EuroHyPerCon

Connectivity Requirements and Network Design for HPC Systems in Europe

Summary Report

The EuroHyPerCon study ("EuroHPC/LUX/2022/OP/01 - Study for hyper-connectivity for HPC resources") is funded by the EuroHPC JU under contract number– LC-02450379. For more information, please check https://EuroHyPerCon.eu.

Disclaimer:

The information and views set out in this report are those of the author(s) (EuroHyPerCon consortium) and do not necessarily reflect the official opinion of the EuroHPC JU. The EuroHPC JU does not guarantee the accuracy of the data included in this report. Neither the EuroHPC JU nor any person acting on the EuroHPC JU behalf may be held responsible for the use which may be made of the information contained therein.

The information contained in this document represents the views of EuroHyPerCon consortium as of the date they are published. The EuroHyPerCon consortium does not guarantee that any information contained herein is error-free, or fully up to date. The EuroHyPerCon Consortium makes no warranties, express, implied, or statutory, by publishing this document. The information and guidance provided in this document should not be considered legal advice. This document is not intended to replace professional legal advice.

ISBN 978-92-95228-04-7

doi: 10.2793/4687161

Table of Contents

1. Introduction

Over the past five years, the EuroHPC Joint Undertaking (EuroHPC JU) has made significant investments in developing, deploying, and sustaining a world-class integrated supercomputing and data infrastructure in Europe, composed of a series of national and multinational High-Performance Computing (HPC) systems. The EuroHPC JU aims to create a secure, hyper-connected and federated, HPC and quantum computing service in the European data infrastructure ecosystem. The goal of this highly competitive and innovative federated HPC infrastructure and associated ecosystem is to put Europe at the forefront of the world, capable of helping both science and industry address major scientific and economic challenges, including global challenges such as climate change, earthquakes, and pandemics. In this way, EuroHPC will foster world-class scientific excellence and industrial innovation driven by European leadership, underpinned by ubiquitous, trusted, secure and easy-to-access production-quality HPC with quantum computation services and data, supported by a secure communication and collaboration environment.

Figure 1: The EuroHPC Joint Undertaking

The EuroHyPerCon study has been a key component of this initiative, with its central focus on identifying and analysing the connectivity needs of the European HPC and data ecosystem. Such an ecosystem includes not only the EuroHPC co-funded HPC systems, but all the relevant European and national supercomputing and data infrastructures, encompassing the related end-user landscape. The project aimed to design a future-proof connectivity service, complete with a set of detailed specifications for the upcoming Hyperconnectivity tender, along with an implementation roadmap for the future.

This summary is organised according to key areas of focus: Stakeholders Analysis, Needs Assessment, Gap Analysis, Analysis of Alternative Solutions, Technical Specifications, and Implementation Roadmap.

2. Executive Summary

2.1. Introduction and Purpose of the Study

The EuroHyPerCon study is a strategic initiative designed to address the growing demand for high-performance computing (HPC) and data infrastructures in Europe. With advancements in research areas such as artificial intelligence (AI), climate science, and other data-intensive fields, the need for a secure and federated hyperconnectivity service is evident. EuroHyPerCon's primary goal has been to develop a roadmap for a European HPC network that meets the evolving connectivity, performance, and security needs of diverse users across Europe, both now and in the coming years.

This study focused on the infrastructure and service requirements essential for creating a robust network capable of supporting HPC resources across Europe. The project scope included identifying connectivity gaps, assessing technical requirements, and providing a phased implementation plan that builds upon existing European network infrastructures such as GÉANT and National Research and Education Networks (NRENs). Ultimately, this initiative sought to establish a cost-effective and sustainable service model to ensure Europe remains at the forefront of global HPC advancements.

First, the Key findings and Recommendations are presented and then the rest of the outcomes of the study.

2.2. Key findings

The key findings that were produced during the 11-month study are directly presented below:

- **Stakeholders database**
  The study identified a comprehensive list of stakeholders through extensive consultation with the HPC ecosystem community and thorough review of the state-of-the-art. A total of 558 stakeholders with 742 entries have been included in the databases, as some HPC providers operate multiple systems. The corresponding database consists of HPC providers (163 providers managing 346 HPC systems), Data providers (149), Network providers (46), HPC users (176) and AI stakeholders (24).

  A variety of tools and methods were employed to ensure the study's success, including:

  - **Workshops**: Three dedicated workshops were organized for HPC providers, HPC users, and network providers.
  - **Focus Group Meetings and Interviews**: These involved EuroHPC hosting sites, major HPC users (e.g., Destination Earth), and AI stakeholders.
  - **Presentations**: Delivered at HPC and other significant events.
  - **State-of-the-Art Review**: This included among others analysing Top-500 HPC lists from recent years.

- **Dedicated Questionnaires**: Tailored for HPC users, HPC providers, and network providers.

In all cases, there was a high level of interest and valuable contributions from participants.

- **Analysis of HPC user connectivity requirements**
One of most important outcome of the study was the collection and analysis of hyperconnectivity requirements for the European HPC ecosystem, encompassing HPC providers, data providers and HPC users. A key tool for this analysis was the HPC Users' Questionnaire, supplemented by complementary and cross-referenced information from the HPC Providers' and Network Providers' Questionnaires. In total, over 600 responses were received. However, due to the complexity of the questionnaires, 165 fully completed responses were ultimately analysed. These included 111 complete responses from HPC users (encompassing data providers and AI users), 32 from HPC providers, and 22 from network providers.

The study produced detailed hyperconnectivity requirements tables, outlining specific data transfer needs across Europe and beyond. These tables included data categorised by site, HPC-data site pairs, and country, offering aggregates and other valuable insights. This information was instrumental in defining the Points of Interest (PoIs) for the Hyperconnectivity service specifications tender.

Given the challenge of obtaining comprehensive responses from all stakeholders and ensuring equal participation across countries, the PoIs list for the tender was further enriched using the stakeholder database developed at the start of the study. For instance, while participation from Germany and France in the HPC Users' Questionnaire was relatively low, these two countries are the top ones in the number of PoIs included in the tender specifications.

- **Gap Analysis**
The study also conducted a Gap Analysis, comparing the current network infrastructure with the desired hyperconnectivity requirements. The vast majority of HPC systems identified during the consultation are currently interconnected via the pan-European research network GÉANT and the National Research and Education Networks (NRENs). Additionally, approximately 90% of the HPC users who responded to the questionnaires are from academia and research, with the remaining 10% representing industry and SMEs. Consequently, the gap analysis focused on the current network infrastructure supporting these users.

HPC users expressed general satisfaction with the services provided by GÉANT and the NRENs. GÉANT and most NRENs operate under long-term fibre Indefeasible Rights of Use (IRUs), i.e. dark fibre infrastructures, typically spanning 15–20 years, and light these fibres with their own equipment at both the optical and Internet layers. These infrastructures are designed to be scalable, with upgrades possible at both the access and backbone levels, supporting speeds of n x 400 Gbps and even Tbps levels. GÉANT is also preparing to offer 400 Gbps connectivity for user access in the near future.

HPC users indicated a demand for both national and pan-European HPC resources and associated connectivity. While many pan-European requirements can be addressed by GÉANT, further upgrades and optimization are necessary to meet evolving demands. Additionally, Destination Earth was identified as a flagship user and data provider, requiring a tailored solution to ensure redundant, multi-homed connectivity for its data sites.

Regarding the gap analysis in terms of access speeds, it was estimated that EuroHPC hosting sites currently have average requirements of approximately 200 Gbps. These needs are projected to reach around 350 Gbps by 2025, requiring implementation by 2026. This represents a 1.75-fold upgrade in average access speeds. EuroHPC sites are identified as the most critical for hyperconnectivity, as they have the highest connectivity demands. In contrast, other HPC and data sites require relatively smaller upgrades compared to EuroHPC hosting sites.

Potential infrastructure gaps were identified, particularly in the "last mile" of the end-to-end network, which refers to internal infrastructure within HPC or data centres. For EuroHPC sites, no significant bottlenecks were identified for access speeds of 400 Gbps or lower. However, for speeds exceeding 400 Gbps, upgrades to border elements such as routers, switches, and firewalls may become necessary. Notably, some EuroHPC hosting sites, such as Leonardo and Vega, are already equipped to handle speeds of 800 Gbps or higher.

No major technology or standard gaps were identified. The IEEE standard for 800 Gigabit Ethernet which is important for the upcoming hyperconnectivity tenders, was approved in February 2024, and the availability of such interfaces in the market is steadily growing. Still, the combination of 100, 200, 400 and 800 Gigabit Ethernet interfaces will be sufficient to meet the study's requirements for 2025.

Administrative, Policy, and Legal Gaps were also analysed. Since the hyperconnectivity service must terminate within the premises of the proposed end sites (PoIs) included in the tender specifications, it is essential for these sites to agree to be connected. This will require formal approval from the end sites, including legal actions undertaken by their designated signatories. A dedicated process will need to be established during the tender contract implementation phase to secure these agreements.

This aspect should not be underestimated, as the majority of the PoIs are universities and research institutes, which often have slow and complex processes for providing such legal confirmations. Careful planning and coordination will be required to mitigate potential delays and ensure timely compliance.

- **Analysis of Alternative Solutions**
  An analysis of alternative solutions for the hyperconnectivity network have been analysed. These include:
  - *A. Expanding the existing Infrastructure for research*
    A network based on an upgraded GÉANT's – NREN network infrastructure
  - *B. IP Transit, based on commercial providers*
    A network based on Internet connections to each of the PoIs, plus Cloud providers and IXPs.
  - *C. Implementation of a new network (optical Infrastructure & IRUs)*
    A network based on creating a new infrastructure, including circuits and IP/optical equipment
  - *D. Mixed solution of A and B/C*

A hybrid solution expanding existing infrastructure & introducing commercial services for some segments of the network.

The assessment criteria for the analysis constituted the following:

- *Fitness for purpose:* Adequacy of the proposed solution to fulfil the HPC hyperconnectivity objectives.
  Support of efficient massive data transfers for HPC applications: "Jumbo" frames[1] + overprovisioning and peerings with international research networks / infrastructures.
- *Impact:* Contribution to the achievement of the EuroHPC JU strategy:
  Hyperconnected (towards terabit) and federated European HPC and Quantum ecosystem; open and inclusive ecosystem recognising diverse user groups, e.g. academia, SMEs, AI stakeholders.
- *Performance*: Guaranteeing/monitoring the performance and the quality of the service (e.g. throughput, delay, jitter, availability).
  Pro-active based measurements, efficient performance monitoring mechanisms, well-defined upgrade rules.
- *Security:* Preventing unauthorized access, misuse, or denial of service, and detecting, monitor and remediating such events.
  Balance between performance and security, preventing attacks/DDoS, facilitating innovative security architectures.
- *Support for Innovation:* Accommodating innovative services and technologies.
  Advanced networking initiatives, active research community, environment for applying new technologies, training.
- *Sustainability:* long-term maintainability of the proposed solution, incl. ease of operation, support, environmental footprint.
  Existing management infrastructure, reduced complexity, O&M cost, proven service provisioning, energy efficiency.

Solution A (expanding the existing infrastructure) holds a unique technical position, as it can support the type of HPC/big data traffic required with "Jumbo" frames end-to-end. Unlike oversubscribed commercial networks, it operates as an overprovisioned network, fostering innovation. It is the only solution that already interconnects almost all Points of Interest (PoIs) and has the potential to connect approximately 10,000 more. This will enable seamless integration within the ecosystem.

In this scenario, commercial providers would need to rely on Internet Exchange Points to interconnect with the rest of the ecosystem. By contrast, Solution A is built on publicly funded fibre infrastructure for both backbone (GÉANT) and access (NRENs) networks, providing a long-term, sustainable solution for 15–20 years. It leverages past investments and expertise while maintaining end users' trust.

[1] What are jumbo frames? | Definition from TechTarget. A jumbo frame is an Ethernet frame, or data packet, with a payload greater than the standard size of 1,500 bytes. This allows the transfer of high-volume data, relevant to the HPC and AI applications in better performance.

NRENs have a history of collaboration with HPC and data sites, demonstrating their capacity and reliability in managing the complex and large-scale network operations required for EuroHPC hyperconnectivity.

Solution A will allow:

- Fast implementation and deployment: Support for innovation through an established connectivity ecosystem.
- Cost Efficiency: Solution B is approximately 3 to 4.5 times more expensive than Solution A Furthermore, it offers reduced operating costs and long-term sustainability.
- There are peerings with all the major research networks of the world.
- Security and DDoS mitigation services are available.
- This solution will maintain the current IP address space (moving away would be complex and challenging).

Solution B (IP transit service by commercial providers) comes with the following advantages:

- Well known and mature service offering open to commercial word.
- Allows easier commercial access to HPC centres.

However, there is also a long list of disadvantages:

- Not suitable for the HPC traffic patterns (overbooked commercial networks vs overprovisioned academic networks)
- Likely cannot support Jumbo frames end-to-end
- IP addressing issues (many HPCs already use IP space provided by NRENs)
- Current investments/ infrastructure not utilised
- Increased network configuration and OAM complexity with each HPC/data centre having multiple upstream providers
- Sites may refuse to connect
- Interconnection with other research networks will be done through commercial peerings
- Not able to use existing last mile circuits (typically IRU based) and might involve a lengthy implementation period for last miles

Solution C (New optical and IP network) comes with the following advantages:

- Dedicated for HPC needs - allows more control to EuroHPC JU
- Long term solution covering to 2032 and beyond.

However, there is also a list of disadvantages:

- Not clear who is going to operate such an infrastructure. There may need to be a dependency on a new service provider.
- Current investments/ infrastructure are not utilised. This will also include a duplication of existing infrastructure and potential waste of public money.
- Lengthy implementation with multiple risks. It is not clear what will happen after the end of the hyperconnectivity tenders in 2032.
- This is the least cost-efficient solution, possibly an order of magnitude more costly than Solution A.

Solution D (Mixed solution of A and B/C) comes with the following advantages:

- Combines solutions A and B or C, with pros applying for the corresponding segments of the network from A and B or C.

However, there is also a list of disadvantages:

- Complex to manage and implement, as it involves multiple solutions and contractors
- Requirements for continuous co-operation between GÉANT/NRENs and any commercial providers (and EuroHPC JU managing it).
- GÉANT / NRENs already have pan-EU footprint; it would have made sense if some parts were not covered.

Cost/Pricing of Solutions A-D

The pricing comparison has been based on market analysis of existing commercial services and cost pricing from recent procurements. The study team used Data from TeleGeography, a specialized telecommunications market research and consulting firm that primary focus in providing data, analysis, and insights about the global telecommunications market, including pricing. And data regarding GÉANT and the NRENs. Expertise from a set of experts from the Athens University of Economics and Business, including professors of accounting, was used.
Based on the two hyperconnectivity phases, Phase A 2025 – 2029 with hyperconnectivity tender 1 (60M EUR) and Phase B 2029 – 2032, with Hyperconnectivity tender 2 (100M EUR).
Taking into account all the sites identified (all priorities, i.e. mandatory and optional) the following results have been produced:

- Solution A – GÉANT and NRENs: Phase A 61M EUR, Phase B 84M EUR; Total: 145M EUR
- Solution B – Commercial IP transit services: Phase A 269M EUR, Phase B 366M EUR; Total: 635M EUR
- Solution C – New network from scratch: Phase A: an order of magnitude higher (possibly > 1B EUR)
- Solution D – Mixed solution: Depends on portions of the network in A or B. If half each, then:
- Phase A: 165M EUR, Phase B: 225M EUR; Total: 390M EUR

Solution A (GÉANT/NRENs) scores better in almost all the assessment criteria. Solution D (Mixed solution of A and B/C) comes second. Solutions B (IP transit) and C (New network) follow, depending on the priorities. However, the cost is prohibitive, especially for Solution C.

- **Compilation of Service Specifications**

  Following the consultation with the EuroHPC ecosystem, the analysis of requirements and the gap analysis, the study team came up with the development of service specifications. The main components of the services specifications are presented below:

  - *Network-as-a-Service (NaaS) approach*

    The EuroHPC JU will not build or operate the network, rather pay for the hyperconnectivity service. The hyperconnectivity network will operate on a NaaS model, where the service provider's infrastructure is treated as a "black box", focusing on the interface between Points of Interest (PoIs) and the network provider.

  - *List of Points of Interest (PoIs) for the Hyperconnectivity network*

    The hyperconnectivity network will serve a wide range of PoIs across Europe, i.e. HPC and Data Providers, while additional peerings will be needed with the NRENs to serve the wider user community and with public(-ly available) Cloud Providers where several HPC users are storing their data. Finally, connection with major Internet Exchange Points (IXPs), EU and national, would be needed to provide interconnection among the industrial and research communities. Regarding the list of PoIs, these have been categorized based on priority. Priority 1 are the EuroHPC hosting sites, Priority 2 are other major HPC and data sites and Priority 3 are HPC and data sites with connectivity requirements above 10TB per day. These priorities will be mandatory in the tender. Priority 5 and 6 PoIs are other non-major HPC and Data sites with requirements lower than 10 TBs per day and are optional. At the end, there is a list of 400 points to be hyperconnected, both mandatory and optional, including PoIs, NRENs, Cloud Providers and IXPs.

- *Plain IP connectivity to PoIs*
  The study team proposed to use plain IP connectivity for the hyperconnectivity network. Overlays (i.e. virtual provide networks) are not practical for this scale. A set of access speeds ranges based on access classes have been defined. The backbone network must ensure that services provided meet functional, quantitative, and qualitative criteria to support HPC traffic, based on non-overbooked, high-performance infrastructure. The so-called “Jumbo” frames, i.e. Ethernet frames usually with 9000 bytes of payload, are mandatory for the hyperconnectivity network, to enable large data transfers needed for HPC and AI applications.

- *Resilience, Redundancy and Latency Minimisation*
  The design of the network must eliminate single points of failure, ensuring uninterrupted service continuity, supporting automatic failover mechanisms. The network must also be designed to minimise latency by routing data as close to the user base as possible, leveraging strategic interconnections and peering arrangements.

- *Enhanced Performance Monitoring and Service Level Agreements (SLA)*
  The network must be monitored through a dedicated performance monitoring infrastructure. The list of hyperconnected PoIs must be maintained with the related connection speeds. Monitoring must at least include Port & Connectivity availability, round trip time, packet loss and packet drops. Relevant periodic reports must be provided. Finally, proactive network upgrades tailored to HPC traffic need to be accommodated.

- *Security*
  The hyperconnectivity network must include built-in security features such as DDoS detection and mitigation mechanisms, traffic filtering, and intrusion detection and mitigation systems. These features must be provided as an integral part of the service to protect the EuroHPC ecosystem from both incoming and outgoing threats. Traffic filtering will need to be performed at the edges, i.e. at the HPC and data sites.

- *User Support and management*
  One-stop-shop support and management is required, including multilingual support, to streamline operations and ensure seamless service delivery. The selected provider must offer end-to-end support, covering everything from initial deployment and configuration to ongoing management, monitoring, and troubleshooting. Multilingual support is essential to accommodate our diverse workforce and global operations, ensuring assistance is available in multiple languages (at least three) as needed. This integrated approach will enable efficient management of the hyperconnectivity network infrastructure, minimize downtime, and ensure that all service aspects are

handled by a single, accountable provider, thereby enhancing operational efficiency and reliability.

- *Professional services and Bespoke Solutions*
  Although the vast majority of PoIs will require plain IP connectivity, there may be a few cases, such as Destination Earth, where bespoke solutions will be needed. These may include Layer 2 or Layer 2 Virtual Private Networks (VPNs). Relevant professional services must be available for provision.

- **Future oriented analysis**
  The study team also developed an Implementation Roadmap and an Evolution Plan for the EuroHPC hyperconnectivity network. The latter will be supported by two main tenders. The first hyperconnectivity tender is expected to be launched in December 2024, with an indicative budget of 60M EUR. The second tender is projected to be launched towards the end of 2027, with an indicative budget of 100M EUR.

  Based on the technical specifications documented the implementation roadmap presents several scenarios for the possible implementation and evolution of the network. The service from the first tender is estimated to run from the second half of 2025 to the second half of 2029 (approximately four years). The second tender service is expected to start in 2028 or 2029 and continue until the end of 2032 (or beyond).

  Two main scenarios for the transition between the two tenders are analysed:
  - *Scenario 1: The second tender follows immediately after the end of the first, and*
  - *Scenario 2: There is an overlap in service offerings from the two tenders.*

  The choice between these scenarios will depend on the legal and financial modalities applicable to CEF funding, the related tender call(s), and the preference of the EuroHPC JU and its Governance Board. Another important parameter is whether the second tender will be awarded to the same service provider or not, and whether it will expand the service to cover a higher number of end sites. Currently, the dominant scenario assumes a slight overlap between the two tenders. It is also anticipated that the second tender will expand service to a greater number of end sites.

  -

2.3. Recommendations

- Solution A, is the optimal solution, i.e. leveraging GÉANT and the NRENs to establish the hyperconnectivity network, as was originally outlined in the tender specifications.
- The number of end sites to be connected is a crucial aspect of the two hyperconnectivity tenders. While the EuroHPC JU prioritises the EuroHPC Hosting Entities, it is important to recognise that the ecosystem also includes a range of other European and national HPC and data centres.

  Key flagship initiatives, such as Destination Earth, depend on the efficient transfer of data across this ecosystem and to/from the EuroHPC sites. This may also involve commercial cloud providers, necessitating their direct interconnection along with a sufficient number of Internet Exchange Points. As more Research Infrastructures and Data Spaces become operational over time, they will feed data into HPC centres for potential simulations and calculations.

