Denmark	Active	Denmark	Center for Biological Sequence Analysis - DTU	Center for Biological Sequence Analysis - DTU, Denmark	—
Finland	Active	—	—	—	—
• CSC, Kajaani, Finland	Active	Finland	CSC (Center for Scientific Computing)	CSC, Kajaani, Finland	—

Register as stakeholder: t.ly/e5FE8

Stakeholders mapping – Nautobot tool (3)



>>> nautobot

Search Nautobot

>>> ORGANIZATION

>>> DEVICES

DEVICES

Devices

DEVICE TYPES

Device Types

FOTIS

Devices

Search Devices

Configure

Filter

Export

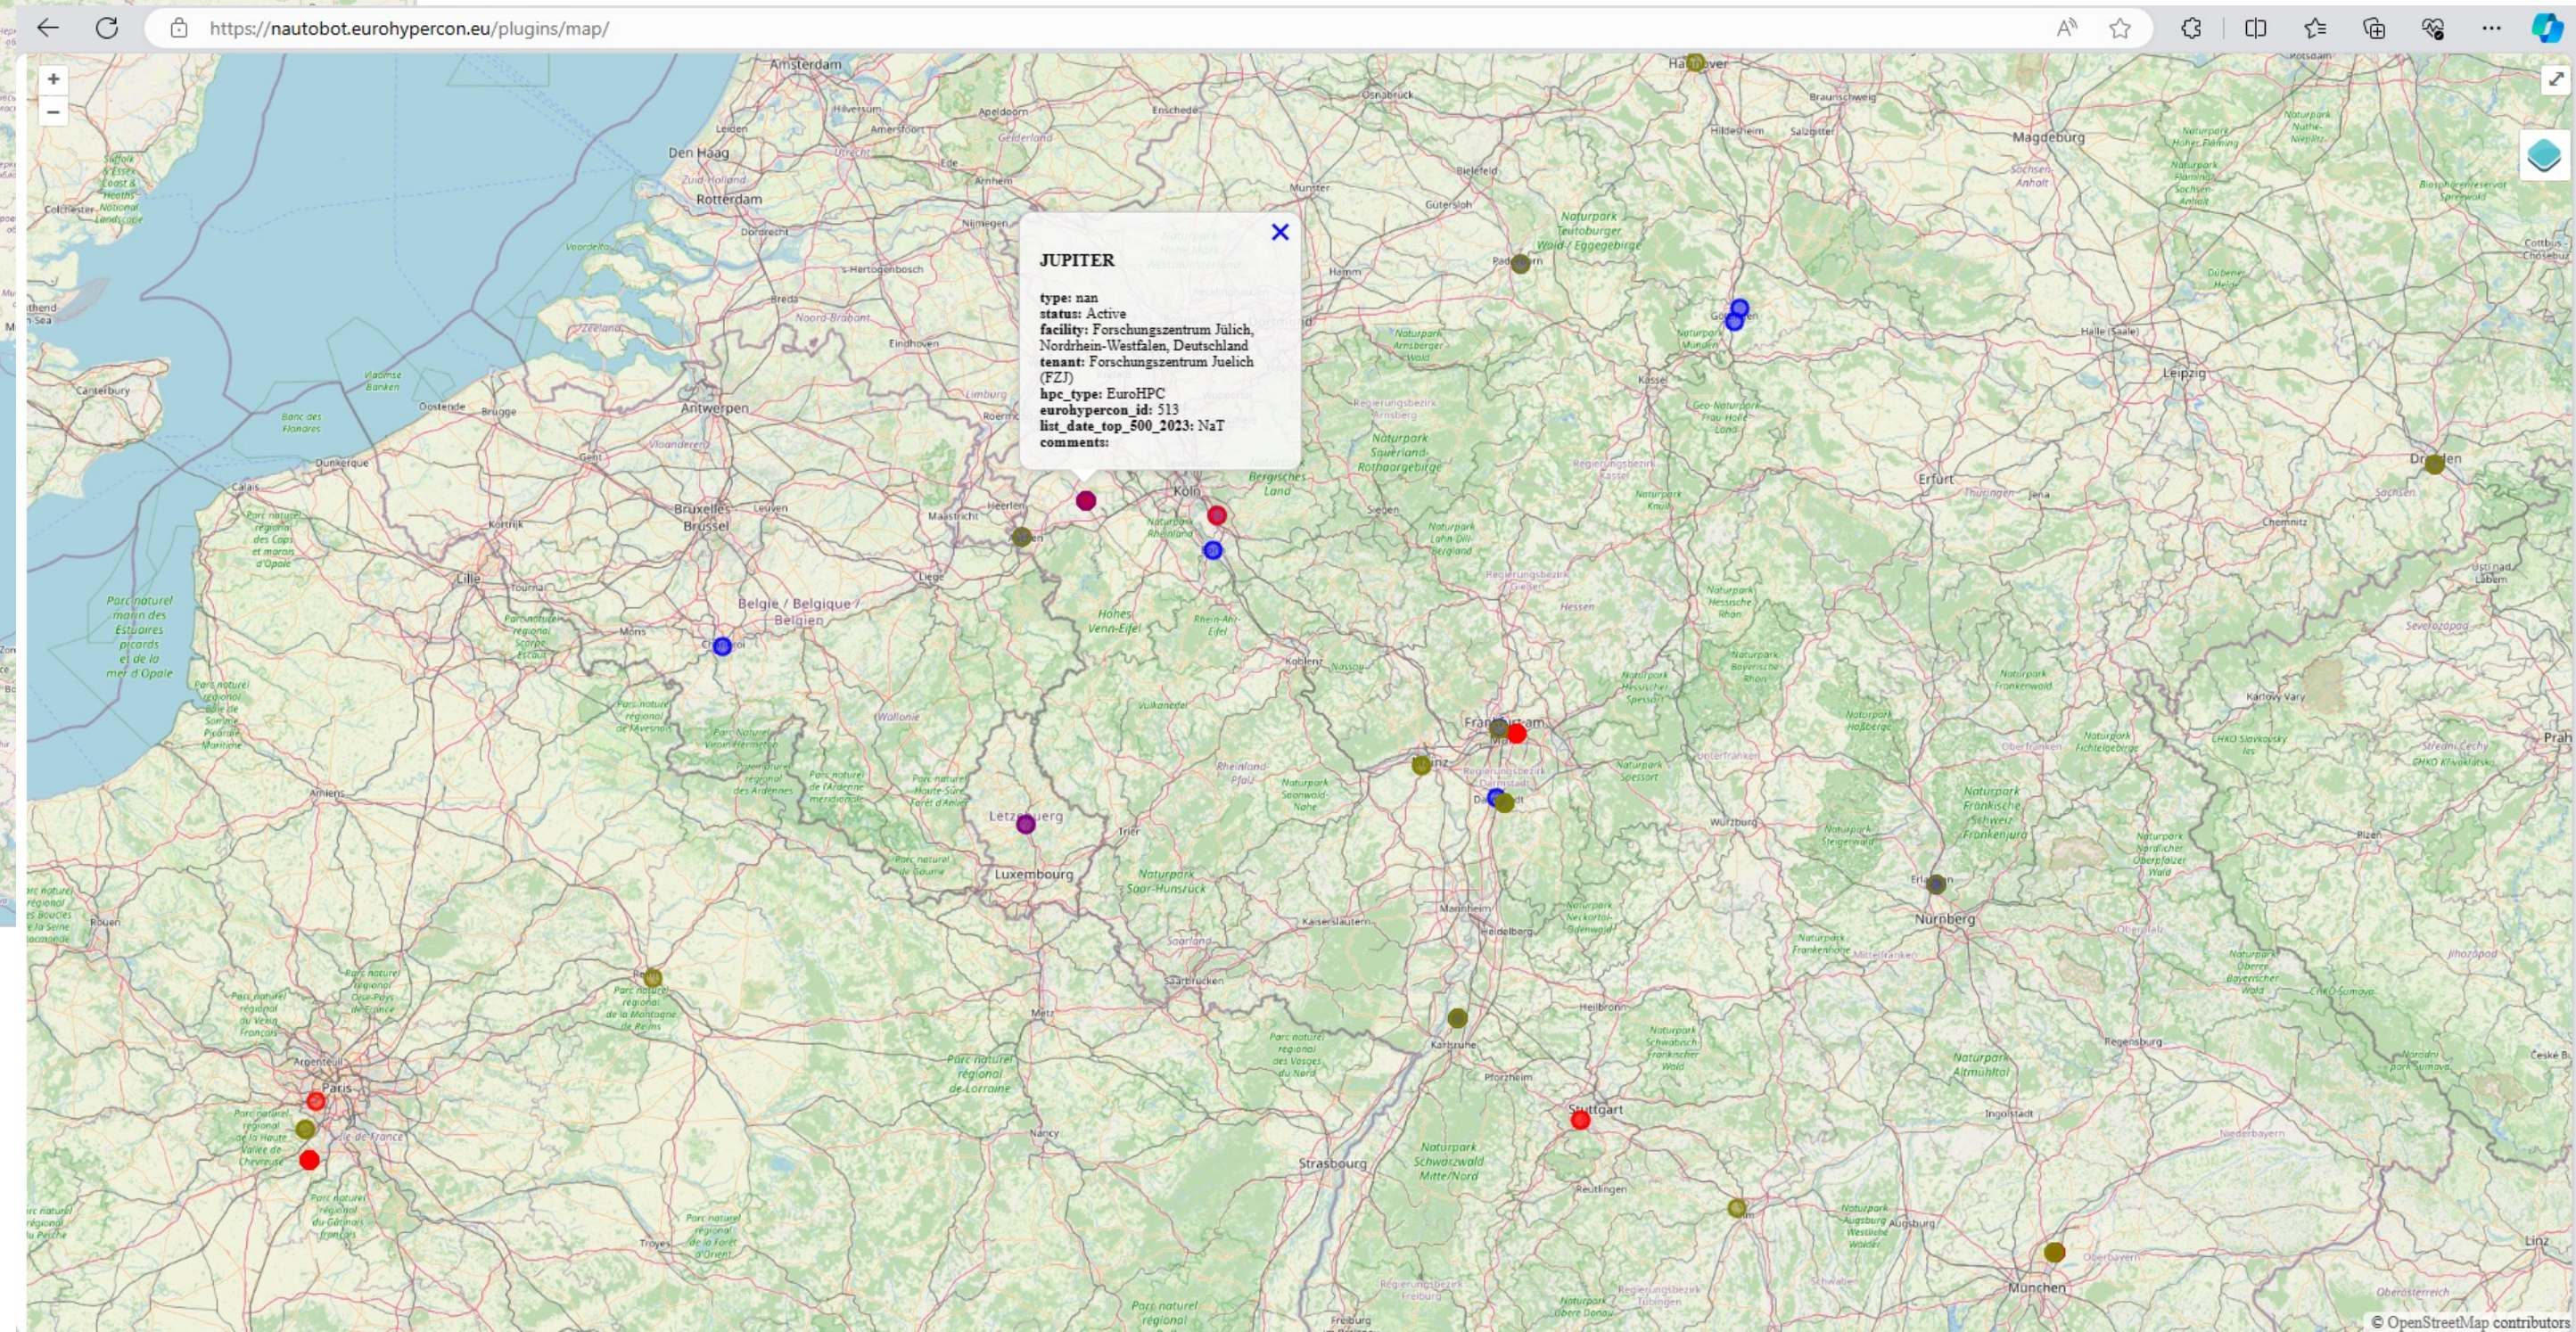
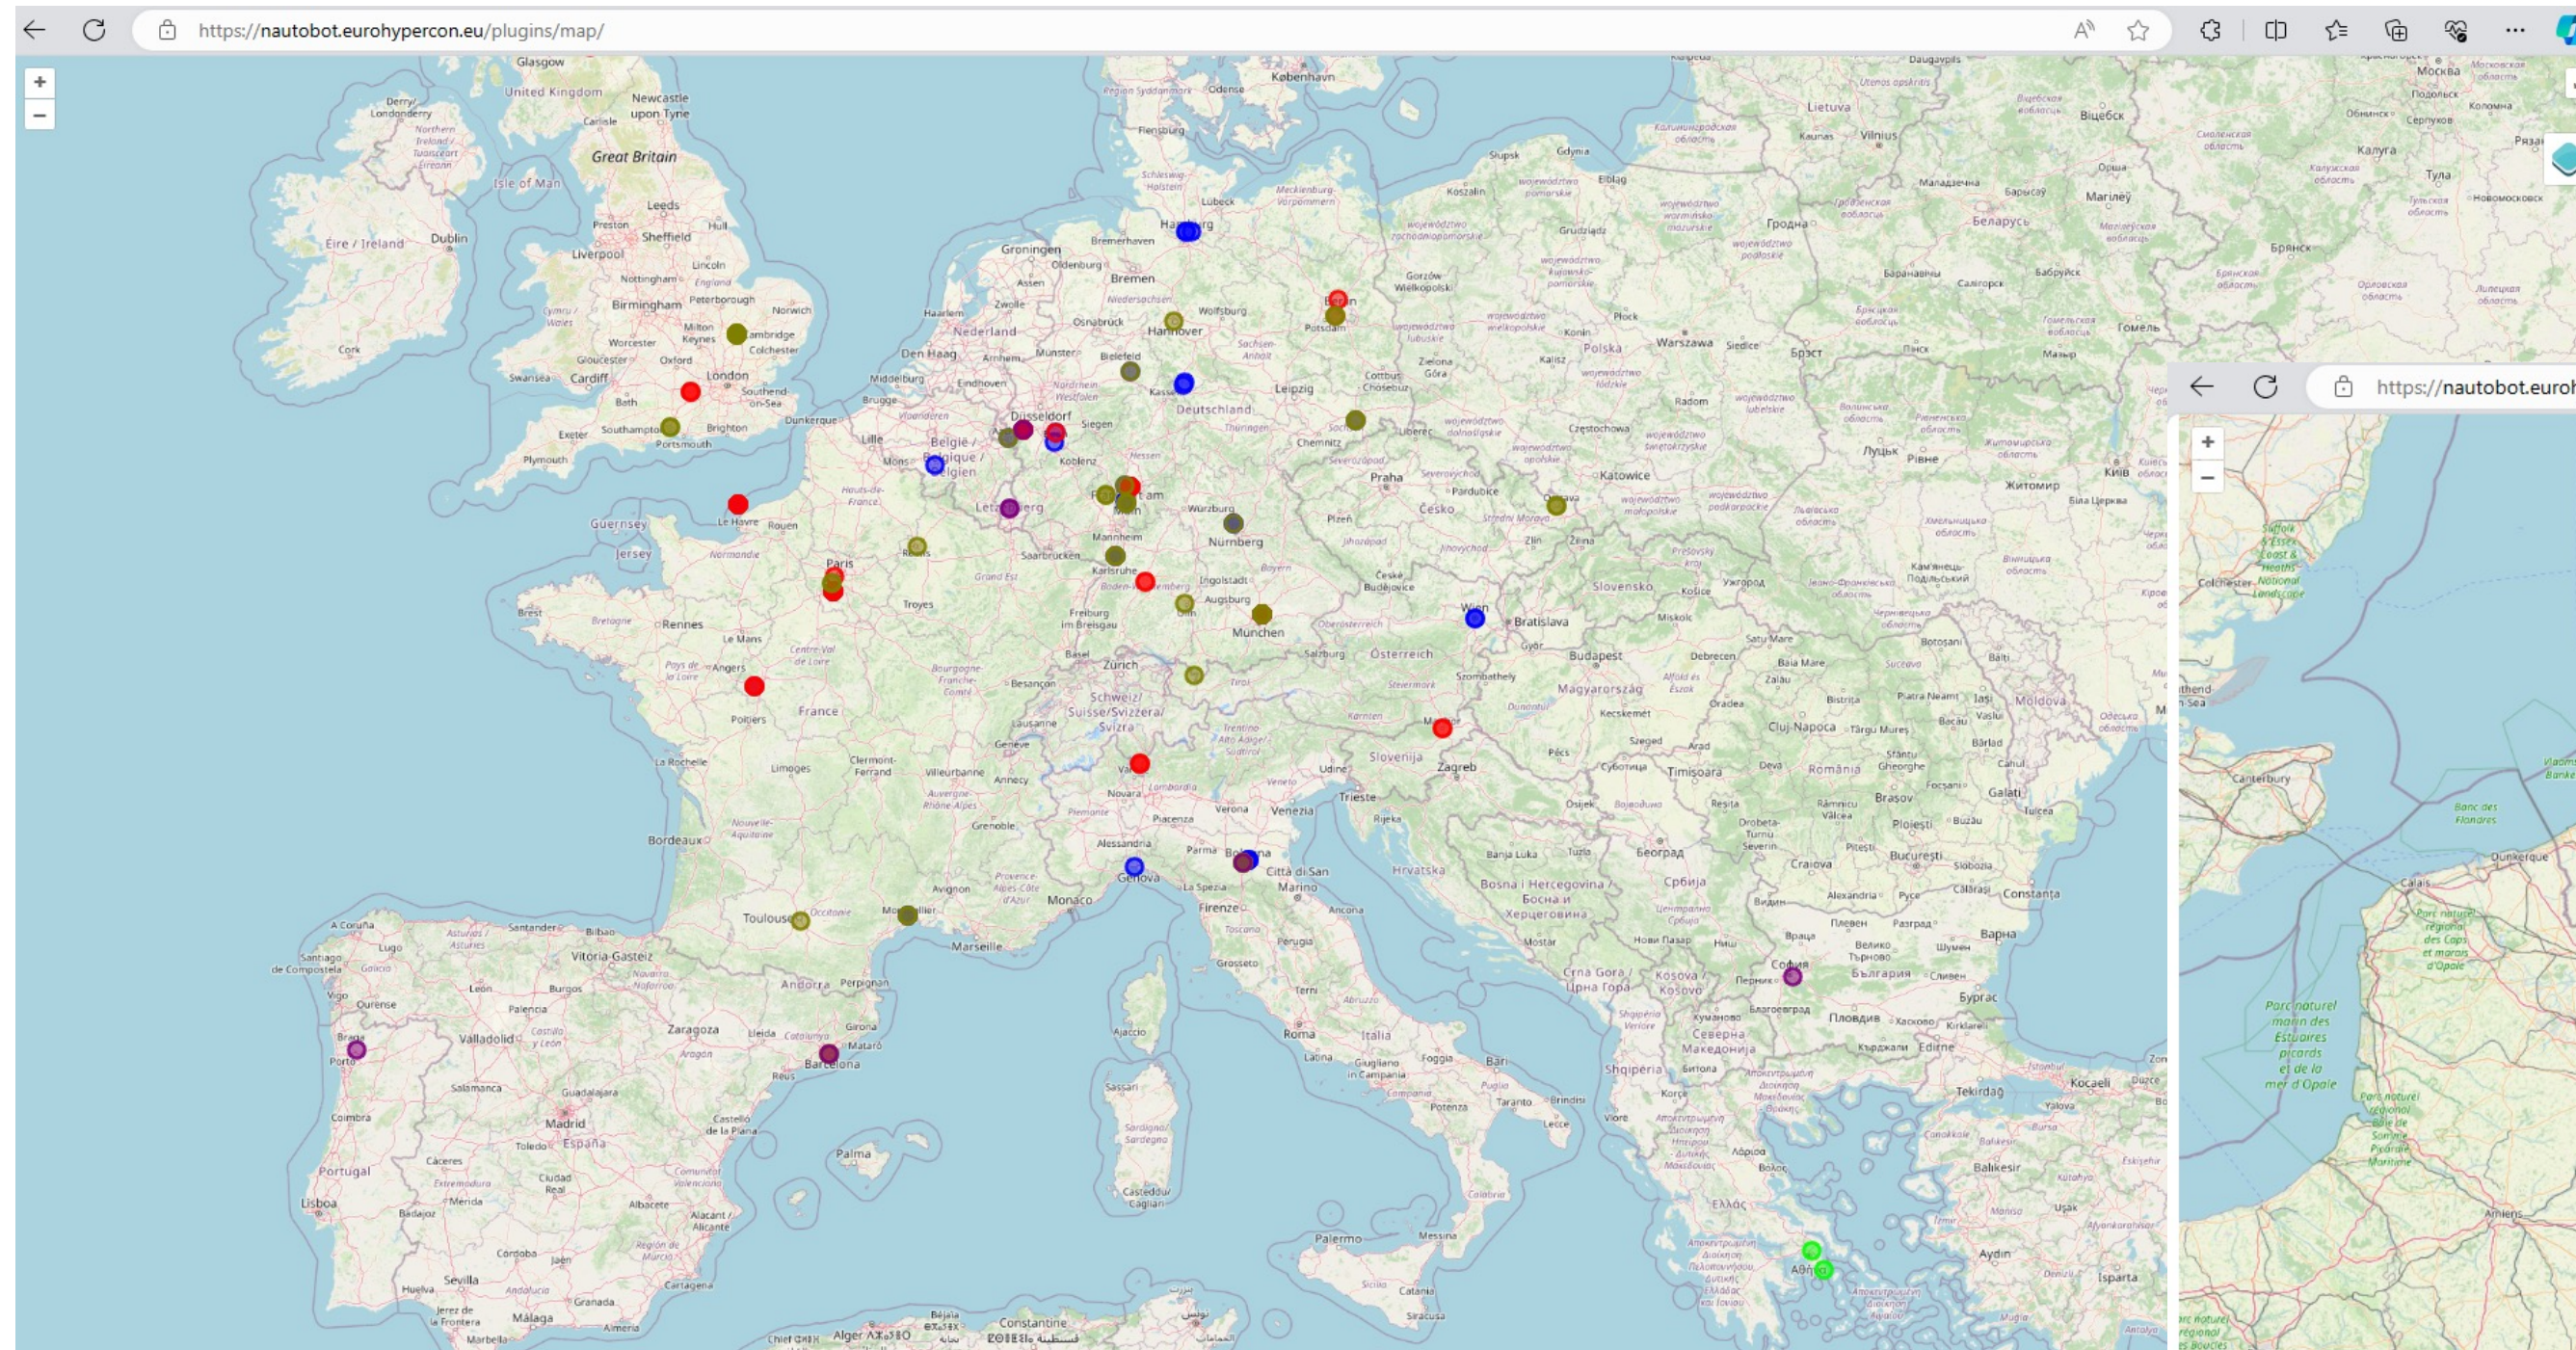
>>> Devices