Therefore, integrating as many end sites as possible is critical for the long-term success of the hyperconnectivity network, supporting the EuroHPC strategy of creating a federated and hyperconnected HPC ecosystem.

- Several risks and KPIs must be proactively monitored to enforce mitigation measures and anticipate the evolution of the network over time. Proactive monitoring will provide insights into the network's trajectory, such as the number of sites being connected and the challenges faced by sites unable to connect.

  This information will determine whether a surplus of budget might be available in the next tender and whether additional services for end sites can be accommodated (such as the purchase of border routers/cards or additional equipment/configurations, including Science DMZs).

- It is particularly important to note that, in the first tender, the hyperconnectivity demarcation point was set at the end site's border router, with the cost of upgrades beyond this point falling outside the scope of the first tender. This creates an opportunity for the second tender to be optimally designed to complement the first, requesting also such additional resources, especially for the EuroHPC sites, thereby maximizing the use of the network end-to-end and also impact of the hyperconnectivity effort.

In the following sections we present further details on the different phases and outputs of the study.

2.4. Methodology

To comprehensively understand the European HPC landscape and its connectivity needs, the EuroHyPerCon team used a multi-faceted approach involving workshops, focus groups, and stakeholder interviews. The data collection effort engaged 558 key participants representing a wide array of stakeholders, such as HPC providers, data providers, AI developers, and network operators across Europe. This stakeholder base comprises prominent entities from the EuroHPC hosting sites, national HPC providers, and leading AI research centres, ensuring a comprehensive insight from major players in the field.

A targeted questionnaire was distributed to gather quantitative and qualitative insights from participants, with over 165 full responses received. This survey captured essential data on current connectivity requirements, network satisfaction, and specific bottlenecks hindering data transfer capabilities. The feedback from these responses enabled EuroHyPerCon to map out the needs and limitations within the existing infrastructure and propose targeted solutions. The following graph provides a breakdown of the stakeholder responses by category.

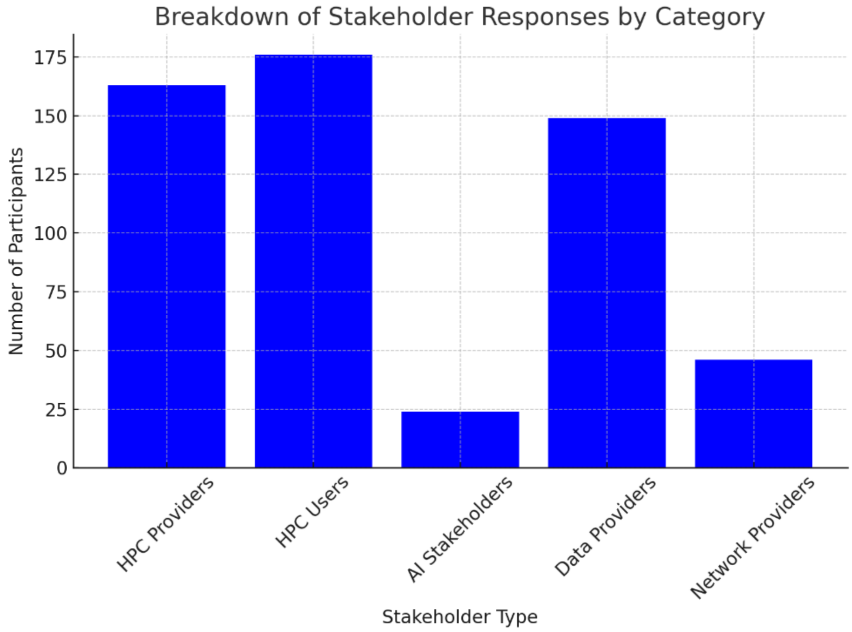

Figure 2. Breakdown of the stakeholders of the study.

2.5. Findings on Current HPC Connectivity Landscape

The EuroHyPerCon study revealed that while stakeholders are generally satisfied with the services provided by GÉANT and NRENs,[2] significant gaps remain in terms of access speed, reliability, and network security. User satisfaction is usually high, but certain critical bottlenecks exist, particularly in access networks and data transfer speeds. The study's analysis of the current landscape shows that while many HPC centres can support speeds up to 200 Gbps, the exponential increase in data demand by 2032 will require significant upgrades in infrastructure.

Furthermore, traffic growth projections indicate a fivefold increase in demand, largely driven by data-intensive applications like climate modelling, AI training, and further scientific research. The study also identified that most of the demand comes from HPC and HTC[3] applications, with AI users making up approximately 10% of traffic. This highlights the need for a federated network that can scale to meet diverse user demands while maintaining high security standards and data integrity.

The following graph shows the expected HPC-related traffic between 2025 and 2032. By 2032 it is expected to increase more than 25 times compared to relevant traffic during 2024.

[2] NRENs - About GÉANT
[3] HTC, short for high-throughput-computing

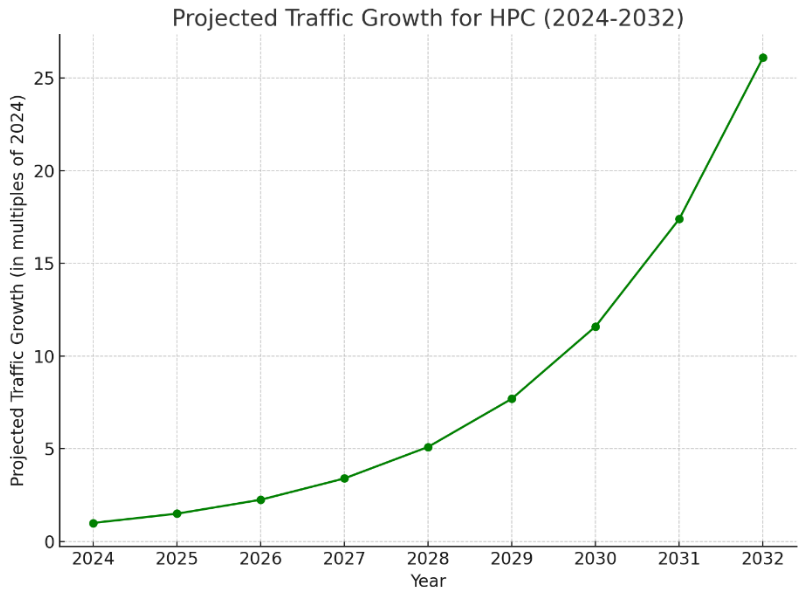

Related connectivity requirements have been organized into 4 classes. The following table provides a summary of the relevant classes, access speeds and their evolution within time.

Tier	Current (2023-2024) connectivity	Phase A (2025) connectivity	Phase B (2028) connectivity	Phase C (2030+) connectivity
Class A	200-400Gbps	400-800 Gbps	800-1.2 Tbps	1.2-1.6 Tbps
Class B	100-200Gbps	200-400 Gbps	400-800 Gbps	800-1.2 Tbps
Class C	10-100Gbps	100-200 Gbps	200-400 Gbps	400-800 Gbps
Class D	1-10 Gbps	10-100 Gbps	100-200 Gbps	200-400 Gbps

2.6. Proposed Solutions and Evaluation

Four potential solutions were evaluated to establish a hyperconnected HPC infrastructure.

- ***Solution A***, which involves upgrading the existing GÉANT and NREN network infrastructure, emerged as the most favourable option. This solution leverages the robust network foundation already in place, supporting massive data transfers and high-speed connectivity.
- ***Solution B*** proposes an IP transit service through commercial providers, but this approach lacks the scalability and network performance required for HPC-specific traffic patterns.
- ***Solution C*** suggests building a new network from scratch, which, although highly customisable, would be costly and resource-intensive, making it less feasible.
- The final option, ***Solution D***, is a hybrid approach that combines elements from solutions A, B, and C. While this could potentially meet the diverse needs of the HPC ecosystem, it would

involve complex management and implementation, making it less efficient than a direct upgrade to the GÉANT and NREN networks.

Assessment Analysis Criteria	Solution A	Solution B	Solution C	Solution D
Fitness for purpose				
Impact				
Performance				
Security				
Support for Innovation				
Sustainability				
Cost Efficiency				

Ultimately, Solution A was recommended as it aligns with the existing infrastructure and offers a fast, cost-effective pathway to hyperconnectivity.

The graph above provides a comprehensive summary of the Assessment analysis of the different scenarios examined. Green indicates full coverage of the criterion, yellow partial coverage, and red not good coverage.

2.7. Financial Analysis

The study includes a comprehensive financial analysis comparing the four proposed solutions, based on recent market research and pricing benchmarks. ***Solution A***, involving the GÉANT and NREN infrastructure, is estimated at 145 million EUR, with 61 million EUR for Phase 1 and 84 million EUR for Phase 2.

In comparison, ***Solution B*** — using commercial IP transit services — would cost approximately 635 million EUR, while ***Solution C***, building a new network from scratch, would exceed 1 billion EUR. The hybrid ***Solution D*** offers flexibility but at a cost of around 390 million EUR.

The financial analysis reinforces the selection of Solution A as the most cost-effective approach. This option not only utilizes existing infrastructure but also minimizes operational costs over the long term, allowing EuroHPC to achieve its goals without compromising on quality or scalability.

The following graph allows an easy comparison of the costs of the different scenarios.

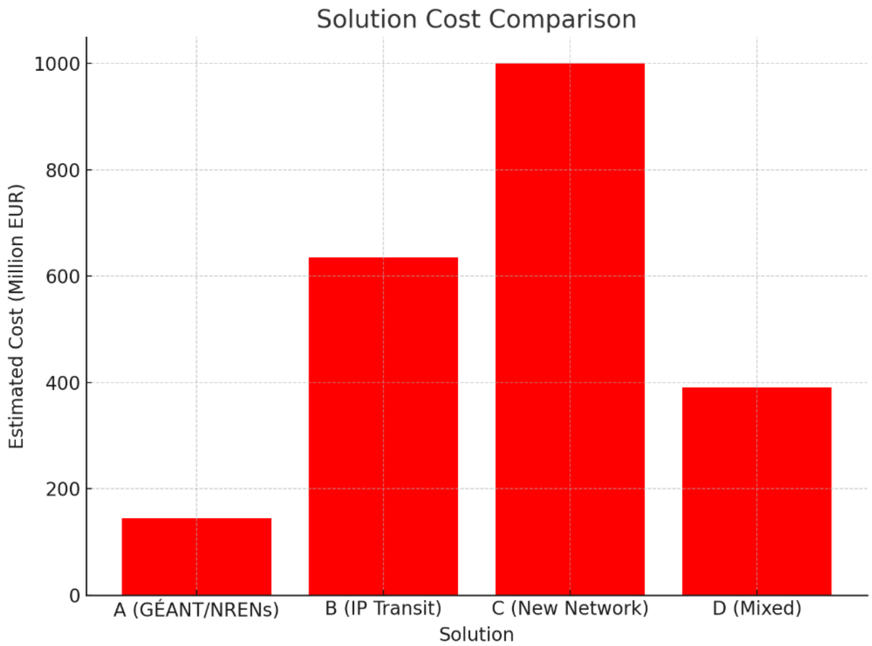

2.8. Implementation Roadmap

The implementation of EuroHyPerCon's hyperconnectivity service will be divided into two primary phases.

- ***Phase 1***, spanning from 2025 to 2029, will focus on initial infrastructure investments and upgrades across priority HPC sites, with a budget allocation of 60 million EUR. This phase will lay the groundwork by establishing high-speed connectivity and building redundancy at critical points of presence (PoIs) across Europe.
- ***Phase 2***, planned from 2029 to 2032, aims to scale the network to accommodate the anticipated increase in traffic, with an additional 100 million EUR budgeted for this purpose.

The roadmap also establishes key performance benchmarks, including a 60% utilisation threshold for link capacity. When traffic consistently exceeds this limit, proactive network upgrades will be triggered to ensure continued reliability and avoid bottlenecks. This phased approach is designed to support both immediate needs and long-term growth while optimising resource allocation.

The two phases are summarised in the following table:

Phase	Timeline	Budget Allocation (Million EUR)	Key Objectives
Phase 1	2025-2029	60	Initial infrastructure upgrades at key sites
Phase 2	2029-2032	100	Expansion of capacity to meet growing demand

2.9. Technical Specifications and Network Architecture

The recommended solution centres on a ***Network-as-a-Service (NaaS)*** model that will support the extensive data transfer needs of the HPC community. This architecture focuses on core requirements such as high access speeds, robust backbone connectivity, and a range of access classes for different site profiles. The network will implement so-called "jumbo" frames[4] to facilitate efficient data transfer and support HPC-specific traffic patterns, while a performance monitoring platform will ensure continuous service quality.

Key technical requirements also include traffic filtering, DDoS[5] protection, and high-availability profiles for critical access points. Security is a paramount concern, and the proposed infrastructure will employ advanced security measures, including redundant pathways for resiliency. Service Level Agreements (SLAs) will specify standards for uptime, latency, and packet loss to ensure that the infrastructure meets the needs of HPC users and data providers across Europe.

2.10. Procurement and Regulatory Considerations

To implement the hyperconnectivity solution, EuroHyPerCon has examined multiple procurement procedures, with a focus on ensuring compliance with EU regulations. Two main options are under consideration: an ***open tender*** process, which allows for competitive bidding, and a ***negotiated procedure without prior publication***. The latter, which is permitted under Directive (EU) 2018/1972 for communication services, offers greater flexibility to negotiate terms and adapt to evolving needs in the EuroHPC network.

The open tender process has certain limitations, including lengthy timelines and restricted flexibility. By contrast, the negotiated procedure would allow direct discussions with the provider to finalize requirements, which is particularly valuable given the evolving demands in areas like AI. This procedure also supports faster deployment and better integration with existing GÉANT/NREN infrastructure, positioning it as the preferred approach for achieving the hyperconnectivity goals.

[4] What are jumbo frames? | Definition from TechTarget

[5] DDoS, short for distributed denial of service

2.11. Expected Impact on EuroHPC JU Strategy

The implementation of a hyperconnected HPC network will significantly enhance EuroHPC's mission of establishing a federated HPC and Quantum ecosystem across Europe. This infrastructure is expected to foster collaboration and drive advancements in critical research areas, including climate science, health, and AI, by connecting leading research centres and enabling seamless data sharing. By creating a network that can accommodate emerging technologies and diverse user groups, the EuroHyPerCon initiative positions Europe as a global leader in advanced research infrastructure.

Beyond scientific research, the project will also support economic growth and innovation across sectors by facilitating access to HPC resources. This infrastructure will provide SMEs, academic institutions, and research entities with the tools they need to remain competitive in a data-driven global economy.

2.12. Conclusions

In conclusion, the EuroHyPerCon study, based on the analysis performed, recommends Solution A as the most viable and efficient pathway for establishing a hyperconnected European HPC infrastructure. By leveraging the GÉANT and NREN networks, EuroHPC can implement a scalable, resilient, and cost-effective solution that aligns with its strategic objectives. The recommended next steps include initiating the tender process, finalising the list of prioritised sites, and progressing towards the implementation deadline of 2025.

As a final remark, the EuroHyPerCon study highlights the critical importance of a timely and coordinated effort to develop a robust and sustainable network infrastructure, adapted to the HPC needs. In other words, this is an excellent opportunity for HPC providers, data providers, HPC users and network providers to acquire networking services and build a hyperconnectivity network optimised for the HPC ecosystem. This initiative not only addresses immediate connectivity challenges, but also strengthens Europe's position as a global leader in scientific research and innovation for decades to come.

3. Stakeholders Consultation

The goal of the Stakeholders Consultation was to identify and analyse the connectivity requirements of the EU and national High-Performance Computing (HPC) systems in Europe and to subsequently specify a future-proof connectivity service, along with its implementation roadmap. To complete this rather complex task, it is of utter importance to identify, involve, and consult with the relevant stakeholders.

An extensive list of stakeholders has been created to perform an exhaustive analysis, gathering as much feedback from as many stakeholders as possible. In total, **558 stakeholders** relevant to HPC-related initiatives (users and providers) have been identified. Out of the 558 stakeholders, **163 are HPC Providers** (from the biggest EU systems down to smaller national ones), **176 are HPC Users**, **46 are Network Providers,** and **173 are** "**Other important stakeholders**", **24** of which have been identified as **Artificial Intelligence-related users,** and **149 are data providers**.

A comprehensive and meticulous approach was employed to ensure the identification of all key stakeholders and relevant categories, leaving no stone unturned in the process. It should be noted that information on Common European Data Spaces has been included later in the study, as this information became available only at that stage. All identified stakeholders were invited to contribute. Still, several of them replied that they don't have specific HPC needs or cannot estimate their needs, while several others did not provide an answer.

A corresponding database was populated including detailed information per stakeholder, depending on the type, i.e. for HPC systems, more information is collected, including geolocation information. Although 558 stakeholders have been identified; the database consists of **742 entries**. This is because HPC providers can appear in the list with multiple HPC systems. Within the database, there are **163 HPC providers with a total of 346 HPC systems**, which may be spread across different geographical locations. In addition, contact information may vary between HPC Systems. These aspects were taken into account during the database compilation, so that a new system of an HPC provider is documented as a separate entry.

Overall, the **number and geographical distribution of the stakeholders is considered sufficient**. This is valid for all identified categories, namely HPC users, HPC providers and Network Providers. As expected, the numbers vary across countries; some countries have many HPC users declared as stakeholders, while others less. This is normal, not only because there are differences in the actual numbers in each country, but also, because in some countries there was better dissemination of the relevant efforts (and this is more relevant in this case). For example, the National Competence Centres (NCCs) of some countries in the EuroCC2/Castiel projects[6] may have done better dissemination within their countries. Notably, several communication channels were used to approach potential stakeholders.

[6] https://www.eurocc-access.eu/ & https://www.linkedin.com/company/castiel-project

The identified stakeholders were then invited to participate in the study and played a vital role throughout the whole consultation process to ensure maximum involvement. The main tools used in this process were a series of workshops with a high number of attendees, multiple focus group meetings going deeper into the details with subgroups of important stakeholders, interviews, and well-thought questionnaires for each of the identified stakeholder groups, namely HPC users, HPC providers and Network Providers. Privacy and compliance with the General Data Protection Regulation (GDPR)[7] were taken into account during the whole process, including the registration of the stakeholders (via both the registration form[8] and the ConceptBoard[9] environment), the workshops registrations and the LimeSurvey tool that was used for the surveys.

3.1. Objectives of Stakeholder Consultation

To provide the best possible recommendations and solutions for future-proof connectivity throughout EuroHPC/European ecosystem, it is vital to understand the current situation of the network, as well as all current and possible future requirements like those of big data users or AI ones. The only way to fully understand this is by asking all those involved – HPC providers, HPC users, Network Providers, data providers and other experts. Combining the knowledge, requirements, and future plans of a broad set of stakeholders ensures the best possible chances to foresee possible requirements and problems. The bigger the set of stakeholders involved, the higher the chances that both current and future needs will be considered and that extreme cases will be properly addressed.

The study has focused on having as many representatives as possible from the groups involved in the information-gathering process. This was successfully achieved through various means, including not only the already established connections of the project partners, but also existing communication channels like EuroCC2, CASTIEL2, GÉANT, FF4EuroHPC[10,11] etc. who acted as multipliers. On top of this, a database of stakeholders was created through combining individual research with filtering through existing public lists that already contained some relevant stakeholders, like the TOP500 list.

3.2. Defining Stakeholders

To generate any relevant findings, it was important to first define the list of relevant stakeholders that can give viable input to the study questions. In our definition for this study, stakeholders are groups of entities with a significant interest in how the future of the hyperconnectivity landscape will evolve. While there are certainly many individuals as well as groups and entities that fit that definition, emphasis was given into the most important ones, especially those with major connectivity

[7] https://eur-lex.europa.eu/eli/reg/2016/679/oj
[8] https://edocs.hlrs.de/nextcloud/apps/forms/s/K8Xn5B27pGtLzJmo95YiXfQg
[9] https://conceptboard.com/
[10] https://www.eurocc-access.eu/, https://geant.org/, https://www.ff4eurohpc.eu/
[11] Through HLRS' involvement with EuroCC 2, CASTIEL 2 and FF4EuroHPC, a collaboration was made to distribute several newsletters to all of their relevant stakeholders, making them aware of the study, the workshops, the surveys, and everything of importance. GÉANT also used its channels to disseminate the study to its members.

requirements or with extensive experience that can also provide the best foresight and see some of the trend patterns that might define the next decade of connectivity earlier than most.

3.3. Stakeholder Categorization

In accordance with the tender specifications, the stakeholder groups relevant to the study are clustered in four (4) main groups, namely HPC users, HPC providers, Network Providers, and other important stakeholders, namely AI users and big data providers. In addition, special HPC users of particular attention are the Artificial Intelligence (AI) ones along with the (big) data providers. This is why the latter two (AI and big data providers) are also flagged separately as a sub-category. The proposed categorisation of stakeholders is thus outlined as follows:

- HPC Users
- HPC Providers
- Network Providers[12]
- Other Important Stakeholders: AI Users and big data providers

While it is common to divide stakeholders into primary (those most directly relevant), secondary (indirectly relevant) and tertiary (least relevant) stakeholders, the study team felt that all of these key groups are relevant in major ways by the outcome of the study or at least the results of future plans and implementation of the network. It was important to hear from all of them to fully understand their current and future challenges and requirements, so as to come up with the best possible solutions on how to go forward with the network. Still, as identified in the tender specs, connectivity requirements from HPC users will be aggregated for the relevant segments of the network to get an estimate of the required network capacity. Thus, HPC users are given appropriate emphasis.

Since these groups are very heterogeneous within themselves, it was important to better understand the different requirements and needs, even within their stakeholder group. For this reason, different questionnaires were prepared for each of the groups, catering to the specific needs and nuances within them. With respect to an optimal solution (leaning on multivariate optimization theory), it is our goal that the offered solution will comply with a number of requirements. The main requirements will be around connectivity performance, including the access and backbone connectivity speeds (bandwidth) and cost efficiency, while other ones will be taken into account, such as other performance and technical characteristics (guaranteed throughput, latency, security) as well as non-technical ones (fitness for purpose, impact, sustainability). The aim is to come up with an optimal solution and this will be derived following Analysis of Alternative Solutions.

For HPC Users, it was important to get a good understanding of what the regular use cases, requirements and processes look like ("User Journeys"), which was derived from discussions with both HPC Users and HPC Providers alike. As the latter group deals with multiple users, they have a good

[12] Notably, network equipment vendors are not included, because network providers have the freedom of selection. Basically, focus was given on basic/standards-based ISP's transport IP technology at top tier speeds (i.e. >=400 Gbps).

understanding of the current usage and processes of their users. In contrast, the users can better foresee where their needs and processes will go in the future, especially thinking of how they might use AI or other upcoming technologies for their purposes.

Any participant of those groups was welcomed and included, and only some distinctions were made in some groups so that the study team can better prioritise their inputs, as explained in the power-interest matrix (Figure 3). For HPC users, high-end cases, such as Destination Earth (DestinE), have been included i.e. those requiring very high amounts of bandwidth for moving data for their HPC needs; and in some cases where the current network and infrastructure available are either already a bottleneck or where the users themselves already anticipate them to be one in the future. While these are a very small percentage of the total user base, they still need to be properly considered and can have quite different requirements than the average user requiring bespoke solutions or at least dedicated attention. For the purpose of stakeholder mapping, these users are called “power users”, while regular users are presented as “regular users”. For HPC users, the relevant area they belong was also identified (e.g. Research, AI, Engineering, etc.), and also used this as a column in the database.

Another important group of HPC users are those who already actively use AI in their calculations, like AI weather forecast systems, which require training on vast historical datasets (e.g. with data from the last 80 years), or Large Language Models for services similar to ChatGPT. Furthermore, the new AI Factories initiative will serve as a one-stop shop for start-ups, offering computing power combined with access to data centres, access to talent through academic partnerships and additional services. Overall, the European Commission and Member States will invest an estimated 2.1 billion Euro in AI Factories and AI-dedicated supercomputers until 2027[13]. Training in AI would also be needed. These use cases are growing rapidly, and it is very common to deal with large data sets requiring a high amount of HPC resources and corresponding connectivity. With both of these requirements and the importance of AI going forward, it was essential to consult with stakeholders in this field and pay attention to how their requirements might look in the future.

For HPC providers, a distinction could be made in several ways. For example, in the past, PRACE used a tier-based structure to distinguish between 10 Tier-0 systems with a peak of up to 77 PF/s and 12 Tier-1 systems that allow a peak of up to 16,21 PF/s [PRACE][14]. Another distinction could be made by dividing systems by funding source, ranging from EuroHPC JU co-funded (which range from mid-range and peta- to pre-exa- to exascale) to nationally or institutionally funded, with the reality often being mixed funding from European, national, and regional/federal levels. Another option would have been to cluster the whole European and national landscape by Petaflops and/or other metrics and draw arbitrary lines in between. It is important though to focus both on all the big EU centres, including the EuroHPC JU hosting sites, and then on the smaller national ones. Institutional ones (e.g. university HPC installations) are outside the direct scope of the study, as the connectivity requirements are in the vast majority of the cases within the corresponding (e.g. university) campus LANs, thus not contributing to

[13] AI Factories | Shaping Europe’s digital future
[14] PRACE HPC Infrastructure - PRACE (prace-ri.eu)

the EU hyperconnectivity requirements. Despite that, some answers to the questionnaires for institutional needs were provided.

With respect to Network Providers, two stakeholder groups have been identified, the research ones, namely GÉANT and the EU NRENs (who are the owners and stakeholders in GÉANT) and commercial Network Providers. Given the long investments by the NRENs and GÉANT over several decades, the tender specs highlight the importance to leverage these stakeholders in the study. Furthermore, GÉANT and the NRENs rely on fibres and wavelengths (i.e. lease) from commercial providers. So, the two groups are closely linked to each other. GÉANT and the EU NRENs have a high interest in the future of the hyper-connectivity services and were very important to include in our stakeholder consultation process. Information from commercial providers, directly or indirectly, are also being sought, about services availability and pricing across Europe.

Another important stakeholder group that was paid significant attention to, were (big) data providers, including the Commons EU Data Spaces[15] and world-class Research Infrastructures[16]. These are an important part of the European Research Infrastructure ecosystem and create huge amounts of scientifically relevant data. These can include big facilities in different research fields such as physical sciences and engineering infrastructures such as accelerators, synchrotrons and laser facilities, along with different types of telescopes, health and food infrastructures such as biobanks, bioimaging, clinical research facilities and structural biology ones, energy infrastructures such as fusion, renewable or other energy reactors, environment-related infrastructures on atmosphere, carbon, river-sea, sea-floor or plate-observatory systems, social sciences and humanities databases, archives, tools and surveys, and finally digital research infrastructures on experimental facilities, social mining and big data analytics, etc. Not only do they use and compute that data for their own research, but they also work towards sharing such data to other researchers from other disciplines. The more this happens, the more often this data needs to be transferred across the network and between HPC sites. With the increased use of AI, and more interdisciplinary research projects, these troves of data could become more and more important, which makes it necessary to include them in this study. On top of this, even though many data providers are currently only using High-Throughput Computing, within the data provider entities themselves, the use and demand for HPC, especially in the future, seems to be growing and needs to be considered. On the other hand, for the same reason, several such research infrastructures facilities could not provide direct inputs to the study, as they could not estimate their future connectivity requirements relevant to HPC analysis.

3.4. Stakeholder Mapping

After identifying these relevant stakeholder groups and subgroups, the framework of stakeholder mapping[17] was slightly adapted to our requirements and used to first rank them **by interest** (i.e. how interested the stakeholder in the outcomes of the study and the future evolution of the network is)

[15] https://digital-strategy.ec.europa.eu/en/policies/data-spaces
[16] https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe/research-infrastructures_en
[17] https://aisel.aisnet.org/icis1981/20/

and influence (here defined as the ability to influence the future of the network as well as other stakeholders), and then further arrange them into a matrix.

This matrix ultimately not only provides a better understanding of how interested and affected they are, but also gives information on how our communication strategy must be adapted. Figure 3 shows the stakeholder groups placed in the **influence-interest matrix**. The four quadrants correspond to a subsection of the matrix and recommend how to weigh and communicate with each set of stakeholders, given non-unlimited resources. The quadrants and the recommended behaviour towards the stakeholders placed within them is relatively straightforward and ranges from "monitor (loosely)" to "manage closely", with the two quadrants of "keep satisfied" and "keep informed" somewhere in between this spectrum (see Figure 3). While the framework is useful in many regards and it was used as a baseline, attention was paid to all the stakeholders during the consultation process.

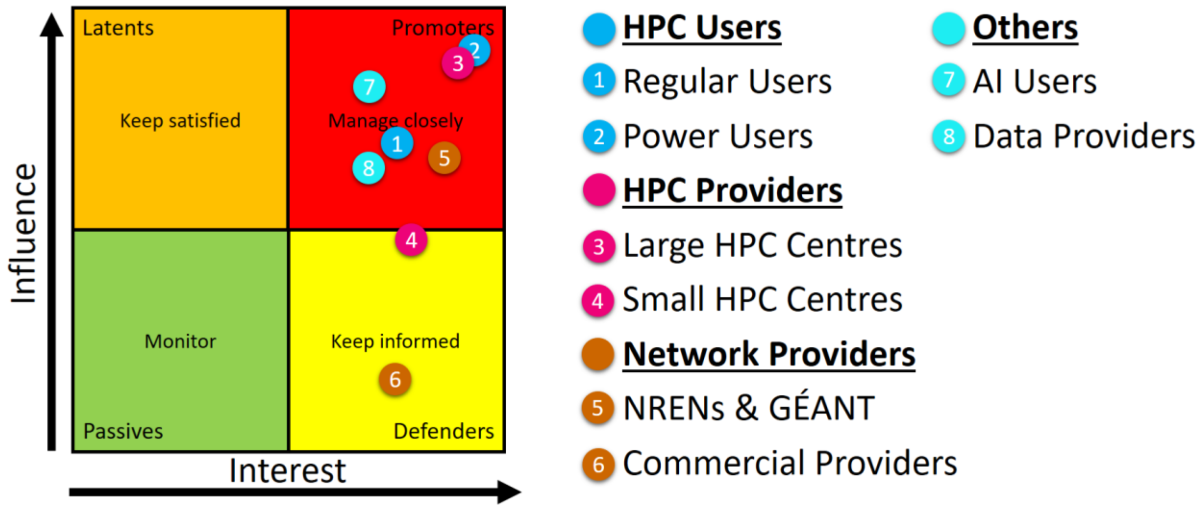

Figure 3. Stakeholder Mapping with power-interest matrix of stakeholders

The identified stakeholder subgroups relevant for the matrix are:

- HPC Users:
  - Regular Users (Medium-High Interest, Medium-High Influence)
  - Power Users (High Interest, High Influence)
- HPC Providers:
  - Large HPC Centres (High Interest, High Influence)
  - Smaller HPC Centres (Medium High Interest, Medium Low Influence)
- Network Providers:
  - NRENs & GÉANT (High Interest, Medium-High Influence)
  - Commercial Network Providers (Medium-High Interest, Low Influence)
- Other Important Stakeholders:
  - AI Users (Medium-High Interest, High Influence)
  - Data Providers (Medium-High Interest, Medium-High Influence)

3.5. Planning and Executing the Consultation Process

Our methodology was developed focusing on how to increase the amount and the quality of the feedback received by the stakeholders, allowing an extended and complete collection of relevant data. The consultation plan was divided into three phases and incorporated everything from creating a stakeholder database, a set of "User Journeys" (which describe different usage scenarios at the HPC centres), a set of workshops for each of the stakeholder groups, a corresponding questionnaire for each of the stakeholder groups, as well as further individual feedback instruments like one-on-one interviews with power users or focus groups with several of the larger stakeholders.

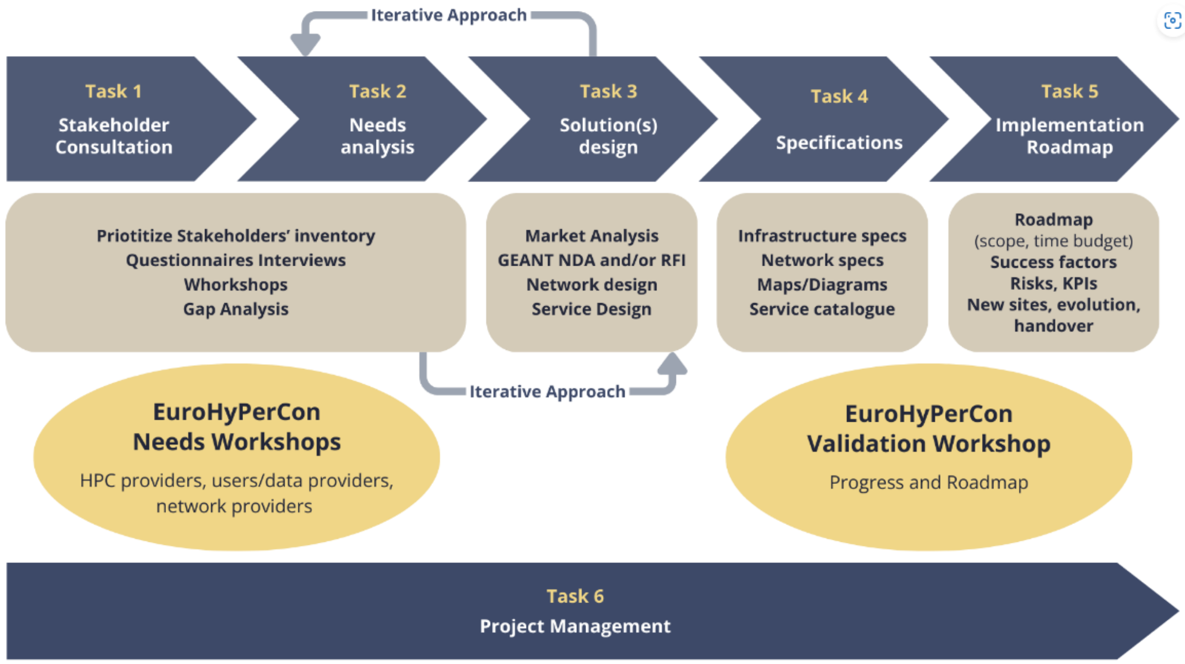

Figure 4. Methodology of the Study

The three phases for the Stakeholder Consultation were initially defined as follows:

- **Phase 1**
  - Compile an initial list of relevant stakeholders
  - Select representative and geographically balanced stakeholders
  - Address additional stakeholders through an open registration form
- **Phase 2**
  - Generate individual user journeys for those stakeholder groups
  - Compile questionnaires for HPC providers, HPC users, and network providers
  - Organise a series of workshops with the identified stakeholders
- **Phase 3**
  - Set up focus groups
  - Conduct bilateral interviews

The plan already defined a preliminary timeline as well as ideas for the workshops with the individual stakeholder groups.

On top of that, several conferences and events were selected and later attended, with the plan to first identify, reach and consult with additional stakeholders, and later to validate the plans and findings. Among them were:

- ***Quantum Industry Summit***, Stuttgart, 10-11.10.2023
- ***PRACE Workshop – Kick-off of the Scientific and Innovation Case 2024 Edition***, 20.10.2023, HLRS
- ***European Big Data Value Forum***, Valencia, 25-27.10.2023
- ***Supercomputing 2023***, Denver, 12.-17.11.2023
- ***HiPEAC conference***, Munich, 17-19.01.2024
- ***EuroHPC Summit***, Antwerp, 18-21.03.2024
- ***ISC High Performance conference 2024***, Hamburg, 12-16.05.2024 (Pending)

The plan was followed strictly and with only minor deviations, which led to the successful execution of a number of milestones:

- Workshops – All well attended with 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
- Filled in questionnaire & interview with ***Destination Earth/ECMWF*** – 18 December 2023
- Focus group with Exascale & Pre-exascale research Network Providers – 18 December 2023
- Meeting with EC (DG CNECT), EuroHPC JU and ***GÉANT***, Amsterdam, 11 January 2024
- Focus group with Exascale & Pre-Exascale HPC providers, 5 February 2024
- Follow-up interviews with the DestinE consortium, including ***ECMWF***, ***ESA*** and ***EUMETSAT*** - 9 February 2023
- Follow-up interview with ***EUMETSAT*** on the involvement of their data (and connectivity) provider (***CloudFerro***), 1 March 2023
- More than 150 full questionnaires were completed among all categories (HPC users, HPC providers and Network Providers), while more than 400 partial questionnaires have been completed (although there are several duplicates and dummy answers).

3.6. Stakeholder Database Design, Implementation & Analysis

The development of the database is a carefully coordinated process designed to ensure that the resulting database is seamlessly aligned with the project's specific requirements.

The final database that has been created during the first phase of the project is based on various sources of information and enables the achievement of several key objectives for the project, including:

- identifying the HPC connectivity requirements and needs from the different standpoints
- identifying synergies with other connectivity initiatives or solutions
- assessing the technical feasibility of the solution(s)
- identifying possible improvements to the solution(s)

3.7.1 Database Purpose and Requirements

In order to achieve the above-mentioned objectives, including a comprehensive analysis of the connectivity requirements for EuroHPC systems and other major European and national supercomputing infrastructures, it is essential to establish a database of stakeholders relevant to the EuroHPC and the Hyperconnectivity framework.

A spreadsheet has proven to be a suitable choice for handling small to medium-sized data sets and uncomplicated databases, which is why it was chosen as the database platform. In the area of data management, Excel serves several purposes, including storing and organising data, providing a clear overview through its tabular structure and user-friendly interface, facilitating data analysis through statistical operations, and allowing the creation of reports and visualisations.

However, certain requirements have been carefully considered to ensure the integrity and security of the data. Access to the Excel file is restricted, adding a layer of protection, and routine backups are systematically performed.

Identifying Stakeholders for the Database

An initial set of stakeholders was identified using a solid knowledge of the EuroHPC and e-Infrastructure ecosystem, using contacts and synergies with other projects such as EuroCC2, EXCELLERAT P2, other Centre of Excellence (CoE) contacts, CASTIEL2 and FF4EuroHPC[18]. The following lists and documents were created in this initial stage: Contact list for CoE leaders, Contact list for EuroCC2 Partners, EuroHPC-Centres Tier-0 (Exa and Pre-Exa scale), EuroHPC-Centres Tier-1 (Peta/Mid-range), HPC-Centres Other Tier-1, EuroHPC Contacts (Technical Staff of Hosting Entities / Stakeholders provided by EuroHPC).

The first version of the database, titled 'Organisations', was compiled through further proactive secondary research on the following websites to gather even more stakeholders:

- https://www.eurocc-access.eu/about-us/meet-the-nccs/
- https://eurohpc-ju.europa.eu/index_en
- https://hpc-portal.eu/
- https://prace-ri.eu/
- https://forschungsdaten.info/themen/speichern-und-rechnen/hochleistungsrechnen/

[18] Through HLRS' involvement with EuroCC2, CASTIEL2 and FF4EuroHPC, a collaboration was made to distribute several Newsletters to all of their relevant stakeholders, letting them know about the study, the workshops, the surveys, and everything of importance. GÉANT also used their channels like E-Mails to disseminate the study to their members

- https://gauss-allianz.de/de/about_ga/
- https://www.top500.org/

Additional information and an open form for potential stakeholders to register can be found here (https://eurohypercon.eu/eurohpc-ju-welcomes-eurohypercon/). Thanks to this open form for stakeholder registration and engagement, this process culminated in the creation of the "Contacts" list, which was one of the inputs for the final stakeholder database.

In the stakeholder identification process, an in-depth analysis was conducted on the identified and/or registered stakeholders to assess their involvement's significance and prioritise them as needed. Special emphasis was placed on guaranteeing the representativeness and geographical balance of the stakeholder group, ensuring inclusion from all European countries.

A range of tools and methods were used to develop the final version of the database, including questionnaires, workshops, interviews, and active internet research. Due to the availability of personal data, the full version of the database, *the full database_EuroHyPerCon_02_2024_final*[19]*,* is intended for a limited project group. For public access, an extract of the database has been created, which is presented in the Annex. To give you an overview of the document, an abridged part of the database was taken, which gives a general overview of the important stakeholders.

3.7.2 Stakeholder Database Analysis

The following categories of stakeholders relevant to the study were identified and are marked in the "Organisations" table, column *Stakeholder Group*: *HPC Providers, HPC Users, Network Providers* and *Other Important Stakeholder. Other important stakeholders* include stakeholders that may not fully fit into the above categories, i.e. Data provider and AI users. Further specifications can be found in the *Area* column: e.g. Academia, Industry, Research, AI, and in the *Sector* column: e.g. Energy, IT, Engineering.

[19] https://docs.google.com/spreadsheets/d/1WPM2Oq_q8kXDbWNDT3kN--iJUqk0TIhcqZOqRVrvQMw/edit?usp=sharing

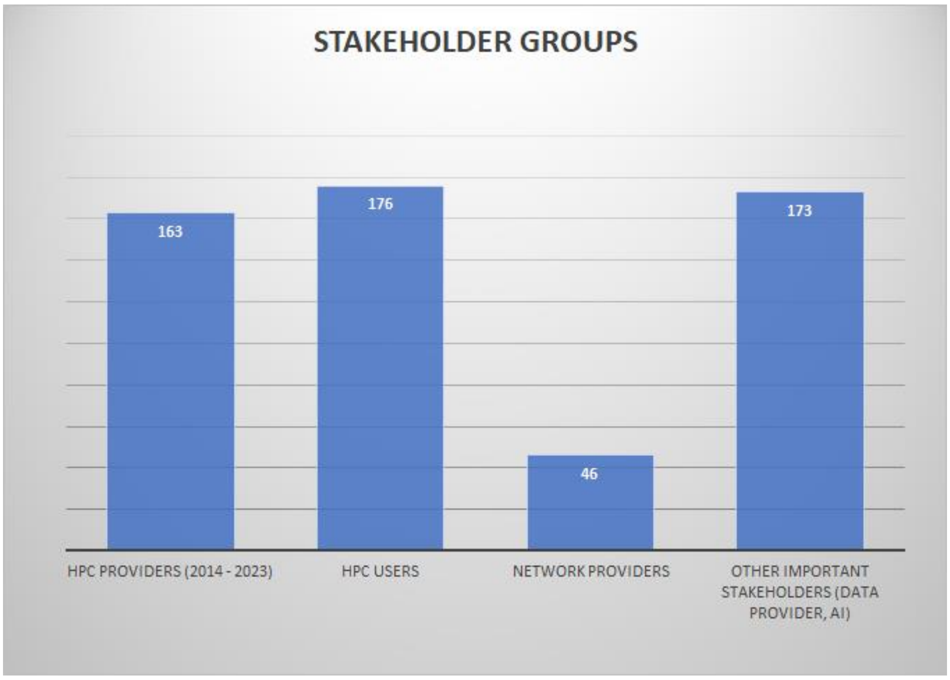

Figure 5. Analysis of Stakeholder Groups, 31.03.2024

The following diagram entitled "Stakeholder groups" provides an overview of the four main stakeholder groups. The diagram shows that all groups are equally represented in terms of numbers. By 31 March 2024, 558 stakeholders had been identified as relevant to EuroHPC and the Hyperconnectivity initiative:

1. **163 HPC Providers** with 346 HPC systems (as each HPC provider has usually more than one system). HPC systems in Europe from 2014 to 2023 were taken into the database. 57 HPC providers (with 118 HPC Systems) come from the industry/vendor sector.
2. **176 HPC Users**, of which 18 HPC Users are from the industry. 158 HPC Users are from education/research.
3. **46 Network Providers**
4. **173 "Other important stakeholders"**, of which **24** are from the AI area, and **149** are data providers.