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, Maroyi, 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



1st Workshop - Stakeholder Identification/ User Journeys



- **Rely on existing knowledge and solutions:** *GÉANT/NRENs, HPC centres, CERN/WLCG, ESFRIs.*
 - *NRENs and GÉANT are capable of delivering the connectivity that EuroHPC needs.*
- Pay attention also on **services on top of the network**. *Use network and move data in a smart way!*
 - *Hyperconnectivity has to be **integrated with storage, federation/authentication and security!***
 - *Services staging data transfer are critical for users who must transition workloads across EuroHPC sites.*
 - *For public shared data (e.g. AI models) consider a content delivery network (CDN) (proxies).*
- Consider **user communities**, their data sources and capabilities, as well as “novel technologies”
 - *How SMEs connect to academic networks will vary according to country regulations.*
 - *Discuss with the **AI community** the needs and availability of large data sets on HPC.*
- **Leverage (public+other) Clouds, taking into account data-related costs** (ingress, egress, storage)
 - *Connectivity and coordination **with large repositories/data providers** (national, thematic, e.g. ESA for earth observation data) for data gathering and for storing processed data from HPC.*



2nd Workshop – HPC users and providers



- *The **current networks are sufficient** and do not face any issues*
 - *Still, need to evolve/upgrade, taking into account **big user requirements***
 - ***Destination Earth (DestinE) is a champion user***
- ***Security** is key and may affect the (perceived) network performance one way or another*
 - *Need solutions that can utilise the full potential/capacity of the network required*
 - *Different challenges and approaches at networking or application levels discussed*
- *The EuroHPC JU **federation call** & access across sites are very relevant to this study*
 - *EuroHyPerCon will try to take into account related developments and inputs, as well as provide input to the federation call! (e.g. recommendations)*



3rd workshop – Meeting with GÉANT and EuroHPC JU/ Focus groups with Exa-/Pre-Exa NRENs



- GÉANT relies on **commercial connectivity providers** (dark fibre, spectrum) & **equipment vendors**
- **Optical layer - Future proof:** At the DWDM layer, GÉANT dark fibre links are expected to provide an average of 24Tbps of potential aggregate capacity. All GÉANT contracts with commercial entities for spectrum or dark fibre based on 15-year IRUs, plus 3x2 year extension options (up to 21 years).
- **IP layer – Future proof - Network upgrade planned (Nokia) - Easy expansion:** “GÉANT awarded NOKIA (June 2023) the contract to replace the IP/MPLS network. **400Gbps** will become the baseline capacity of all GÉANT **backbone** links between IP/MPLS devices and **400Gbps** user **access** ports will be provided at every routing node in GÉANT. The devices provide 36x 400Gbps per card. The line cards have 800Gbps-capable interfaces. While the Ethernet **800G** standard is still in development, such interfaces can be used to provide **2x400Gbps** in the meantime”.
- GÉANT/NRENs interconnect vast majority (if not all) of **academic/research users in Europe** and beyond; At least 80% of the computing time will be granted to European R&E users. NRENs also interconnect **all EuroHPC hosting sites and quantum sites** (and via GÉANT to EU/world).
- GÉANT and the NRENs are **dedicated overprovisioned networks for research**.
- Cost information received from GÉANT (anonymous) for EuroHyPerCon network cost estimation.