Some statistics per country and type of stakeholder are presented in the following maps and diagrams. In the next step, the "Stakeholders" diagram shows the distribution of all stakeholders across the European countries in the four focus groups. The number of HPC users, HPC providers and other important stakeholders varies greatly from country to country. The three maps below also show the exact distribution of stakeholders across countries. The maps are structured in such a way that the intensity of the colour indicates the density of the stakeholders. The actual number of stakeholders is also shown in each country. Each category is represented with a different colour.

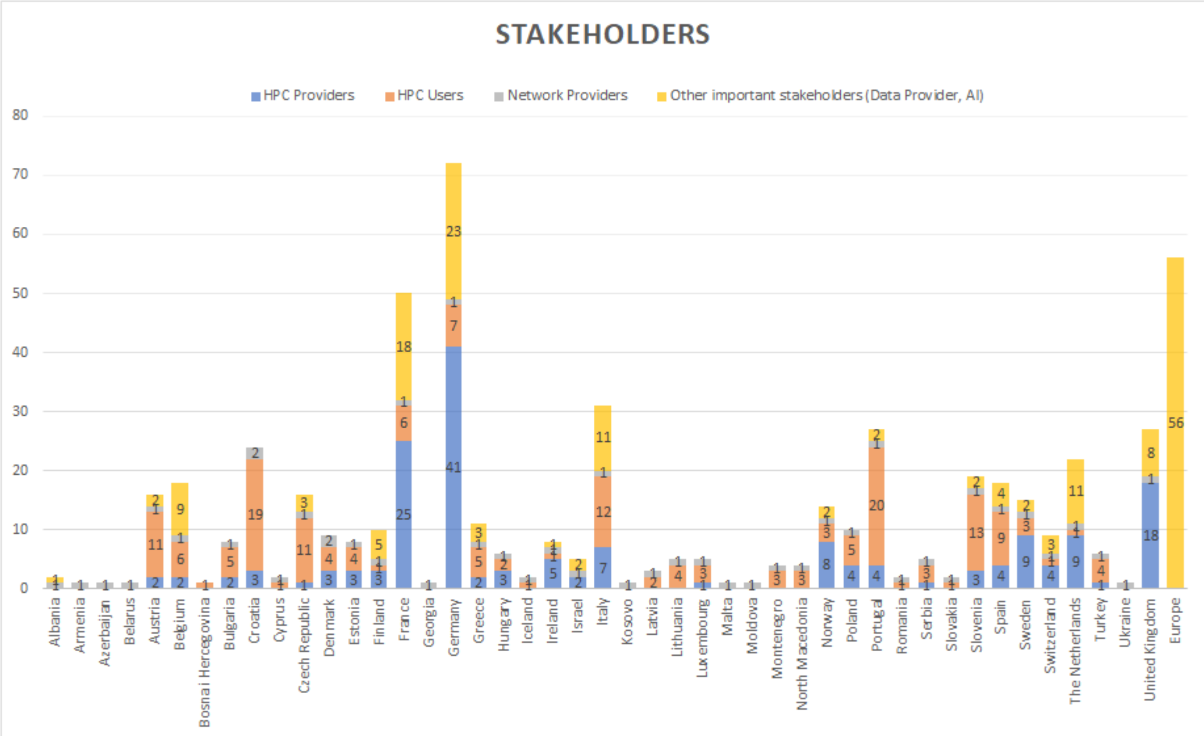

Figure 6. Overview of Stakeholders by Type and Country, 31.03.2024

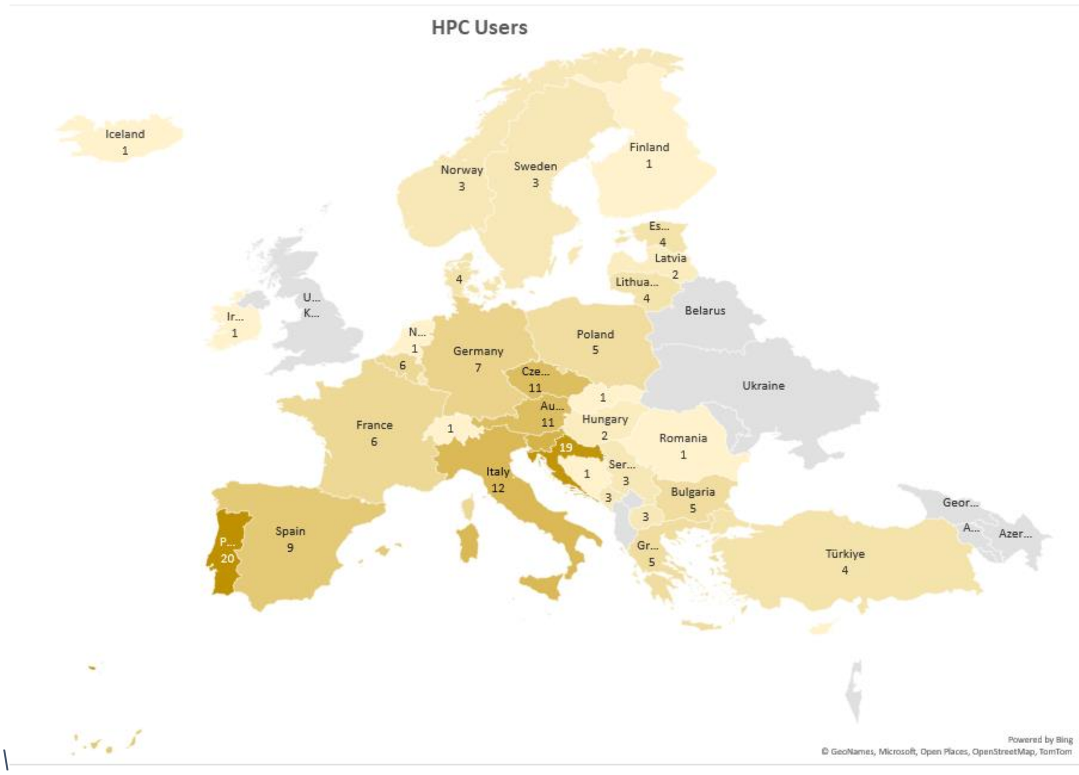

Figure 7. HPC Users density map by country, 31.03.2024

In some countries, there are many HPC users listed as stakeholders; in others, fewer. These figures may possibly be explained by the fact that not all countries were involved in the same way, although all communication channels were activated. In particular, some of the EuroCC countries may have been more active in the dissemination process than others (as this was one main communication channel). The personnel factor may have also played a role. It can be assumed that in some countries there are less personnel dealing with such requests. The most active participation was observed in Portugal, twenty (20) and Croatia nineteen (19), then Slovenia, thirteen (13), Italy, twelve (12), Austria, eleven (11). In the other countries, the presence of stakeholders is lower, but well-representing most countries.

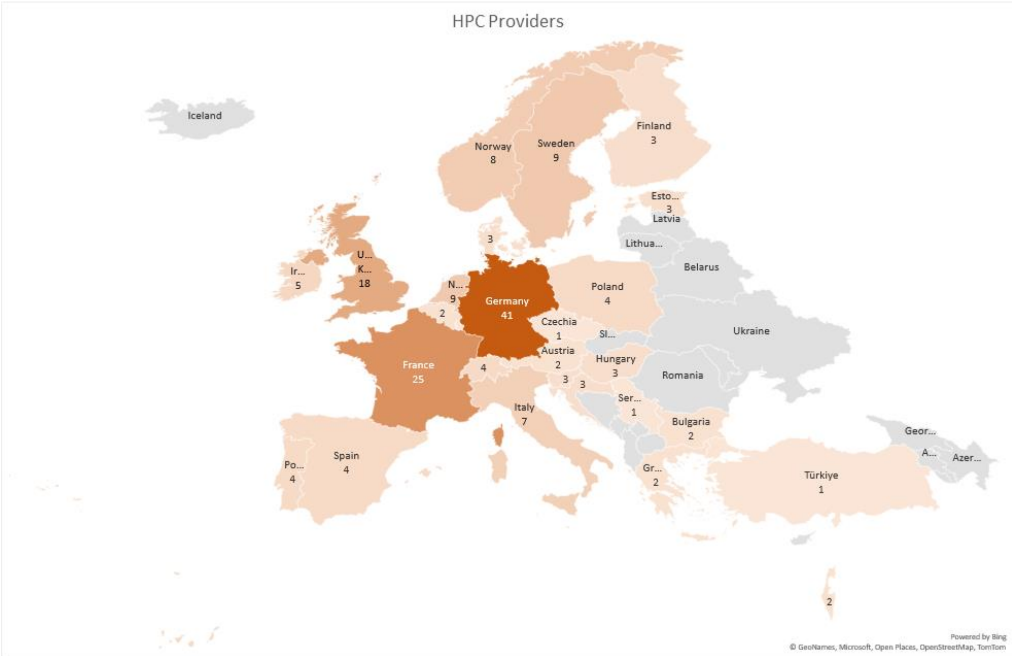

Figure 8. HPC Providers density map by country, 31.3.2024

HPC providers are also not evenly represented, as the map shows. When selecting the HPC providers, time parameters were introduced into the database in order to highlight, in particular, the most powerful HPC systems for further analysis, i.e. the HPC systems installed after 2014. For this reason, several HPC systems from the first phase of the research were removed from the database.

It was later realised that this meant that a number of countries (mainly Eastern Europe, such as Romania and Lithuania, etc.) were not included in the analysis, as the available systems were outdated.

The largest database of HPC providers was compiled for Germany, forty-one (41), France, twenty-five (25) and the United Kingdom eighteen (18), which is expected, given the advanced state of related technology, and also funding for such systems.

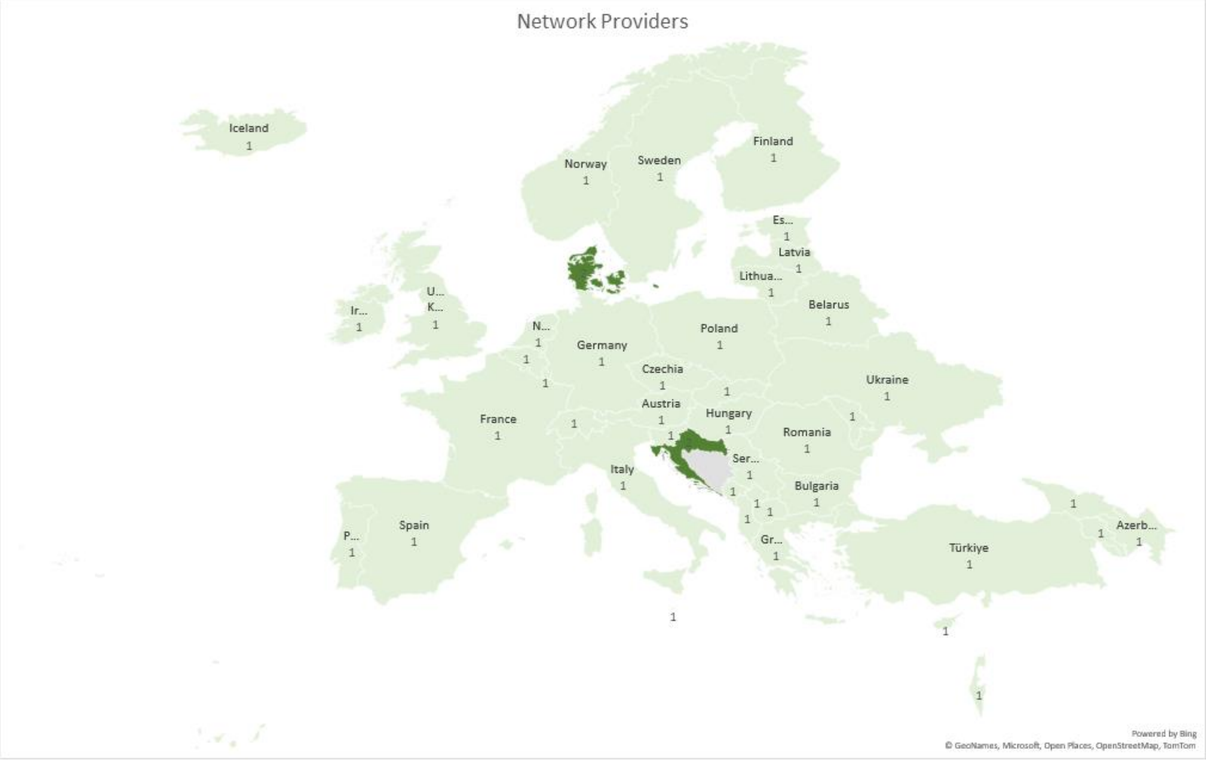

Figure 9. Network Providers density map by country, 31.03.2024

As can be seen from the map, the distribution of network providers is balanced, with the exception of Bosnia and Herzegovina (marked in grey). The Network Providers database was created on the Basis of the National Research and Education Network (NREN) list[20] as these are currently used to interconnect the HPC systems. Given the very short timeframe of the study and the clear direction in the tender technical specifications that the study should leverage the existing GÉANT and NRENs networks, it was decided not to proceed with a consultation with commercial providers. This is further justified by the fact that 80% of the EuroHPC capacity (and also for other systems) and anyhow the GEANT and NRENs are already using fibres and circuits from commercial providers. GÉANT and NRENs networks form a dedicated and advanced Internet infrastructure and service provider to the research and educational communities within a country.

[20] NRENs - About GÉANT (geant.org)

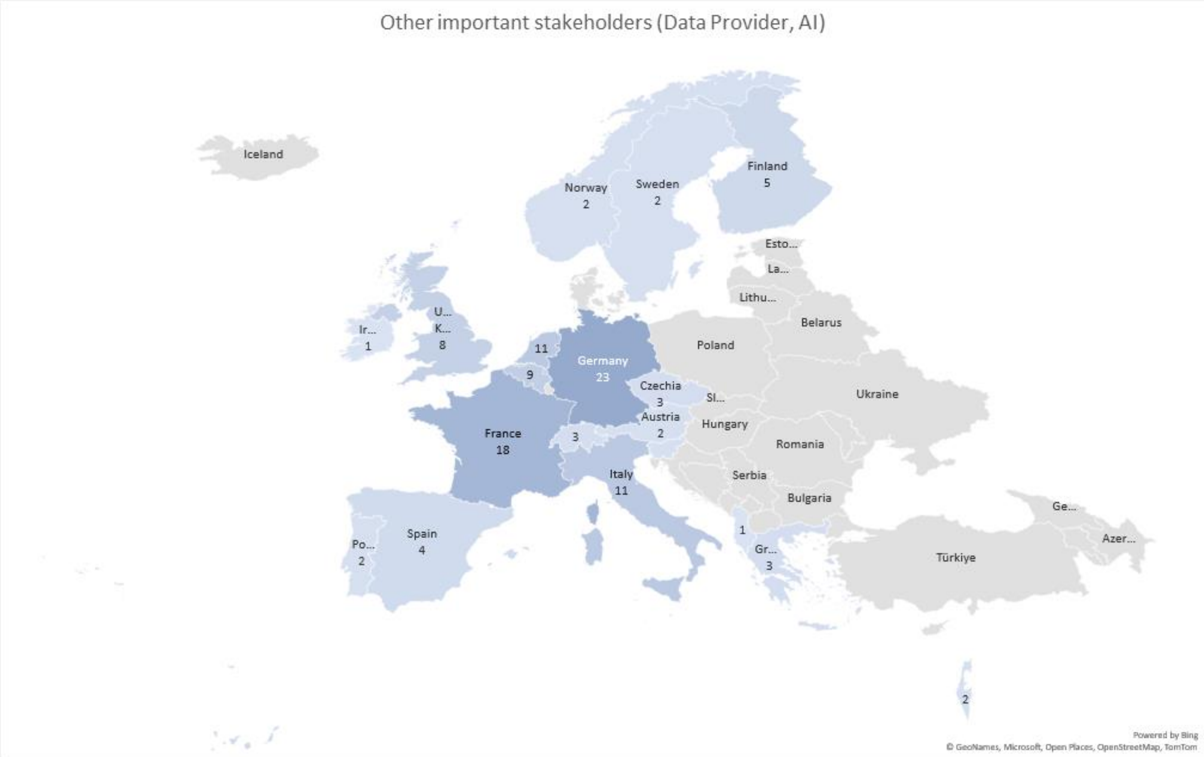

Figure 10 - Other important stakeholders density map by country (includes 56 Pan-European Data Providers), 31.03.2024

The database of "Other Important Stakeholders" includes 24 AI users and 149 data providers representing different sectors such as agriculture, arts, sustainability, energy, environment, etc. The map shows that the distribution of stakeholders is not even. Eastern Europe, in particular, is not well represented in this group, at least according to the map. However, it should be noted that 56 data providers belong to pan-European organisations that cannot be assigned to a country due to their internal organisational structure. The research revealed that many representatives from Eastern Europe belong to these organisations. Furthermore, several of the data providers (such as big ESFRI infrastructures or ERICs) have national nodes in many countries, but these are not depicted on the map. All these stakeholders were summarised in a separate European pillar (Stakeholders diagram). It is not possible to show these stakeholders by country on the map.

3.7.3 Data Privacy and Security Considerations

The data collection for the database design was carried out in accordance with the General Data Protection Regulation (GDPR)[21] on the basis of a privacy statement detailing the processing and use of personal data (e.g. of stakeholder contact persons) used only for the purpose of the study. In the online stakeholder registration form, stakeholders were informed that by submitting their information, they agree that all information provided will be stored in the database (located in the EU) for a specific period of time to be agreed with the EuroHPC Joint Undertaking. The data collected was processed by

[21] Regulation - 2016/679 - EN - GDPR - EUR-Lex (europa.eu)

for the sole objectives of the study contracted to the EuroHPC Joint Undertaking (JU). The data will be kept by the consortium for the duration of the study (which will end in 2024).

In addition, stakeholder contact persons were asked if some (non-sensitive) information (such as full name, organisation and possibly some basic information about their relevance) could be shared with other stakeholders to facilitate communication between them, especially in the context of the workshops.

4. Needs Analysis

Based on the data collected primarily through the questionnaires and supplemented by additional inputs, **several key observations and recommendations** have emerged regarding the needs and future developments for HPC connectivity across Europe: For all three (3) surveys, the HPC users, the HPC providers and the Network providers survey, the input was very satisfactory in both quantity and quality, including geographical distribution across Europe. As expected, there were some countries with more responses than others, which was rather random. In our case, this was influenced by the level of activity among the EuroCC partners (EuroCC and Castiel served as the main communication channels for the questionnaires).

Regarding the HPC users' questionnaires:

- The collected data demonstrate that the survey effectively reached individuals who are active users of HPC services, with a strong representation from both organisational and project representatives, spanning various thematic fields across all sciences. Additionally, there is a substantial sample from AI and Data Science-related fields.
- The questionnaires were most appealing to individuals in academia and the research sector, with a good reach into industry/SMEs, a bit more than ten percent (10%). In addition, several data providers, including Research Infrastructures, Data Spaces, Libraries participated, around ten percent (10%). The participation from academia and research is close to the eighty percent (80%) HPC usage for research that is typical for most HPC systems, including the EuroHPC hosting sites.
- The respondents' roles range from junior researchers and academics to senior managers and projects/infrastructures leaders.
- The majority of the surveyed population is currently closely connected with HPC projects, approximately seventy percent (70%).
- National-level resources are currently the most accessible or preferred for HPC use.
- The expectation of HPC resource usage by 2030 and beyond heavily favours continued national HPC usage, combined with EuroHPC JU-funded resources, both reaching nearly the same level, approximately seventy to seventy-five percent (70-75%). Meanwhile, Regional or Institutional Centres are expected to account for about half of that usage, around forty percent (40%).
- The frequency of High-Performance Computing (HPC) resource usage among respondents indicates a significant preference for high-frequency usage, with nearly fifty percent (50%) of users engaging with these resources on a daily basis, and an additional thirty percent (30%) declaring weekly usage.

- A diverse array of computational needs among users has been identified, with a significant portion requiring tens of thousands to hundreds of thousands of node hours per year. This underscores the critical role of HPC resources in supporting a wide range of scientific and industrial activities. The presence of users across all usage brackets from very low to very high emphasises the varied nature of computational work in different fields and at different scales.
- Users are preparing for growing computational needs, which must be accommodated in the planning and scalability of HPC resources. The questionnaire provides an indication of the potential growth of computational resources in the future, estimated at around sixty percent (60%) for 2025 and approximately one-hundred-fifty-five percent (155%) for 2030 and beyond.
- There is a clear preference for batch processing among users, with the majority finding it fully suitable for their needs. Interactive processing or other types are less common, and typically reserved for more demanding applications. In more detail, the majority of users run "plain" jobs, i.e. upload their data, run their job, and collect their data. Only a small percentage have other types of jobs, such as interactive, i.e. uploading data while the job is running (e.g. data assimilation). There is no clear pattern with regards to the different types of users/usages. One big user using interactive jobs/workflows is Destination Earth, with high requirements, but again, there is no clear pattern on the needs based on different user typologies.
- The majority of users, i.e. more than eighty percent (80%), do not handle sensitive data such as personal health information, which would necessitate stringent data sharing rules. Most users do not have additional security needs beyond those previously addressed. However, security and privacy are critical factors for certain fields, especially those dealing with personal or sensitive data (e.g. health or copyrighted data). In particular, sectors recorded with security requirements are biomedical research (health, genomics, clinical), biodiversity-related, military/export control, and observatory data (from SKA) given their access and security policies. There are also industrial users with security requirements from different sectors, namely agri-food, biomedical, earth observation, and engineering, referring to copyrighted protected material and sensitive data from industry partners and third parties.
- The largest volume of data transfer, i.e. approximately sixty percent (60%) occurs after executing workflows, in particular during the outputs download phase. The most time-consuming phase of HPC jobs is the "Runtime on HPC", i.e. over fifty percent (50%), while other phases (data preparation, upload, download, data analysis) account for the remainder of the time (ranging from seven to eighteen percent).
- The majority of the respondents, i.e. sixty-two percent (62%), did not identify any bottlenecks related to data transfer to HPC centres. However, twenty-eight percent (28%) of respondents reported encountering bottlenecks, such as access bandwidth limitations, data storage issues, or data transfer interruptions, some of which may be attributed to the application software. The first (access bandwidth limitations) is the most frequent, while others are more or less equally reported.

- The respondents perceive the downloading phase (i.e. from the HPC site to the data or user site) to involve a significantly larger portion of data transfer compared to uploading or during loop updates in their HPC jobs[22].
- The majority of users' transfer requirements are less than 100GB per workflow execution; however, there is a significant demand for transferring large datasets in HPC environments. Additionally, for users handling data at the extreme high end (up to and possibly beyond 1PB unfiltered), specialized solutions and support[23] may be needed to facilitate their workflows effectively. Users with more than 10TB largely originate from the climate and weather communities, AI/LLM and the telescope/observatory communities.
- Regarding the data growth from the recorded HPC users, by 2028 there is a data growth of approximately 230% compared to the current ones declared, and for 2030 close to a 500% growth. There are also cases with much more than 500% increase for 2030 and beyond. Characteristically, the input from the interview with CERN/WLCG, which declared a pessimistic scenario of 1,500% increase for 2030+, compared to the optimistic ~500% (the latter is in line with the average data growth increase from the questionnaires).
- Most users are satisfied with current data transfer times (see more details, including the ones reporting issues, in Section 2.24). This suggests that while improvements could be beneficial, the current research networking infrastructure is generally adequate for the majority of users.
- Most workflows have an execution frequency of around one day, with some exceptions.
- A small percentage of users, less than ten percent (10%) use Virtual Private Networks (VPNs).

National and International traffic requirements

- Summary tables of the national and international traffic exchange requirements have been presented, as recorded in the questionnaires.
- The analysis of the requirements has been a very complex and tedious process. The questionnaires included several workflows, each with different HPC run frequencies and data sizes needing to be moved across Europe. In addition, the questionnaires requested the HPC and data centre locations for each workflow, with many workflows having multiple locations. Some workflows even had ten (10) or more locations. Furthermore, the execution frequencies, locations and data sizes declared were not standardised. There were similar (non-standardised) responses for locations (different names for the same locations, HPC/data centres acronyms without a city or country, etc.) and data sizes (e.g. varying data sizes, such as from 10-100 GB). These required cross-checking with the respondents. Several clarification questions were sent via email for all these, and even a few short calls or interviews were conducted to clarify some open issues.

[22] In all cases, irrespective of the type of job, the patterns are similar. This means that for interactive type of jobs (e.g. data assimilation), the size of data being uploaded during runtime is relatively small and the same goes for the size of data during the upload phase.

[23] E.g. application software that can use multiple lines simultaneously (e.g. 2x 100 Gbps or 2x 400 Gbps), special data architectures such as the one from Destination Earth (with storage at different levels, i.e. edge, central, syncing among them), content delivery networks (CDNs), and appropriate IT and network support including (MTU size jumbo frames) to optimise the data transfers.

- Total traffic requirements are 6.5 PBs per day currently, 8.4 PBs in 2025, 14.1 PBs in 2028 and 42 PBs in 2030 and beyond.
    - From the 6.5 PBs of current requirements (2023-early 2024), the sum of the current daily national traffic requirements is around 2.28 PBs, while the sum of the international traffic requirements is around 4.23 PBs.
    - From the 8.4 PBs in 2025, the sum of the current daily national traffic requirements is around 3.4 PBs, while the sum of the international traffic requirements is around 5 PBs. Compared to the current requirements, the growth is 126%.
    - From the 14.1 PBs in 2028, the sum of the current daily national traffic requirements is around 6.1 PBs, while the sum of the international traffic requirements is around 8 PBs. Compared to the current requirements, the growth is 213%
    - From the 42 PBs in 2030+, the sum of the current daily national traffic requirements is around 19.1 PBs, while the sum of the international traffic requirements is around 22.9 PBs. Compared to the current requirements, the growth is approximately 635%. This exceeds the average 500% growth recorded in data sizes, mostly due to the higher growth rate of national traffic.
    - With regards to the different countries:
        - In terms of outgoing daily traffic requirements, Italy stands out first, with a total of 1.55 PBs, most of it being international (1.29 PBs).
        - Spain is second in terms of international traffic requirements, i.e. 1.26 PBs, and a total of 1.37 PBs.
        - Germany is third in terms of international traffic requirements, with 0.67 PBs, and a total of 0.80 PBs.
        - Finland follows with 0.63 PBs of international traffic requirements, matching its total of 0.63 PBs.
        - Portugal is first in terms of national daily traffic requirements, with 1.51 PBs, and has a very low amount of international traffic requirements recorded (0.29 TB per day).
    - With regard to different HPC sites (see also Annexes 10.2, 10.3 and 10.4):
        - The HPC site with the most traffic requests[24] is CINECA/Leonardo[25] in Bologna appearing 62 times in the workflows. In terms of data exchange requirements, approximately 1.41 PBs have been recorded for Leonardo, most of which are outgoing.
        - The next HPC site with the most traffic requests is BSC in Barcelona, appearing 46 times in the workflows. In terms of data exchange requirements, approximately 1.46 PBs have been recorded for BSC, most of which is outgoing.
        - LUMI in Kajaani appears 41 times in the workflows. In terms of data exchange requirements 0.6 PBs have been recorded for LUMI, most of which is outgoing.

[24] This refers to the number of traffic requests in the workflows received in the questionnaire.
[25] In general, in terms of hyperconnectivity the access is provided to a site, not necessarily only to the supercomputer. However, in this case there are two CINECA sites with different links. The older one hosting the previous machines (including Marconi) and CINECA/Leonardo. So, CINECA/Leonardo clarifies which site it is.

    - Other sites with considerable traffic demands is KIT in Karlsruhe (1.2 PBs total traffic, most of it outgoing), INCD in Lisbon (1.5 PBs total traffic, most of it outgoing and within the country), and INFN-CNAF (0.6 PBs, most of which is incoming).
    - Many of these are EuroHPC hosting sites, although BSC requirements pertain to previous MareNostrum installations.
- Note that with an 800 Gbps link (i.e. 2 x 400 Gbps currently as there is no 800 Gbps interface/link yet), approximately 0.34 PBs of traffic can be moved in one hour and with a 1.2 Tbps link (i.e. 3 x 400 Gbps), approximately 0.5 PBs of traffic can be moved in one hour. This estimation does not account for header overhead and assumes the support of jumbo frames (see Section 5.4). Furthermore, to be able to stream data across two or three different links (2x400 Gbps or 3x400 Gbps), special software at the application layer or specific network features are needed. More information is provided in the Network Providers section, as well as the overall architecture. In other words, 1 PB can be moved in a little over 2 hours with a 1.2 Tbps link, and in more than 3 hours with an 800 Gbps link.
- Not all connectivity requirements in the questionnaires are complete or usable and some requirements may not have been gathered. On the other hand, not all these requirements are currently recorded in the research networks where the vast majority of the research and academic centres are connected. Given the above considerations, it is recommended that an overprovisioned network be designed.
- There is a delicate point that needs to be considered concerning the type of computing paradigms, namely whether it is High Performance Computing (HPC) or High Throughput Computing (HTC) (see Section 2.28 for more details). There are also other types of applications, such as the ones related to Artificial Intelligence, including Large Language Models (LLM) and beyond, which require GPU-based systems (see also Section 8 on the impact of AI). Given that the study extended beyond the EuroHPC hosting sites to include national HPC systems and major data centres (e.g. research infrastructures and data spaces), all diverse requirements, including those for HTC, have been incorporated into the study and the aggregated requirements.
- Based on the above considerations and the requirements' tables per period, it is evident that the top country sites should be upgraded to approximately Terabit levels (800 Gbps or 1.2 Tbps) by 2025 and/or 2028.

Regarding the HPC providers' questionnaires:

- There was a very good sample of national (~50%), EuroHPC (~30%) and regional (~20%) HPC centres that responded to the questionnaire.
- The vast majority (97%) are involved in HPC-related projects, as expected. A broad range of around 60 projects has been recorded.
- In terms of numbers of users, this varies due to the presence of all types of HPC centres, from regional/institutional to national and European.
- Around 25% of the HPC providers declared having commercial users.
- Data protection and compliance with relevant regulations, including GDPR, are taken extremely seriously. Some centres mentioned specific certifications.

- HPC providers confirmed that batch processing is the dominating use case among HPC centres, consistent with the feedback provided by HPC users.
- The feedback received on security and VPNs from HPC providers is similar to that from HPC users, namely, that only a small number of their users require security and VPNs.
- There is a difference between the end users' responses, with a significant percentage being unaware of their internet provider, and the responses from HPC centres, whose managers are well aware of their internet provider, in most cases being the NRENs.
- Most of the HPC centres reported 2x100Gbps as access connectivity, with some cases ranging from 10 Gbps to multiple 400 Gbps. Regarding the projection of access connections in the future, these are expected to increase by around 60% (but in some cases up to more than two times more) by 2025 and around three to four times more for 2028.
- HPC providers declared that their network traffic is expected to grow significantly, around 30-50% annually, with some cases recording even higher rates.
- Most HPC providers (~70%) are not connected through other network providers than NRENs to the public internet. Only 15% of HPC providers are connected through other network providers.
- The vast majority of HPC providers (~90%) declared collaborations with other HPC providers or research institutions to facilitate data and resource sharing.
- A little over half of the HPC providers declared some kind of bottleneck in their infrastructure. These bottlenecks range from connectivity-related aspects such as access bandwidth and IP routing to storage hardware performance limitations, application, software bottlenecks, etc. They are ordered from more to less relevant (with the most relevant being connectivity-related aspects and the rest equally).

Regarding the Network providers' questionnaires:

- NRENs provide access to the vast majority of HPC users and providers in cooperation with GÉANT, which provides interconnectivity at pan-European and global levels. Only research branches or departments of commercial entities can be connected to the NRENs based on their Acceptable Usage Policies (AUPs).
- The number of HPC providers served by NRENs varies from one (1) HPC centre up to more than six (6) HPC providers. In fact, around half of the respondents serve five (5) or more HPC centres.
- The vast majority (more than 90%) of NRENs serve or plan to serve commercial R&D users, i.e. the research branches/departments of commercial entities. Close to 80% serve them without any special requirements, such as VPN.
- Through the NRENs questionnaires' responses, services are provided to seventy-two (72) different HPC centres. It is estimated that the R&D HPC centres connected via the NRENs are in the order of hundred (100) HPC centres. Thus, the sample received is excellent. Around 40% of these (72) HPC centres are connected at speeds ranging from 10-100 Gbps. Given the higher sample size (compared to the full responses from HPC providers, which is less than half), smaller HPC centres are included, resulting in lower access speeds than those declared in the HPC providers' responses. Additionally, around 10% have even smaller speeds (1-10 Gbps), while another 45% have speeds exceeding 100 Gbps.

- Most of the HPC providers (70%) have dual access to either the same NREN point-of-presence (PoPs) or to two PoPs (dual homed).
- Other services provided to HPC centres are Layer 2 VPNs, Layer 3 VPNs, and protection against Denial-of-Service attacks.
- In the majority of cases, (85,51% of the HPC centres) the relevant fibre infrastructure connecting the HPC provider is owned and/or operated by the NREN. The remaining cases are owned by the HPC provider. Regarding the equipment interconnecting the HPC centre with the NREN, there are several different choices declared, including L1, L2 and L3 equipment, and in several cases there is also no customer premises equipment.
- Most network providers (~85%) provide internet connectivity to the HPC centres/users.
- The majority of NRENs have completed a series of upgrades allowing their core to handle 400 Gbps up to 800 Gbps. Some are already planning to expand their core network to ensure Terabit connectivity. Furthermore, some Network Providers replied that they could complete upgrades towards 400 Gbps links within a few days (less than ten) and all the network providers can deliver upgraded links within six (6) months.

Regarding the EuroHPC Federation Platform (EuroFP):

- To fulfil EuroHPC JU's goal to develop, deploy, maintain, and extend in the European Union a world-leading federated and secure supercomputing, quantum computing service, a market study was conducted during the first half of 2023 to understand and explore the existing federation solutions and platforms in the European ecosystem. The results of this five-month study were presented to the EuroHPC JU Governing Board and consequently, based on the Governing Board decision, a call for the procurement of the EuroHPC Federation Platform was launched in September 2023. The call for procurement was closed in December 2023 and the evaluations of the submitted proposals will be carried out in 2024.
  The public procurement action is to build and deploy a fully operational infrastructure federation solution for EuroHPC JU supercomputers, providing access to a rich portfolio and professional-quality services in all relevant domains from access to pre-processing, computing, analysis, post-processing and storing. The procurement also covers operations, maintenance and support of the federation solution and services for five years.
- The EuroHPC Federation Platform should be robust, secure, scalable, flexible, and user-centric. It should constantly be improved and upgraded following user feedback and the state-of-the-art of the underlying core technologies. It should be open enough to be able to accommodate all EuroHPC JU supercomputers and quantum computers including the future industrial supercomputers. It intends to offer high quality of service management compliant with industrial standards. The proposed federation should also provide a superior user experience, usability, and ease of use for a large number of users.
  Regarding potential connectivity requirements from the EuroFP call, it is not possible to make estimations based solely on the tender specifications, especially without examining the specific proposed solutions. Nevertheless, taking into account the connectivity requirements raised, the upcoming AI factories and the EuroFP project, the hyperconnectivity network for the EuroHPC hosting sites should be at the Tbps level.

Recommendations on the overall architecture choices of the hyperconnectivity network:

- *Recommendation - Network-as-a-Service (NaaS) approach: Based on the analyses performed, it is recommended to adopt a Network-as-a-Service approach. This approach will leverage existing infrastructures and networks already operating across Europe, which have proven their capacity in servicing high-demanding research infrastructures.*
- HPC traffic differs significantly from commercial traffic. In the case of HPC traffic, the network is required to handle a small number of flows, mostly elephant ones, i.e. extremely large (in size) continuous flows.
- A wide range of access speeds in HPC and data centres has been recorded. Therefore, it is believed that each HPC/data provider needs to upgrade at their own pace, taking into account the traffic exchange requirements they receive.
    - *Recommendation – Implement Access classes (Tiers) for HPC/Data Centres Access: Indicative access classes (more like averages) are recommended for the access speeds of HPC/Data centres. Actual speeds will be determined based on traffic exchange requirements.*
- Due to the nature of the HPC user scenarios, especially in bulk data transferring as described in Section 2.25, Maximum Transmission Unit (MTU) size plays a significant role when comparing network overhead between different MTU sizes, such as 1500 bytes and 9000 bytes (also known as "jumbo frames"). With an MTU size of 1500 bytes, the overhead (packet headers) is approximately 10%, while with an MTU size of 9000 bytes, it is only around 1%.
    - *Recommendation – MTU size 9000 bytes (jumbo frames): Given the HPC traffic patterns and data transfer sizes of more than 10GB, and in several cases reaching several hundred TBs, it is recommended to use "jumbo frames" (MTU size of 9000 bytes) in the hyperconnectivity network.*
- Based on the requirements gathered, an overlay network or VPN-based network is not justified and can introduce several challenges. Given the overprovisioned type of the network, a plain IP-network is recommended.
    - *Recommendation – Plain IP network: A plain IP network is recommended for the hyperconnectivity network. This can take the form of an IP network with MPLS[26] encapsulation and segment routing capabilities. A vendor-neutral implementation provides sufficient tools for either traffic engineering needs for elephant flows or even (if need be) a closed group communication (i.e. MPLS VPN).*
- Based on the analysis in this document, there is no need for special security measures on the core network layer, such as encryption on the optical or transmission layer. These measures would significantly increase service costs and complicate the leveraging of existing infrastructures. Current technology trends, suggest that security measures should be placed at the edge of the network, at the application layer (see more details in Section 5.6 and recommendation below).
    - *Recommendation – Security at the Edge/Basic Transport Security: It is recommended not to use encryption at the optical or transmission layer, as this would increase costs and may impose performance limitations on hyperconnectivity networks. Leveraging established*

[26] Multiprotocol Label Switching - Wikipedia

industry protocol stacks such as SSH/SSL at the edge would provide the necessary privacy for data in transit. Different measures should be deployed for data at rest (in storage). Basic transport security including protection from Denial-of-Service attacks and Intrusion Detection Systems (IDS)/Intrusion Prevention Systems (IPS) techniques should be implemented. Firewalling and access lists, among others, should also be considered.

Recommendations

- *Projected Computational Needs – EU and national HPC providers should compare the growth rates recorded with their own plans, particularly looking ahead to 2030 when relevant planning information may be limited. This comparison will help ensure alignment between projected computational needs and available resources.*
- *Enhanced Security Measures – Given the critical nature of many HPC applications, a key recommendation would be to strengthen security protocols for data transfer and storage. This may entail implementing advanced encryption methods, secure authentication processes, and enhanced monitoring of network traffic*
- *HPC Infrastructure Upgrades: The demanding needs from users towards HPC Centres and the foreseen data growth require ensuring that HPC centres have, and more importantly will have in the future, the necessary infrastructure to support efficient data storage, management, and transfer at scale to meet these demands. This entails corresponding strategic planning and investment to ensure that computational resources keep pace with these demands.*
- *Network Infrastructure Upgrades: Upgrading existing network infrastructures to support higher data transfer rates, reduced latencies, and increased bandwidth is recommended. This would ensure that the network can handle the growing volume of data and the complexity of computations required by modern HPC applications.*
- *Hyper Connectivity network overprovisioning– The study team recommends an overprovisioned network: On one hand, HPC requirements are best satisfied with overprovisioned architectural patterns, which differ from existing commercial ones that are typically highly overbooked. On the other hand, there are several parameters related to the study needs which are not fully known, including the AI Factories needs, the EuroHPC federation needs and the EuroHPC Industrial system(s).*

Other recommendations from the additional material are summarised as follows:

- *Development of New Technologies: Encouraging the development and integration of new technologies, such as quantum computing elements or advanced data compression algorithms, is recommended to improve the efficiency of data transfers and processing capabilities.*
- *Collaborative Frameworks & EuroHPC Federated Infrastructure: Establishing or enhancing collaborative frameworks among European HPC centres, including the federation of EuroHPC infrastructure, is recommended to foster better sharing of resources, expertise, and best practices. This could also involve forming partnerships with private sector entities to leverage their technologies and networks. This point is already being addressed by the EuroHPC JU with the federation call and is relevant to the upcoming EuroHPC federation project.*

- *Standardization of Protocols: Offering standard protocols and interfaces to ensure compatibility and interoperability between different HPC centres across Europe is recommended. This would facilitate easier and more efficient data sharing and collaboration. This point is relevant to the upcoming EuroHPC federation project and beyond.*
- *Training and Support Programs: Developing comprehensive training and support programs for users of the HPC infrastructure is recommended to ensure they can fully leverage the available resources. This includes training on computing and network technologies, best practices for data management, and security awareness. This point is already being addressed by the EuroHPC JU with the EuroCC and Castiel project series and is relevant to the ongoing EuroCC2 and Castiel2 projects.*
- *Regular Assessment and Feedback Loops: Establishing mechanisms for regular assessment of network performance and user satisfaction is recommended to allow for continuous improvement based on user feedback and evolving requirements.*

These inferred recommendations aim to address the strategic, technical, and operational aspects identified through the needs assessment, setting a roadmap for future developments in HPC connectivity across Europe.

4.1. Questionnaires' overview and statistics

Three (3) questionnaires were diligently prepared in an iterative process and in close consultation with all relevant stakeholders. Draft versions of the questionnaires were presented in a series of workshops, where valuable feedback was provided. This feedback was then integrated, and the questionnaires were optimised and finalised. The three (3) questionnaires correspond to each of the three stakeholder groups, namely HPC Users, HPC Providers, and Network Providers, while the inputs from the "Other important stakeholders" identified, namely AI stakeholders and Data Centres, are included in the HPC Users questionnaire. The questionnaires ended up with well over twenty (20) questions each, subdivided into different sections. The LimeSurvey tool was finally chosen to conduct the survey as the software complied with all of the requirements, especially when it came to the design and evaluation capabilities.