Interviews with DestinE (ECMWF, EUMETSAT, ESA)

- **DestinE simulations relevant** to EuroHyPerCon; ECMWF weather forecast models (multiple daily runs) not directly relevant (are considered institutional needs)
- **Simulations** on extreme events and digital twin climate adaptation
 - More simulations will come (volcanos, tsunamis, etc.)
 - Simulation events are **1 PB per run**; 1 PB for whole earth; extreme events in parts of the earth; Filtered to **100TB**, possibly even to **10TB** (trickier)
- Working on **AI weather models** in cooperation with industry
 - AI weather models trained on huge datasets with historical weather analyses (reanalyses); **30 PBs** of data (which may need to be moved around). **Not time-critical.**
- **Discussed the DestinE data and network architecture**
 - Collocation of DestinE/Eumetsat data servers in/near EuroHPC hosting sites; commercial provider solution (CloudFerro)
 - With Terabit network can consider on-the-fly data movement



EuroHyPerCon questionnaires



- Questionnaires → 154 full responses (>400 partial)
 - HPC Users: 102
 - HPC Providers: 32
 - Network Providers: 22
- Initial deadline was 19/1 (first phase) – More inputs received, e.g. AI stakeholders (iterative approach)

The screenshot shows the LimeSurvey 'Survey list' page. At the top, there's a navigation bar with the LimeSurvey logo, a '+ Surveys 3' button, and links for 'Help' and 'Configuration'. On the right, it displays '9718' responses, '12.9 / 1024' progress, '3' notifications, and a user profile for 'eurohypercon' with an 'Upgrade plan' button.

Below the navigation bar, the 'Survey list' section has two tabs: 'Survey list' (active) and 'Survey groups'. There are search filters for 'Search:', 'Status:' (set to '(Any)'), and 'Group:' (set to '(Any group)'). A 'Search' button and a 'Reset' button are also present.

The main table lists three surveys:

<input type="checkbox"/>	Survey ID	Status	Title	Group	Created	Owner	Anonymized responses	Partial	Full	Total	Closed group	Action
<input type="checkbox"/>	283724	Active	EuroHyPerCon Questionnaire for Network Providers	Default	11.12.2023	eurohypercon	No	52	22	74	No	...
<input type="checkbox"/>	727227	Active	EuroHyPerCon Questionnaire for HPC Providers	Default	08.12.2023	eurohypercon	No	74	32	106	No	...
<input type="checkbox"/>	788866	Active	EuroHyPerCon Questionnaire for HPC Users	Default	07.12.2023	eurohypercon	No	329	102	431	No	...




Questionnaires' analysis – Work in progress

- **Feedback from users' questionnaires show mostly national/regional needs**
 - Some active countries have more responses
 - Some cross-border (EU) needs
 - **Some countries could benefit from more answers**
- Additional inputs from all other means (workshops, interviews, focus groups, etc.)
 - Inputs are analysed and cross-checked
- **Last chance to influence the study – Questionnaires will be closing end of March**
 - <https://eurohypercon.eu/surveys/>


Users Questionnaires – Some statistics (1/2)




Type of HPC usage - Now




HPC usage timeframe




Security requirements?