While the analysis of the questionnaires started after the initial deadline of questionnaires (31.01.2024), the questionnaires were left open also for the next months. Given that information on the Common European Data Spaces[27] was published only around mid-March, the questionnaires were left open until the first week of April, to allow three full weeks for inputs. Initially, the final deadline set was end of March, but after a set of reminder emails, it was shifted to the first week of April. Due to the low input from data providers, both Research Infrastructures (ESFRIs[28] and ERICs[29]) and EU Data Spaces, a set of emails was sent in an effort to understand the reason for this.

[27] Common European Data Spaces | Shaping Europe's digital future (europa.eu)
[28] Projects and Landmarks | ESFRI Roadmap 2021
[29] ERIC Landscape - European Commission (europa.eu)

The main question was whether they do not have any HPC requirements, and/or if they cannot foresee future HPC requirements. If so, they were asked to declare this by responding to the email, as this would help the study in drawing conclusions. In the end, approximately 10 responses were received: a few indicated that they do not have HPC requirements; a few mentioned that they may have HPC requirements in the future but cannot foresee them; a few completed the questionnaire, either fully or partially; and a few provided some information via email regarding their HPC needs, data transfer sizes and locations.

The final numbers of responses to the questionnaires are provided below.

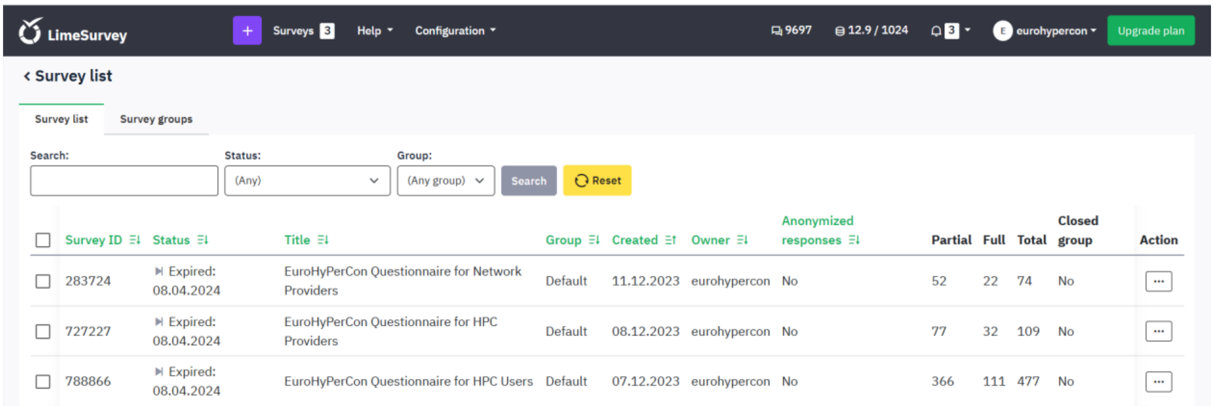

Survey ID	Status	Title	Group	Created	Owner	Anonymized responses	Partial	Full	Total	Closed group	Action
283724	Expired: 08.04.2024	EuroHyPerCon Questionnaire for Network Providers	Default	11.12.2023	eurohypercon	No	52	22	74	No	...
727227	Expired: 08.04.2024	EuroHyPerCon Questionnaire for HPC Providers	Default	08.12.2023	eurohypercon	No	77	32	109	No	...
788866	Expired: 08.04.2024	EuroHyPerCon Questionnaire for HPC Users	Default	07.12.2023	eurohypercon	No	366	111	477	No	...

Figure 11. Final numbers of replies to the questionnaires

The following responses have been finally recorded:

- HPC Users: 111 full/completed responses; 366 partial/ongoing, total 477.
- HPC Providers: 32 full/completed responses; 77 partial/ongoing, total 109.
- Network Providers: 22 full/completed responses; 52 partial/ongoing, total 74.

Thus, a total of 165 full/completed responses and 495 partial/ongoing have been received, giving a total of 660 entries. Although this is an impressive number of inputs, a significant number of partial surveys do not have enough information to satisfy the most important requirement (data sizes moved across Europe to be able to calculate the bandwidth requirements). Furthermore, a large number of surveys were not actually “submitted” because the submit button was not pressed. For these cases, the study team decided to contact the respondents via email to obtain permission to use their data, as pressing the submit button included the necessary consent. Only a very small number of responses was received - three (3) to be exact.

4.2. Overall Architecture Choices of the Network

In the following Section a high-level description of the design principles to be deployed in the Network Architecture is presented.

4.2.1 Network as-a-Service solution

The scope of this study is to determine the characteristics of a network that would interconnect all HPC ecosystem actors across the continent performantly and efficiently. It is advised that the EuroHPC

JU should not build a new parallel infrastructure to achieve this. Furthermore, it is recommended not to create a related organisation or division to support a brand-new infrastructure. **So, a Network-as-a-Service (NaaS) approach is preferred that will leverage existing infrastructures and networks, that are already operational across Europe (GÉANT and NREN networks) and have proved their capacity in servicing high-demand research infrastructures. Note the current GÉANT and NREN networks are interconnected to the global internet and industrial users can reach the HPC centres via Internet Exchange Points or direct peerings, either from the first (NRENs) or via the latter (GÉANT).**

The figure below depicts a high-level approach to a tiered Network-as-a-Service infrastructure.

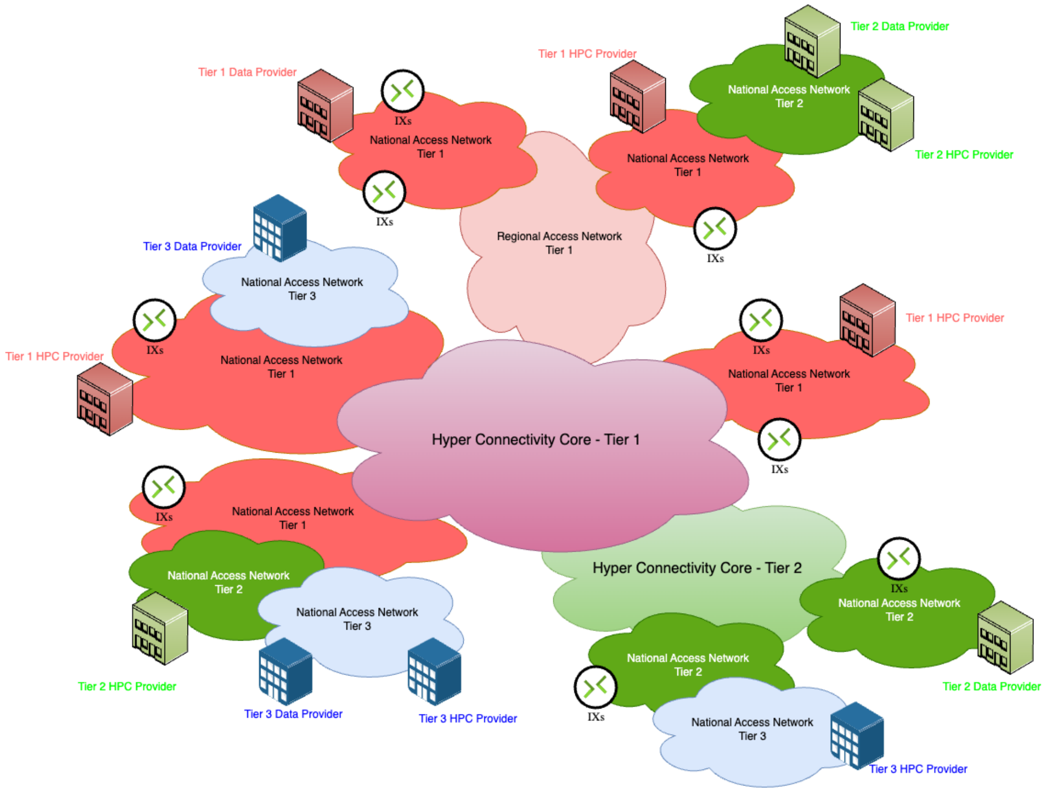

Figure 12: High-level approach of a tiered NaaS infrastructure

GÉANT is considered a successful example of such an infrastructure that has been servicing the European Research and Academic community for over three (3) decades now. It represents a consistent investment of the European Unit towards a common, inclusive network infrastructure that allows the European R&E community to collaborate and thrive.

4.2.2 HPC vs commercial traffic patterns

As ordinary commercial ISPs and telecom networks focus on common file transfer applications, they tend to overbook the sum of input capacity for the corresponding required backhauling. This is absolutely viable and acceptable since they rely on the fact that those networks serve hundred thousand customers (if not millions) since it is assumed that not all customers will demand their access link concurrently to perform at its highest throughput. Moreover, commercial internet traffic consists of millions of small to medium (volumetric) flows (i.e. "*mice flows*[30]") for a short duration of time. "Elephant flows[31]" are rare.

In research networks[32], traffic patterns differ greatly. Although the number of users is far less, each one may put excessive stress on the network infrastructure by requesting either long-lasting huge volumes of data (multiple PBs) to be transferred to a remote location or even by deploying low-latency demanding applications.

The following usage patterns and conclusions have been identified in the narrower context of HPC networking needs:

- *Most mission-oriented workflows are elephant flows with large duration (i.e. >5 hours) and an execution frequency of around once per day, with some exceptions.*
- *There is a wide range in the amount of data to be transferred, including some workflows that transfer large amounts of data. In this case the median is 300 GB and the average 1.37 TB.*
- *There is no apparent correlation between the amount of data to be transferred and the recurring window for transfers, indicating that transfer schedules are set independently of data volume.*
- *The workflows predominantly involve batch processes, which implies they are designed to regular jobs, rather than interactive or other ones. Still there are some interactive user scenarios, indicating a variety of workflow types being executed.*

The above conclusions indicate that HPC traffic differentiates a lot when compared to commercial traffic as the network is requested to serve a small number of flows but mostly *elephant ones*. Since the multiplexing entropy is kept low, network-wide hashing techniques deployed across parallel links (either Equal Cost Multi Path-ECMP[33] or Link Bundling) are not always efficient[34]. This is also valid for AI traffic requirements (in particular relevant to the AI factories). These will be co-hosted at the EuroHPC hosting sites, which are already connected to the NRENs. For the industrial site internet exchange points or peerings can be used, while additional commercial links may be provided if needed. <u>Such diverse requirements are implemented with overprovisioned architectural patterns, a totally different implementation principle from existing commercial ones which are typically overbooked.</u>

The modus operandi of a desired research network (yet interconnected with the global internet) will need to adhere to a different capacity upgrade rule from the existing commercial Telecom networks which upgrade capacity when there is > sixty percent (60%) <u>average</u> utilization for a typical observation period. In research networks and especially on HPC/HTC networks there is a need to focus primarily on <u>peaks</u> and secondarily on <u>averages</u>. Keeping backbone links below sixty percent (60%) <u>peaks</u>, while also maintaining thresholds for ninety-fifths (95ths) averages is the rule of thumb. Repeated traffic

[30] https://www.sciencedirect.com/science/article/pii/S1389128605001118
[31] Elephant flow - Wikipedia
[32] https://www.sciencedirect.com/science/article/pii/S0167739X18302322
[33] C. Hopps. Analysis of an equal-cost multi-path algorithm
[34] https://marchiesa.bitbucket.io/docs/chiesa/ecmp.pdf

peaks over sixty percent (60%) of the nominal link capacity during a week period led to a respective capacity upgrade of a backbone link.

4.2.3 Access Classes and Resiliency

HPC and data centres are critical research resources that work on time-sharing basis. They are shared across many research groups that perform their computational jobs according to the slots that have been allocated or store their research data. Any disruption to the availability of the HPC and data centre resources caused by network unavailability or bad performance cause delays and pushbacks on those groups' schedule. During consultation, some users stated that they face challenges related to low bandwidth, both in-house (or at the data centres they store their data) and at HPC Centres, leading to slow transfer rates, timeouts, and difficulties in handling large datasets have been categorised according to their capacity.

Another important note is that many HPC centres are developed in constellations (in proximity between them - HPC campuses) where many systems are hosted in the same area. This makes network availability of paramount importance. Thus, dual access is the most common connecting paradigm of HPCs or NRENs.

A wide range of access speeds in the HPC and data centres has been recorded. It is thus believed that each HPC/data provider needs to upgrade at their own pace, taking into account the traffic exchange requirements received. The following table presents an indicative classification of the HPC / data centres and the proposed connectivity bandwidth for each one of them. This is more a possible classification in classes (or tiers), but again this will depend on the exact requirements and can range within the boundaries of each class (in average). But if justified by the needs, it can go outside the boundaries of a class, jumping to higher class, or even going above the highest class. This is also aligned with the AI factories and the federation project requirements (see Section 8), which are expected to provide additional traffic beyond 2025-2026 (towards the 2028 phase and beyond).

Tier	Currently	Phase A (2025)	Phase B (2028)	Phase C (2030+)
Class A	200-400 Gbps	400-800 Gbps	800-1.2 Tbps	1.2-1.6 Tbps
Class B	100-200 Gbps	200-400 Gbps	400-800 Gbps	800-1.2 Tbps
Class C	10-100 Gbps	100-200 Gbps	200-400 Gbps	400-800 Gbps
Class D	1-10 Gbps	10-100 Gbps	100-200 Gbps	200-400 Gbps

In terms of topology a general rule of thumb favours the ring topology as it can accommodate failures more effectively. This is a common architectural choice also in commercial networks. It is proposed that every access PoP connects to two backbone (i.e. East-West) Points of Presence. The East West capacity should be forty percent (40%) more than the access capacity (overprovisioning) thus allowing some extra "space" for traffic engineering of one direction vs the other.

Resiliency to network failures will be achieved by deploying redundancy access schemes such as Dual-Access-Dual-Homed (DADH) or Dual-Access-Single-Homed schemes (DASH)[35]. Those options are depicted in diagram bellow.

A day 0 network topology map is foreseen with multiple 400Gbps in the core to accommodate a 1.2 Tbps aggregated capacity if interface HW is not yet readily available to be deployed. Multiples of 100s Gbps can replace the existing 10 Gbps and 10 Gbps link can replace 1 Gbps access links, thus a 10-fold capacity increase will be readily available for the research community (academic and commercial) of Europe.

Taking into consideration the responses from Users Questionnaire 9b: "Can you estimate how the access connectivity of your HPC centre will look like in the following years (in Gbps)" (The year in question: 2025), user's needs can be accommodated by adopting the following indicative topology and capacities.

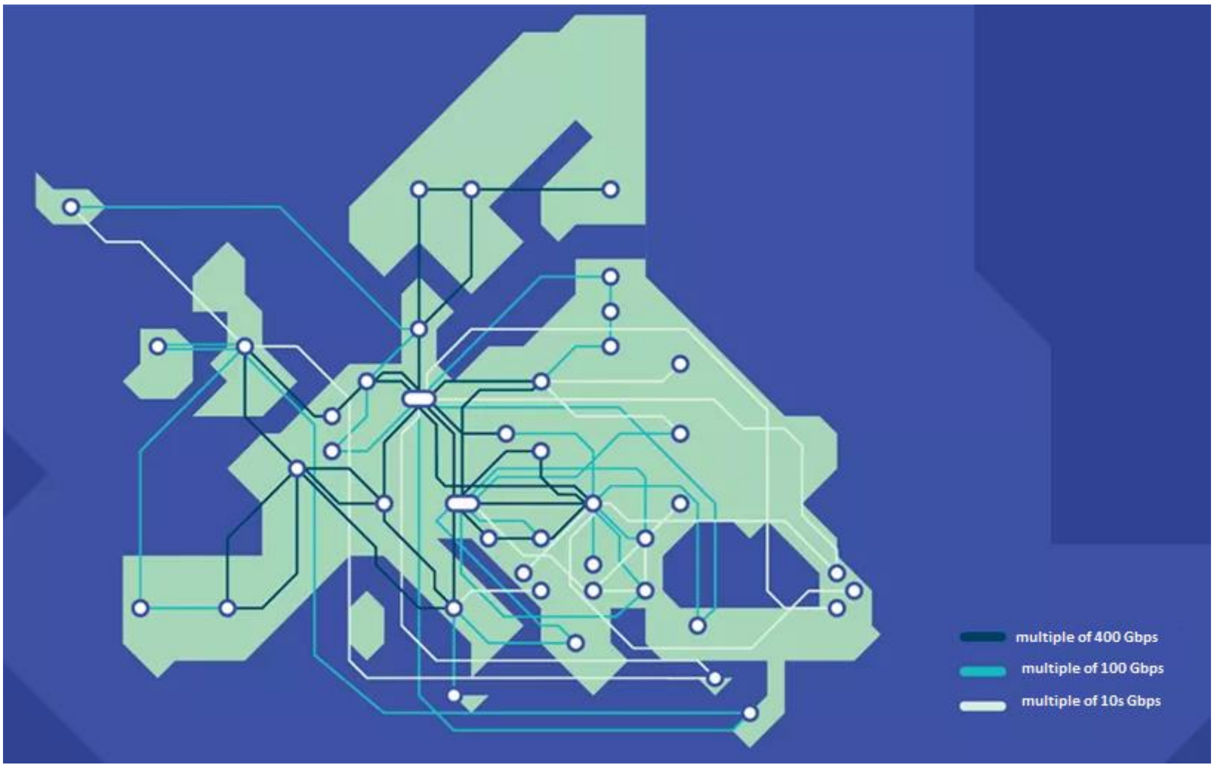

Figure 13: Indicative topology

4.2.4 Plain IP routing

Consultation with all HPC ecosystem actors such as:

- End Users,
- HPC providers,
- Network Providers and

[35] Multi-homed vs. single-homed network | DataPacket.com

- Data Providers

showed that in current implementations no overlay networks construct is used. Communication between all actors is achieved over public Internet. The diverse nature, location and implementation details make using the public Internet the only choice.

Creating an overlay (i.e. security VPN) would put a significant burden in the implementation phase of the network. The volume of the problems that need to be resolved regarding:

- IP Addressing and possible re-numbering,
- remote access (such as enterprise entities) difficulties,
- performance issues in case that massive tunnelling or encapsulation techniques are deployed

is very high and make the overlay option unjustifiable.

This conclusion is supported by the experience gained by GÉANT as in the past it created overlays to support analogue infrastructures such as PRACE or LHC, which soon migrated to plain IP routing. A plain MPLS encapsulation with segment routing[36] capabilities as a vendor neutral implementation provides enough tools for either closed group communication (i.e. VPN) or traffic engineering needs even for elephant flows.

4.2.5 Security as an edge function

During the consultation phase, dedicated parts were allocated to security concerns regarding the data transmitted over the network and the security of the infrastructure as a whole. The main outcomes of those discussion are the following:

- While security and privacy are critical factors for certain fields, especially those dealing with personal or sensitive data, they may not be a universal concern across all HPC users. It is important for HPC Centres to be equipped appropriately to handle the specific needs of users who do require enhanced security measures, even if they represent a smaller segment of the overall user base.
- There are specific cases, mainly from industry that express security concerns especially regarding to copyrighted protected material or sensitive data.

Based on the above conclusions enforcing special security measures on the core network layer such as encryption on the optical or transmission layer is not recommended. Those would increase the cost of the service significantly and would make leveraging existing infrastructures more difficult. On the contrary, current technology trends (i.e. AVX-512 with elliptic cryptography)[37], suggest that **security measures should be placed at the edge of the network in the application layer**. Leveraging on established industry protocol stacks such as SSLs would provide the necessary privacy for data in transit, while different measures should be deployed when data are rest in storage.

[36] https://datatracker.ietf.org/doc/html/rfc8663
[37] https://networkbuilders.intel.com/docs/networkbuilders/intel-avx-512-and-intel-qat-accelerate-wireguard-processing-with-intel-xeon-d-2700-processor-technology-guide-1647024663.pdf

Since the HPC servicing infrastructure will be exposed to the public internet, special measures should be deployed to secure the infrastructure against malicious actors. Those measures should include:

- **Protection against Denial-of-Service attacks**
- **Protection against intrusions by leveraging Intrusion Detection Systems (IDS)/Intrusion Prevention Systems (IPS)[38] techniques.**
- **Securing perimeter by deploying firewalling, application protection and access lists**

It is important to note that the proposed techniques and measures could not be out-of-the-self solutions. As the network in study will offer Terabit access, each tool should be carefully examined regarding its ability to support the throughput requested with regard to its cost weight portion.

5. Gap Analysis

The Gap analysis starts with identifying and describing existing connectivity initiatives used by the research and academic community, such as GÉANT and the NRENs. This is because the NRENs and their upstream provider, GÉANT, are used to interconnect the vast majority of the identified in this study HPC and data Centres. In fact, the only sites that are not connected via the NRENs are the SMEs or other private or industrial sites in the list of end sites that was prepared in this study, plus one special case, i.e. Destination Earth, where connectivity has been provided as part of its data (cloud) architecture.

There is currently an updated list of around 350 sites (HPC and data centres), which have been prioritised based on the strategic importance of the site and the corresponding connectivity requirements identified in the study. The EuroHPC hosting sites (HPC and Quantum) feature in the first priority of the list, while the next two priorities include other major HPC and data sites and other sites with significant connectivity requirements, i.e. >= 10 TB daily traffic (which corresponds to ~1 Gbps). The first three priorities are proposed as mandatory for the upcoming tender and encompassed around 150 sites. The next priorities correspond to other non-major sites or sites without significant identified requirements and are proposed as optional.

This list of end points (HPC and data centres) has been used to evaluate the complementarity and gaps between related EU-funded projects (e.g. GN4 and GN5[39] Framework Partnership Agreements) and the needs projected by this study for the upcoming period, as outlined in the tender specifications. This ensures the fulfilment of unmet requirements specific to the HPC domain. In other words, the gap analysis framework presented herein systematically identifies and addresses differences between the current network capabilities and future requirements. By comparing the target situation with the current state for each requirement, the analysis highlights areas where existing or planned solutions already meet specific needs and identifies gaps where further development is necessary. This process allows for the informed allocation of resources to enhance HPC performance and achieve strategic

[38] Intrusion detection system - Wikipedia
[39] GÉANT Project | GÉANT (geant.org)

objectives effectively. Special emphasis is placed on avoiding redundancy with other EU-funded initiatives and focusing on the unique needs of the HPC environment.

Additionally, it presents a detailed analysis of the infrastructure, connectivity, regulatory and policy gaps identified in the framework of this study. In particular, possible infrastructure gaps at the end sites have been identified in the last mile of the end-to-end network, specifically inside the HPC or data centres. As the hyperconnectivity network is expected to reach at the border of the HPC or data centres, possible gaps inside the campus network of the latter may affect the provided service, acting as bottlenecks. A set of questions was sent to the EuroHPC hosting sites on this point, requesting information about the border elements of the end sites (routers, switches), the capability/capacity of their internal network and other possible rate-limiting equipment, such as firewalls. The results are presented in Section 5.

Furthermore, technology gaps have been identified, namely in terms of Terabit interfaces for networking equipment, which are currently not available. A related IEEE standard[40] on 1,6 Tbps is only expected by mid-2026[41]. On the other hand, the corresponding IEEE standard on 800 Gigabit Ethernet was approved in February 2024[42], and more and more such interfaces will be available in the market.

Other potential administrative, policy and legal gaps, or at least risks, have also being identified. One of them is the fact that in case of an open tender, all end sites (HPC/data centres) proposed to be listed in the technical specifications of the hyperconnectivity tender, will need to agree to be included in such a list. In other words, the legal representatives of the end sites will need to provide their consent so that the sites they represent are listed in the tender specifications. This will require legal actions from the legal representative (signatory) of the site[43], which may be time-consuming based on the site's decision-making process. Furthermore, although the hyperconnectivity tender is 100% funded, meaning no contribution will be requested from the end sites, there is still some bureaucracy involved in obtaining these signatures and this needs to be taken into consideration.

On the other hand, there may be cases where the internal end site (campus) network has bottlenecks, necessitating some investment from the end sites to fully exploit the hyperconnectivity network. If this is the case, a decision for the end site to participate in the tender may be delayed. Another potential gap is that end sites may not agree to be connected to a provider other than the corresponding NREN, i.e., a commercial network provider or ISP. This is because the end sites have the freedom to select the connectivity provider of their choice, based on the legal type and rules their site is governed by. In the vast majority of cases, the NREN is the de-facto connectivity provider for research and education organisations. And sites may not want to change this. Such cases were reported during the consultation phase of the network. So, if the proposed legal instrument to implement the hyperconnectivity

[40] IEEE P802.3dj 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet Task Force (ieee802.org)
[41] Adopted IEEE P802.3dj Timeline (28 Nov 2023) (ieee802.org)
[42] [802.3_B400G] IEEE P802.3df Standard Approved! (ieee802.org)
[43] In case the hyperconnectivity tender is an open tender, in order for the proposed sites to be listed in such a tender, the legal representative of the site needs to agree, as the hyperconnectivity service will need to be terminated inside their premises.

network is based on an open tender (or other related open process), there is a risk that some sites may not agree to be connected. In case these sites are included as mandatory in the next tender (60M Euros), there is a risk of the tender or the implementation of the contract failing.

Legal compliance and a set of alternative legal options for the upcoming procurement have also been identified. From the related analysis, a prioritised list of procurement options is proposed. If the relevant legal and financial conditions are fulfilled, the negotiated procedure without a contract notice is the optimal choice. This approach is faster and technically safeguarded, given the expertise of GÉANT and the NRENs, without the need of technical assistance from third parties. Moreover, this approach will be much more flexible with regards to the expansion to the second tender in 2027 (100M Euros). However, the feasibility of such a proposal needs to be examined by the EuroHPC JU and GÉANT. If this is not possible, then the next options in priority would be the competitive dialogue and, lastly, the open process.

All in all, no major connectivity gaps have been identified in the GÉANT and NRENs networks. With proper upgrades and expansions, these networks can satisfy the prioritised list of end points, and their identified requirements. The other identified gaps - such as infrastructure, technology, administrative, regulatory and compliance - are also not showstoppers, and the upcoming tender can proceed as planned.

5.1. Overview of Existing connectivity providers

Two main alternative solutions were considered in the context of the hyperconnectivity study. Both rely on the offerings by network providers spanning across the European continent and even beyond. These are networks that have been developed the last decades either by public or private funds and are currently serving mainly other smaller scale networks.

Tier-1 Internet providers[44] or telecom providers offer a range of services tailored to meet the diverse needs of their customers. From IP transit and wavelength services to spectrum sharing and dark fibre, these offerings ensure high performance, scalability, and cost-efficiency, enabling businesses to maintain robust and reliable network connectivity in an increasingly connected world. Through leveraging these services, customers can achieve enhanced global reach, dedicated high-speed connections, cost-effective bandwidth solutions, and complete control over their network infrastructure.

In this context two main categories have been identified. Commercial Tier-1/telecom providers and the GÉANT network.

[44] https://www.thousandeyes.com/learning/techtorials/isp-tiers: Tier 1 Internet providers are the networks that are the backbone of the Internet. They are sometimes referred to as backbone Internet providers. These ISPs build infrastructure such as the Atlantic Internet sea cables. They provide traffic to all other ISPs, not end users. Tier 1 ISPs own and manage their operating infrastructure, including the routers and other intermediate devices (e.g., switches) that make up the Internet backbone. Key Tier 1 ISPs include AT&T, Verizon, Sprint, NTT, Singtel, PCCW, Telstra, Deutsche Telekom, and British Telecom.

5.2. Gap Definition

Typically, gap definition involves identifying and analysing the disparities between the current state of a network and its desired future state. This process encompasses evaluating various network parameters such as bandwidth, latency, and reliability, which are relevant for a specific sector e.g. an HPC centre, cloud providers, etc, and comparing these metrics against the requirements needed to support anticipated future workloads and applications tailored to the users' needs. The goal is to pinpoint specific areas where the network's performance is insufficient, which could impede achieving optimal performance and business objectives. In the context of the hyperconnectivity network, insights on the peculiarities of the network requirements have been gained by conducting thorough analysis on the questionnaires filled in by the Users, HPC centres and the current network operators.

5.2.1 Gap Analysis Framework

Typically, the gap analysis framework contains two elements. The comparison of the current state with the future requirements and the criticality assessment to prioritize the requirements.

Compare Current State with Future Requirements

Identify Discrepancies in Bandwidth

The goal is to conduct a thorough assessment of the existing network's bandwidth to determine if it meets current and anticipated needs. This involves gathering data about the data transfer rates across various segments of the network including the backbone network as well as the campus one.

If available, network latency helps to identify delays that could impact high-performance computing (HPC) tasks, especially during the transfer of huge data sets needed for the HPC tasks.

Last but not least, the reliability of the current network provider's performance by examining historical data on downtime, poor performance and the frequency of network outages should be taken into consideration.

Gap(s) classification

There are plenty of approaches in classifying the gaps in a network evaluation study. Under this work focus was given on the network speed. Based on that and other strategic objectives of the EuroHPC JU (e.g. that the EuroHPC hosting sites have higher priority than all others), the study team came up with a classification index called **Priority** that varies from 1 to 10, 1 being the highest priority and 10 the lowest.

5.2.2 Analysis of the Connectivity Status and Classification of HPC Centres

In the context of this study, the study team has analysed raw data related to **520** Points of Interest (PoI). For each point, a set of attributes was gathered including the current network connectivity offerings and the study team has analysed, prioritized, and characterized them. The following sections provide an overview and an analysis of the first three priority categories, which the study team considers as the main candidates for the procurement strategy.

5.2.3 Priority categorisation

The following table outlines the distinct categories used to classify the PoIs collected during the first phase of the study (Task 1). The first three categories will be considered as mandatory part of the procurement strategy, while the rest of them may be included as optional.

Priority #	Description	Estimated number of PoIs	Mandatory/ optional	Reason for priority
1	EuroHPC hosting sites	~20 sites	mandatory	Top tier HPC sites
2	Other major EU and national HPC and data sites identified in first phase of the study (Task 1)	~100 sites	mandatory	Other major HPC and data sites
3	HPC/data sites identified in the questionnaire with requirements > 10 TB per day (~1 Gbps)	~30 sites	mandatory	Sites with high data transfer needs
5	Other non-major EU and national sites identified in Task 1		optional	Other smaller scale HPC/data sites
6	HPC/data sites identified in the questionnaire with medium and low requirements < 10 TB(~1 Gbps Mbps)		optional	Other sites with low data transfer needs
10	SMEs/Industrial sites		Optional	SMEs/ industry sites will be connected via IXPs/peerings

5.2.4 Access speeds and classes

A wide range of access speeds in the HPC and data centres has been recorded. It is thus believed that each HPC/data provider needs to upgrade at their own pace, considering the traffic exchange requirements received. The following table presents an indicative classification of the HPC / data centres and the proposed connectivity bandwidth for each one of them. This is more a possible classification in classes (or tiers), but again this will depend on the exact requirements and can range within the boundaries of each class (in average). But if justified by the needs, it can go outside the boundaries of a class, jumping to higher class, or even going above the highest class.

Tier	Current (2023-2024) connectivity	Phase A (2025) connectivity	Phase B (2028) connectivity	Phase C (2030+) connectivity
Class A	200-400Gbps	400-800 Gbps	800-1.2 Tbps	1.2-1.6 Tbps
Class B	100-200Gbps	200-400 Gbps	400-800 Gbps	800-1.2 Tbps
Class C	10-100Gbps	100-200 Gbps	200-400 Gbps	400-800 Gbps
Class D	1-10 Gbps	10-100 Gbps	100-200 Gbps	200-400 Gbps

5.3. Detailed Gap Review

5.3.1 Infrastructure Gaps

A critical parameter for designing the network solution is understanding the type and capabilities of the internal network (LAN) and related networking equipment each HPC/data centre owns. This is because in several cases there may be bottlenecks inside the HPC/data centre network, which may act as gaps in fully utilizing the hyperconnectivity. Examples of such bottlenecks are routers or switches with slower routing/switching capacity than the incoming hyperconnectivity link or firewalls having lower line rates than the hyperconnectivity links. During the needs assessment phase, various adopted approaches have been identified, which need to be considered during the design and procurement phase to ensure compatibility between the new solution and existing systems.

However, it was impractical to evaluate all institutions, as this would take several months. Instead, the study team chose to compile and distribute a questionnaire to Priority 1 members with significant networking needs. This approach aimed to help us understand their deployed network topology and the equipment used for interconnecting their HPC centre with the existing upstream network provider (NREN).

The summarized questions are provided below. The questionnaire was kept concise with only a few essential questions to ensure responses from all institutions. The technical focus included:

1. Understanding if there is a Next Generation Firewall (NGFW) at the network perimeter and its mode of operation. The rationale behind this is that these firewalls do not scale proportionally to routers. A 400 Gbps capable firewall is significantly more expensive than a 400 Gbps router and may lack the necessary number of interfaces required by the host, potentially creating a blocker (i.e. bottleneck) for the institution.
2. Determining if the border network element is designed for solutions that require connections to two different upstream network providers. During workshops with HPC providers, they emphasised that they do not intend to drop the existing link with the local research and

academic community via their NREN. Therefore, all parameters for a dual-homing solution should be considered.

3. Ensuring there is adequate internal capacity in the HPC core system to fully utilise the upstream network upgrade. Upgrading the access link of an HPC/data centre is pointless if the internal (campus) network (LAN) cannot effectively utilise this upgrade.

The following data have been collected by the HPC centres, which have been answered by the beginning of July 2024.

Based on the input received, a short analysis is presented below, including the following conclusions:

1. **Next Generation Firewall (NGFW) Usage**: Most centres employ NGFWs; however, not all do. In particular, Daedalus, LUMI and VEGA do not use[45] Firewalls and alternative architectures are used, such as a Science DMZ[46]. In some instances, the NGFW intercepts all incoming traffic, providing comprehensive protection. In other cases, NGFWs are implemented only for certain sensitive paths, offering targeted firewalling capabilities. This variability in NGFW deployment highlights the need to consider different security requirements and operational modes when designing network solutions.
2. **Border Network Elements:** The centres use a mix of routers and Layer 3 (L3) switches as their border network elements, while only one site (BSC MN5) uses a firewall (routed). In particular, Daedalus, EHPCPL, Lumi and Meluxina use (or plan to use) a router, while Karolina, Leonardo and Levente use L3 switches. Furthermore, EHPCL, Karolina, Lumi and Meluxina utilise (or will utilise) routers capable of handling the full internet routing table, making them suitable for a dual-homing solution. However, an equal number of institutions rely on L3 switches that cannot cope with the full internet routing table. For these centres, implementing a dual-homing solution would not be feasible due to the limitations of their current infrastructure.
3. **Internal Network Capacity:** For all the reported cases, no significant bottlenecks were identified with the downlinks up to 400Gbps, i.e., the internal HPC/data centre (campus) network. This finding suggests that a potential link upgrade via the hyperconnectivity tender is feasible without any bottleneck of up to 400 Gbps. By extending or replacing the border element, centres can enhance their network capacity without encountering internal network constraints. Ensuring adequate internal capacity to leverage the upstream network upgrade is crucial for maximising the benefits of the enhanced network infrastructure. For capacities above 400 Gbps, the sites that do not face any issues are Leonardo, Meluxina and Vega, while the rest (from the reported ones) will need some upgrades. Furthermore, the firewalls can also act as bottlenecks if they don't support speeds higher than 400Gbps. This is (or will be) the case for EHPCPL and Karolina, if they don't proceed with an upgrade. Furthermore, also, the ones using firewalls which are capable of speeds greater than 400Gbps (Leonardo and Meluxina and MN5), would need to buy expensive licenses for such speeds. An alternative is the approach to combine both a firewall (e.g. for commercial

[45] In the case of Daedalus (under procurement) it is what is planned.

[46] Science DMZ (es.net) A Science DMZ is a portion of the network, built at or near the campus or laboratory's local network perimeter that is designed such that the equipment, configuration, and security policies are optimized for high-performance scientific applications rather than for general-purpose usage.

traffic flows) and a Science DMZ (e.g. for research traffic flows). This will significantly reduce the firewall burden and reduce (or even eliminate) the firewall upgrade/licensing costs.

These conclusions provide valuable insights into the current state of HPC centre networks and inform the design and procurement phases to ensure compatibility and optimal performance.

5.3.2 Connectivity Gaps

Summary of Gap Analysis for HPC Centres Connectivity

The analysis of the current and expected future connectivity for various HPC centres reveals several key insights. Most centres are selected for substantial upgrades to their connectivity by 2025, with some centres expecting to increase their bandwidth by several hundred Gbps. This is crucial for meeting the increasing demands of high-performance computing (HPC) applications and AI models, and in the near future also the AI factories.

The following table highlights the expected connectivity upgrades for 2025 for the Priority 1 institutions. Note that the link upgrade may take some time to be implemented and in several cases this may be completed in 2026. This will also depend on the timeline and especially the award of the next tender (which is expected to be concluded sometime early in 2025).

Detailed Analysis

Overall statistics

- Overall, the current average connection speed of the EuroHPC hosting entities is ~197 Gbps (without taking into account the sites for which the capacity is not yet known, such as several Quantum sites).
- The average projected speed procured in 2025 (to be implemented late 2025 or 2026) is ~327 Gbps.
- Thus, the projected increase for the procured access speeds is around 166%.

Significant Upgrades Expected:

- Several centres are planned to double or even quadruple their current bandwidth, such as the National Supercomputing Centre IT4Innovation (EuroHPC Karolina) from 100 Gbps to 400 Gbps, and the University of Minho (EuroHPC Deucalion) from 100 Gbps to 200 Gbps.
- The CINECA/LEONARDO site and the Jülich Supercomputing Centre (EuroHPC Jupiter) also expect significant upgrades from 200 Gbps to 400 Gbps.

Current Access Speed Discrepancies:

- One centre currently has relatively low bandwidth compared to their selected network class. This is the HPC centre Sofia Tech Park (EuroHPC Discoverer) that currently has 30 Gbps and expects to increase to 200 Gbps.

- On the other hand, the EuroHPC/BSC Quantum simulator, which will be using the old MN4 chapel site[47], with a current access speed of 20 Gbps, is not currently foreseen to change speed.

Planned Connectivity Enhancements:

- Centres like ICHEC (EuroHPC CASPIr) and (EuroHPC Quantum System) currently have no data on their access speed and are ranked as Class B, establishing significant connectivity by 2025 (200 Gbps).
- The GRNET/EuroHPC Daedalus currently has no specified connectivity but is based on reaching 400 Gbps.

Centres with High Current Bandwidth:

- The Linköping University (LiU) (EuroHPC Arrhenius) currently has a high access speed of 800 Gbps and plans to maintain this speed.

5.3.3 Network Services and Technology Gaps

In terms of network services' gaps, the main point is the end-to-end nature of the hyperconnectivity service, which spans across several administrative domains, namely the HPC/data sites, the national networks, and the pan-European network. Operating and monitoring the end-to-end hyperconnectivity network in an integrated manner, requires coordinated efforts crossing multiple administrative domains with different types of equipment and different vendors. This require using standard protocols and solutions, i.e. IP, for the general-purpose connectivity and VLANs or MPLS for layer-2 tunnels or VPNs. Regarding the hyperconnectivity network monitoring, high-level descriptions of the Service Level Agreements (SLAs) and the potential solutions for measuring the required parameters will need to be proposed (with some examples), allowing other equivalent solutions.

Regarding potential technology gaps, these have mainly to do with the available technologies to use high-capacity interfaces, i.e. towards Terabit level and beyond. Although the 800 Gigabit Ethernet standard was approved in February 2024[48], there are still limited commercial solutions in terms of transceivers and network interface cards, and these are mainly used by the top cloud companies[49]. Thus, the main gap in this case, is that in order to achieve hyperconnectivity speeds in the order of Tbps, multiple 100G or 400G have to be combined, either via link bundling or with clever application-level solutions (such as GridFTP-like[50]), which can use multiple links simultaneously. Furthermore, IEEE is also working on the 1.6 Tbps standard, which is currently expected to be approved mid-2026[51]. The above create some uncertainty, especially in a multidomain network, with different equipment at the end sites and national/EU ones. Still, effort will be made to use the available technologies and vendors'

[47] BSC About to Dispatch MareNostrum 5, for Critical Research - EE Times
[48] [802.3_B400G] IEEE P802.3df Standard Approved! (ieee802.org)
[49] A Quick Introduction to 800G Ethernet | FS Community
[50] GridFTP-Intro.pdf (eudat.eu): GridFTP is currently the de-facto standard for moving large files across the Internet. It can use multiple parallel links (e.g. 4x 100Gbps) to send files.
[51] Adopted IEEE P802.3dj Timeline (28 Nov 2023) (ieee802.org)

equipment at the time of the upcoming hyperconnectivity tender to create a high-class connectivity network.

Services' specifications for procuring network connectivity services for High Performance Computers (HPC) include:

1. **IP Connectivity:** Establish robust IP connectivity to ensure seamless communication between the HPC centres, data providers, user aggregation points, public cloud providers, and global internet. This service must support high-bandwidth data transfers, low latency, and high reliability to meet the demanding requirements of HPC workloads.
2. **IP Jumbo Frames for All Connections:** Implement support for IP jumbo frames (packets larger than the standard 1500 bytes, typically around 9000 bytes) for all network connections. This is crucial for improving the efficiency of data transfers between HPC centres, data providers, and end user aggregation points, including National Research and Education Networks (NRENs) or Internet Exchange Points (IXPs) for direct peerings.
3. **Performance Measurements:** Establish comprehensive performance measurements to monitor and guarantee the performance, availability, and reliability of network services. This includes metrics such as uptime, latency, packet loss, and jitter, with regular reporting to ensure compliance with the agreed SLA parameters.
4. **Committed Uptime:** Guarantee a minimum of 99.9% uptime, ensuring that network services are available for all but 8.76 hours per year. This high level of availability is critical for HPC environments where downtime can significantly impact research and computational tasks.
5. **Physical Access Highly Available:** Ensure highly available physical access to network infrastructure to support maintenance, troubleshooting, and upgrades. This may involve redundant access paths and/or equipment and swift response times to minimise downtime during physical interventions.
6. **Connectivity with Public Cloud Providers:** Provide direct connectivity to major public cloud providers (e.g., AWS, Azure, GCP, OVH) at a minimum capacity of 100-200 Gbps for each link in five locations in Europe. This would need to be achieved with direct peerings with such cloud providers. This ensures that HPC resources can leverage cloud services for additional computational power, storage, and hybrid cloud solutions.
7. **Rapid Provision of Dedicated Connectivity for Public Cloud Providers:** Enable the rapid provisioning of dedicated network connectivity to public cloud providers. This ensures rapid deployment and scaling of connections to AWS, Azure, GCP, and OVH, facilitating seamless integration with cloud services and resources. Examples of such solutions that can connect a Virtual Private Cloud to another cloud provider include Amazon AWS Direct Connect[52], Azure Express Route[53] and Google Partner Interconnect[54].
8. **Presence in Major IXPs:** Maintain a presence in at least five major Internet Exchange Points (IXPs) in Europe to facilitate efficient interconnection with other networks. It is preferable to have a

---

[52] Dedicated Network Connection - AWS Direct Connect - AWS (amazon.com)
[53] Azure ExpressRoute - Private Cloud Connections | Microsoft Azure
[54] Pricing | Cloud Interconnect | Google Cloud

presence in 15 countries and/or IXPs across Europe. This reduces latency and improves the overall performance of data transfers between the HPC/data centres and commercial networks[55].