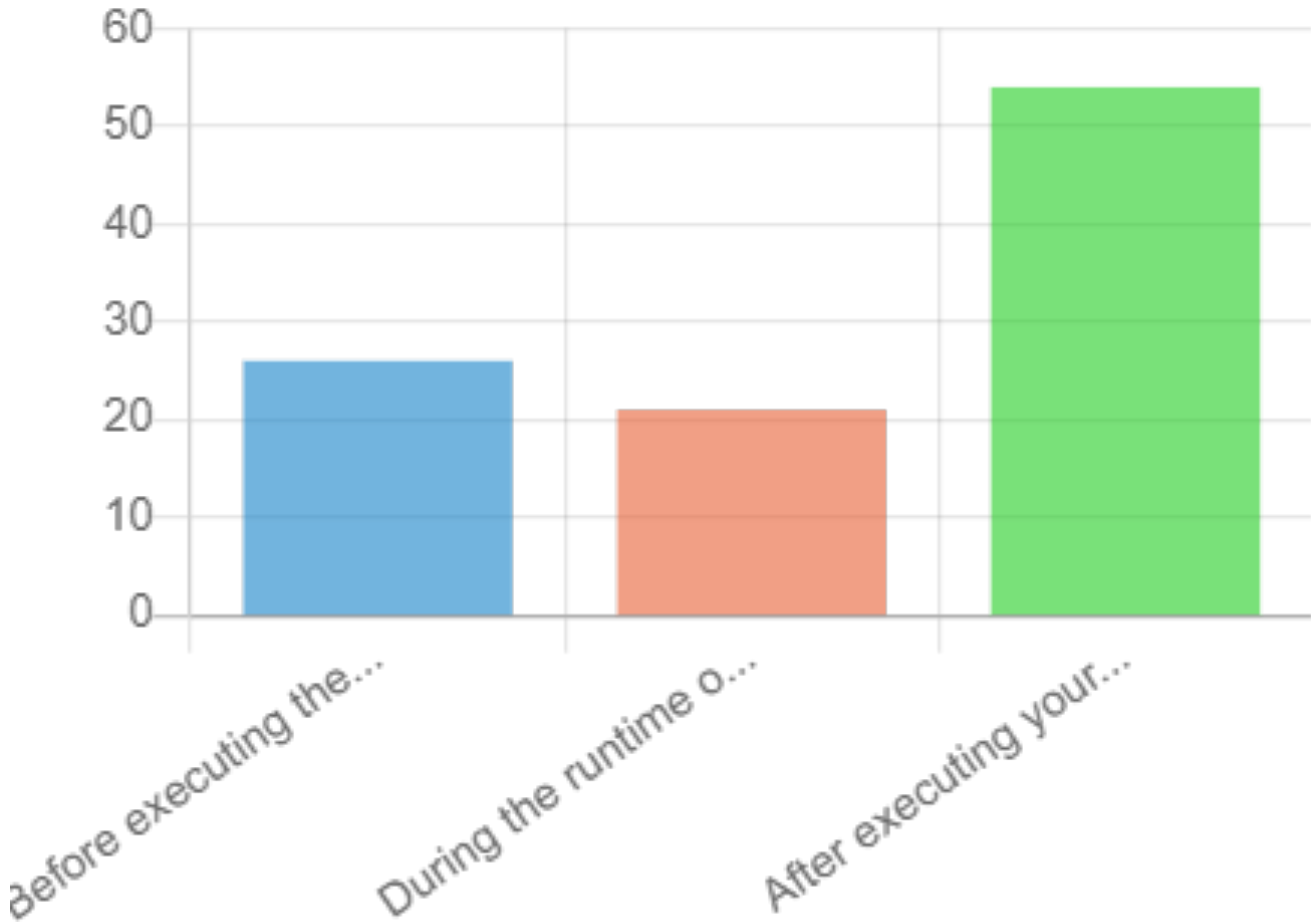
Type of HPC usage - 2030



Amount of resources used



Most data transferred...






Users Questionnaires – Some statistics (2/2)




Data size




Work in progress

Data transfer times




Satisfied with data transfer times?



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

~500%



>200%

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



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



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




- **Approach A - GÉANT / NRENs**

- Cost information from GÉANT (anonymous)
 - Need also cost info from a good sample of NRENs (East/West, North/South, big/medium/small)
 - For both: CAPEX/OPEX
 - Optical (layer 1) and IP/MPLS (Layers 2/3)
 - Cost evolution: 2025, 2028, 2030+
- Proposed approach: Incremental cost based on GÉANT and sample of NRENs



- **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)
- Following the current DestinE model (which is based on a commercial provider)
- Cost evolution: 2025, 2028, 2030+
- Commercial Service based on IP transit + Local Loops (DWDM / Metro Ethernet)





Conclusions and Next steps



- Stakeholders' identification performed
 - Last chance to influence the study with connectivity requirements
 - Surveys will close ~ at the end of the month
- Needs analysis in progress
 - Complete needs analysis
 - Proceed with Gap analysis (from current solutions)
- Alternative solutions design
- Techno-economic analysis ongoing
 - Proceed with technical assessment: fitness for purpose, performance, support for innovation, security, impact..
- Validation workshops (~May)
 - One online and possibly one f2f at ISC 2024



Questions?



More info

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



Thank you!