9. **End-User Support (Levels 1, 2, 3):** Provide comprehensive end-user support across three tiers: Level 1 for basic troubleshooting and customer inquiries, Level 2 for more complex technical issues, and Level 3 for advanced engineering support. This ensures that all user issues are promptly and effectively resolved.
10. **Professional Services:** Offer professional services to assist with the planning, deployment, and optimisation of network connectivity solutions. This includes consulting, project management, and technical expertise to ensure the successful implementation of services.
11. **Reserved or Guaranteed Bandwidth (Destination Earth, CERN, etc.):** Ensure the provisioning of guaranteed on-demand bandwidth allocations for specific high-priority projects such as Destination Earth and CERN. These will be requested on a per-case basis following the relevant Professional Services engagement. This dedicated bandwidth supports critical research and scientific endeavours by providing consistent and reliable network performance.
12. **DDoS Protection Service:** Offer DDoS protection services, for the IP connectivity service provided. to safeguard HPC connectivity. Distributed Denial of Service attacks (DDoS)[56] service should include automated detection AND manual or automated mitigation mechanism against volumetric or more sophisticated DDoS attacks. This ensures the continued availability and performance of network services even during malicious traffic attempts.
13. **Automated Network Upgrade Process:** Establish an automated upgrade process for network infrastructure that is triggered when traffic consistently exceeds 60% of the 95th percentile. This proactive approach ensures that the network can scale dynamically to handle increased demand without degradation in performance.
14. **Latency Requirements:** Ensure that the average latency between any EU end sites and the country's R&E aggregation point is less than 15 milliseconds[57] (which is considered very good)[58]. This low latency is essential for maintaining the performance and responsiveness of HPC applications and services, including Quantum Computing.
15. **Direct Peerings with NREN Networks:** Establish direct peering arrangements with each country's National Research and Education Network (NREN). This enhances connectivity, reduces latency, and improves data transfer efficiency between research institutions and educational organisations.

The above services' specifications will be further analysed as part of the technical specifications.

[55] IXPs is a great solution to interconnect GÉANT/NRENs with SMEs/industrial entities that are usually connected to commercial ISPs. With multiple IXPs across Europe, the SMEs (e.g. AI companies) can reach fast and with low latency the EuroHPC and other HPC sites.
[56] What is a DDoS Attack? DDoS Meaning, Definition & Types | Fortinet; DDoS) Attack" is a cybercrime in which the attacker floods a server with internet traffic to prevent users from accessing connected online services and sites.
[57] The final numbers (between 10-15ms) will be part of the final specifications of the network.
[58] What Are Good Latency & Ping Speeds? | PingPlotter

5.3.4 Other Gaps

The proposed solution is expected to leverage GÉANT and the NREN networks, i.e. expanding the existing infrastructure of GÉANT & NREN networks. The vast majority of the end sites (HPC or data centres) that have been identified as potential beneficiaries of the hyperconnectivity network are currently connected to the NREN networks. Here, two important risks or gaps come into play:

- All end sites (HPC/data centres) included in the hyperconnectivity tender (60M Euros) must agree to be listed, which requires legal actions from the site signatory. This process can be time-consuming due to the site's decision-making process. Although the hyperconnectivity tender is 100% funded, requiring no contributions from the end sites, obtaining these signatures involves some bureaucracy that needs to be considered. Additionally, if the internal end site (campus) network has bottlenecks (see Section 5.1), investments from the end sites may be needed to fully exploit the hyperconnectivity network. This may delay the decision for the end site to participate in the tender.
- Some end sites may not agree to connect to providers other than the NRENs, such as commercial network providers or ISPs. Such cases have been reported during the consultation phase of the network. If the proposed legal instrument for implementing the hyperconnectivity network is based on an open tender or a similar open process, there is a risk that some sites may not agree to connect. In case, these sites are included as mandatory in the tender, there is a clear risk of the tender or the network implementation failing.

Furthermore, procuring the hyperconnectivity network for High-Performance Computers (HPCs) in Europe involves significant investment (the first tender is 60M Euros) and must adhere to stringent legal regulations.

6. Analysis of alternative scenarios

Four different possible solutions for delivering the hyperconnectivity network and the relevant services are analysed in this section:

- Solution A: Expanding the existing Infrastructure of GÉANT[59] / NRENs[60]. This solution utilises the GÉANT/NRENs infrastructure, which currently covers the vast majority of endpoints at sufficient speeds based on user feedback. This infrastructure will be expanded to meet the recorded needs of the hyperconnectivity network.
- Solution B: A service based on commercial IP Transit solutions. This solution utilises commercial networks and infrastructure, relying on IP transit services from commercial providers across Europe. The last-mile connections to the endpoints need to be provided via Internet Exchange Points. This solution requires the establishment of the necessary services and potential upgrades to the relevant infrastructure.

[59] GÉANT Network (geant.org)
[60] National Research and Education Networks (NRENs) - About GÉANT (geant.org)

- Solution C: Implementing a new dedicated hyperconnectivity network based on own fibre and IP/DWDM technologies. This solution involves a new long-term network specifically for the EuroHPC JU and the hyperconnectivity network. For this solution, the entire network needs to be constructed from scratch, including last-mile connections to end users, concentration nodes, core long-term fibre and wavelength leases, core nodes, service provisioning, and management infrastructure. The EuroHPC JU does not plan to implement network management capabilities for this network; therefore, the management and operation of the infrastructure should be provided as a service.
- Solution D: A hybrid solution: Expanding existing Infrastructure & introducing some network segments via commercial services. This solution is similar to solution A, but with the addition of some extra hyperconnectivity segments implemented directly via commercial service providers (and not via GÉANT/NRENs), either at the IP layer or the optical layer. In this approach, upgrades to the GÉANT/NRENs infrastructure are carried out by GÉANT/NRENs, while other network segments are developed through separate procurements from commercial providers, including both network and equipment suppliers.

These solutions were examined based on the following assessment criteria:

- **Fitness for purpose:** Adequacy of the proposed solution to fulfil the objectives of the HPC hyperconnectivity action.
- **Impact:** Contribution of the solution(s) to achieving the EuroHPC JU strategy.
- **Performance:** Guaranteeing and monitoring the performance and quality of the service (e.g. bandwidth, throughput, delay, latency, jitter, availability).
- **Security aspects:** Preventing unauthorised access, misuse, or denial of service, and detecting, monitor and remediating such events.
- **Support for innovation:** Accommodating innovative services and technologies (e.g., quantum communications and computing or on-demand network services if such requirements arise).
- **Sustainability:** Long-term maintainability of the proposed solution, including the ease with which it can be operated and supported, as well as its environmental footprint.
- **Cost efficiency:** Providing the best value-for-money to implement, operate, and maintain the solution.

The assessment criteria outlined above are further analysed into basic key elements. For each key element and solution under discussion, a mark is assigned based on the compliance of the specific solution.

To compare the different solutions, a relevant cost model has been developed. The cost model is based on prices collected through market analysis, with some predictions for higher speed rates, as the majority of hyperconnectivity related services are not widely available in the market today. The developed cost model acts as a budgetary tool to provide evidence and estimates for the cost of the relevant solutions.

6.1. Alternative Solutions

For the realization of the hyperconnectivity service, based on the requirement analysis and the gap analysis, the most appropriate way forward is a service-based solution. This was made clear by EuroHPC JU, EuroHPC JU is not planning to implement or develop a network management infrastructure and capacity, thus, the hyper-connectivity network needs to be provided and management by an external service provider. Nevertheless, an alternative solution was included in the analysis where a dedicated hyperconnectivity network will be implemented based on long term fibre and relevant infrastructure. Such a network should be implemented and managed by a third entity, providing relevant services for the hyperconnectivity network.

Based on the results of requirement and gap analysis a set of Services required by the hyperconnectivity network. These services should be the basis of any implementation and will allow to build the relevant specifications.

Four different possible solutions for delivering the hyperconnectivity network and the relevant services have been analysed:

- Solution A: Expanding existing Infrastructure of GÉANT/NRENs. A solution utilising GÉANT/NRENs infrastructure. Currently, this infrastructure covers the vast majority of the endpoints at sufficient speeds based on the user’s feedback. This infrastructure is expanded to cover the terabit policies of the hyperconnectivity network.
- Solution B: Commercial service based on IP Transit services. A solution utilising commercial networks and infrastructure. Such a solution is based on IP transit services from commercial providers around Europe. This solution needs to cover new last-mile connections to all the relevant end-user points of interest, services, and possible upgrades of relevant infrastructure.
- Solution C: Implementing a dedicated hyperconnectivity network. This solution is based on a new long-term network that is built for the purpose of the EuroHPC JU and the hyperconnectivity network. For this solution, the whole network needs to be built from the beginning, including last miles to end users, concentration nodes, core long-term fibre and frequency leases, core nodes, service provisioning and management infrastructure. As mentioned before, EuroHPC JU is not planning to implement network management capabilities for the reasons of this network, thus, management and operation of the infrastructure should also be provided as a service.
- Solution D: A hybrid solution expanding existing Infrastructure & introducing some network segments via commercial services. This solution is similar to solution A, with the addition of some extra connectivity services, either at the IP layer or the optical layer. In this solution though, the upgrades of GÉANT / NRENs infrastructure are implemented through GÉANT / NRENs, while the other network segments through separate procurements by commercial providers, including both network and equipment providers.

6.1.1 Common Services

The hyperconnectivity network, based on the user requirements and the relevant gap analysis, should provide the following set of services:

- IP Connectivity: Establish robust IP connectivity to ensure seamless communication between the HPC centres, data providers, user aggregation points, public cloud providers, and global internet. This service must support high-bandwidth data transfers, low latency, and high reliability to meet the demanding requirements of HPC workloads. Access speeds vary based on the end point tier – class, according to the following table

Tier	Phase A (2025) connectivity	Phase B (2028) connectivity	Phase C (2030+) connectivity
Class A	400-800 Gbps	800-1.2 Tbps	1.2-1.6 Tbps
Class B	200-400 Gbps	400-800 Gbps	800-1.2 Tbps
Class C	100-200 Gbps	200-400 Gbps	400-800 Gbps
Class D	10-100 Gbps	100-200 Gbps	200-400 Gbps

Table 1: End Points connectivity speeds and tier/class

- IP Jumbo Frames for All Connections: Implement support for IP jumbo frames (packets larger than the standard 1500 bytes, typically around 9000 bytes) for all network connections. This is crucial for improving the efficiency of data transfers between HPC centres, data providers, and end-user aggregation points, including National Research and Education Networks (NRENs) or Internet Exchange Points (IXPs) for direct peerings.
- Performance Measurements: Establish comprehensive Performance measurements to monitor and guarantee the performance, availability, and reliability of network services. This includes metrics such as uptime, latency, packet loss, and jitter, with regular reporting to ensure compliance with the agreed SLA parameters.
- Committed Uptime: Guarantee a minimum of 99.9% uptime, ensuring that network services are available for all but 8.76 hours per year. This high level of availability is critical for HPC environments where downtime can significantly impact research and computational tasks.
- Physical Access Highly Available: Ensure highly available physical access to network infrastructure to support maintenance, troubleshooting, and upgrades. This may involve redundant access paths and/or equipment and swift response times to minimise downtime during physical interventions.
- Connectivity with Public Cloud Providers: Provide direct connectivity to major public cloud providers (e.g., AWS, Azure, GCP, OVH) at a minimum capacity of 200 Gbps for each link in five locations in Europe. This ensures that HPC resources can leverage cloud services for additional computational power, storage, and hybrid cloud solutions.
- Rapid Provision of Dedicated Connectivity for Public Cloud Providers: Enable the rapid provisioning of dedicated network connectivity to public cloud providers. This ensures rapid deployment and scaling of connections to AWS, Azure, GCP, and OVH, facilitating seamless integration with cloud services and resources.
- Presence in Major IXPs: Maintain a presence in at least five major Internet Exchange Points (IXPs) in Europe to facilitate efficient peering with other networks. This reduces latency and

improves the overall performance of data transfers between the HPC centres and external networks.

- End-User Support (Level 1, 2, 3): Provide comprehensive end-user support across three tiers:
    - Level 1 for basic troubleshooting and customer inquiries,
    - Level 2 for more complex technical issues, and
    - Level 3 for advanced engineering support. This ensures that all user issues are promptly and effectively resolved.
- Professional Services: Offer professional services to assist with the planning, deployment, and optimization of network connectivity solutions. This includes consulting, project management, and technical expertise to ensure the successful implementation of services.
- Reserved or Guaranteed Bandwidth (Destination Earth, CERN, etc): Ensure provisioning of reserved or guaranteed on demand bandwidth allocations for specific high-priority projects such as Destination Earth and CERN. Those will be requested on a per case basis following relevant Professional Services engagement. This dedicated bandwidth supports critical research and scientific endeavours by providing consistent and reliable network performance.
- DDoS Protection Service: Offer DDoS protection services, for the IP connectivity service provided. to safeguard HPC connectivity. DDoS service should include automated detection AND manual or automated mitigation mechanism against volumetric or more sophisticated Distributed Denial of Service attacks. This ensures the continued availability and performance of network services even during malicious traffic attempts.
- Automated Network Upgrade Process: Establish an automated upgrade process for network infrastructure that is triggered when traffic consistently exceeds 60% of the 95th percentile. This proactive approach ensures that the network can scale dynamically to handle increased demand without degradation in performance.
- Latency Requirements: Ensure that the average latency between end sites and the country's R&E aggregation point is less than 15 milliseconds. This low latency is essential for maintaining the performance and responsiveness of HPC applications and services.
- Direct Peerings with NREN Networks: Establish direct peering arrangements with each country's National Research and Education Network (NREN). This enhances connectivity, reduces latency, and improves data transfer efficiency between research institutions and educational organisations

6.1.2 Solution A: Expanding existing Infrastructure of GÉANT/NRENs

An obvious solution for implementing the hyperconnectivity network and providing the relative services is by expanding the existing infrastructure of GÉANT and the NRENs. As spotted in gap analysis, there is marginal expansion that could be needed to achieve the relevant classes of service. It should be noted that the current connection speeds for the endpoints went through several rounds of validations, mainly with the NRENs, but in some cases, also with the end sites and are considered sufficient. Thus, any upgrades should consider both the current utilization of the network and the terabit policy initiatives by EuroHPC JU.

The following figures provide a detailed view of the network for Solution A.

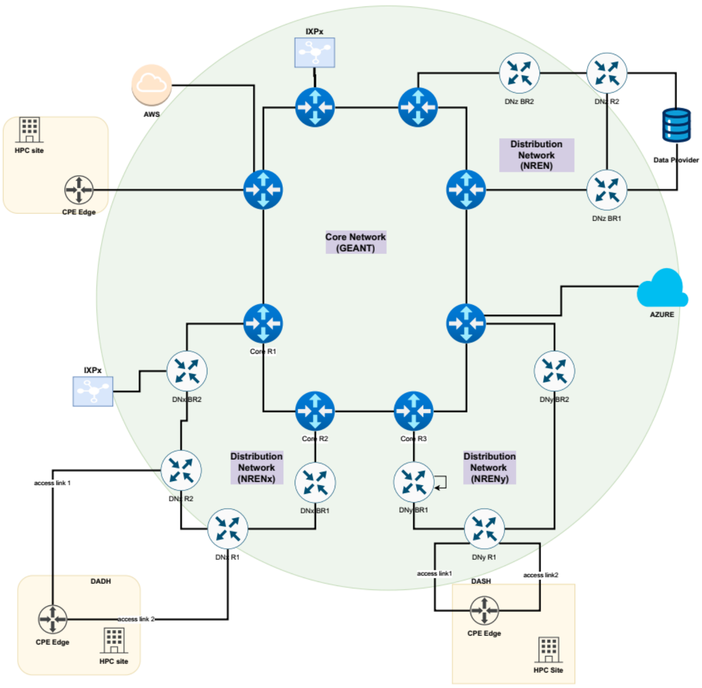

Figure 14: Architecture of Solution A utilizing GÉANT & NREN for the hyperconnectivity network

In this scenario the NREN provides the local/national network and GÉANT provides the international links. Every EuroHPC facility, included in the initial build, requires a last mile in order to be connected to the nearest NREN PoP and from there to the hyperconnectivity network. As in the majority of the cases this last mile is already in place sufficiently covering end point's requirements for bandwidth, from day zero existing HPC facilities can be given access to the hyperconnectivity network through multi 100G gigabit connections.

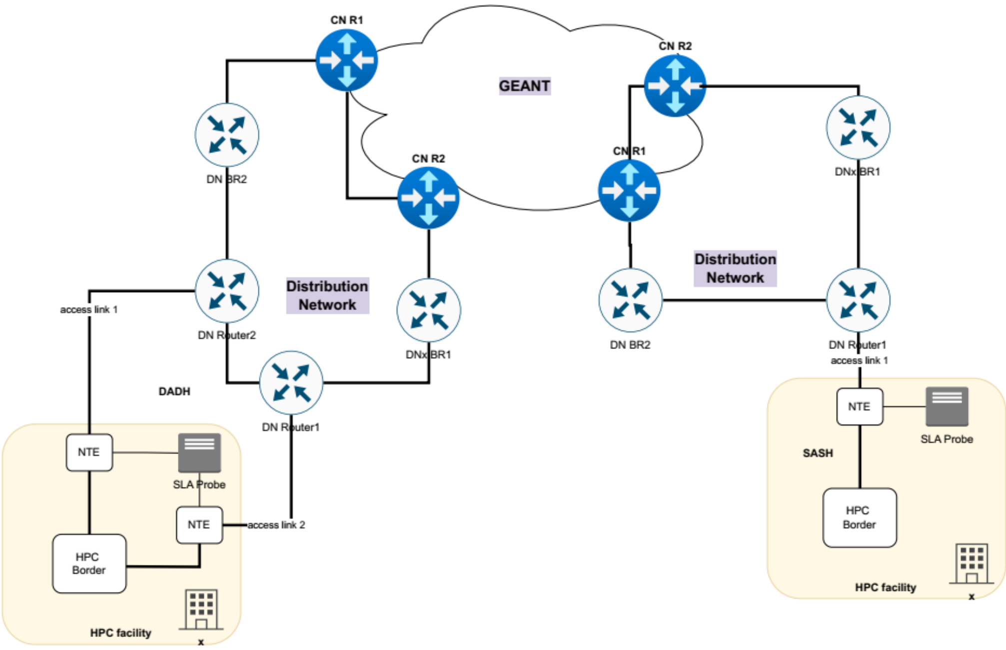

Figure 15: Architecture of Solution A with performance monitoring at end sites.

The strong points of this solution include:

- Enough resources provided from the very beginning after upgrading necessary end side equipment.
- A single cooperative close related to the European research community and HPC providers.
- Extra capacity can be easily allocated through upgrades of equipment and marginal upgrades of links between nodes.
- An already operating streamlined operation, providing support and SLAs.
- A proved high-capacity unified backbone.
- Users are already connected.
- As marginal upgrades needed it should be the most cost-efficient solution.
- Minimal risk regarding the timeframe and the implementation of the network.
- The opportunities of this solution include:
- One stop shopping via GÉANT.
- Enlarging the user community with industry research multiplying the effects of the hyperconnectivity network.
- The hyperconnectivity network will provide advanced networking infrastructure both for the HPC related community and the whole research and educational community of Europe.
- The weak points of this solution include:
- GÉANT is a cooperation of all the NRENs providing networking services to research community of Europe and is not close to commercial practices (i.e. risk-taking during implementation or/and operation).

- As a cooperation of public bodies there might be additional administrative burdens that could delay implementation of specific requests.

The threats of this solution include:

- GÉANT is a cooperation of all the NRENs of Europe based on a long-forged agreement. The introduction of hyperconnectivity network as a pan European service might introduce friction points that need to be resolved and settled.

Estimating the budget for this solution involves costing the necessary upgrades of existing access and backbone network as well infrastructures and administration teams of the GÉANT's – NREN network Cost of the solution. These costs were estimated through the costing methodology described previously. Speeds were upgraded for the main points of interest in priorities 1 to 3, while for the rest existing speeds were used that cover users' needs based on our questionnaires, interviews, and requirement analysis, presented in requirements and gap analysis. Based on this hypothesis price tags were estimated for the first period between Q3/2025 and Q2/2029 and for the second period between Q3/2029 and Q4/2032. The relevant results are presented in the following table.

	Stage 1 (2025 to 2029 in M€)	Stage 2 (2029 to 2032 in M€)
GÉANT	21,75	28,03
NREN – Backhauling (towards backbone network)	32,10	45,45
NREN Last mile (towards access network /end site)	7,05	9,98
Total	**60,9**	**83,46**

Table 2: Price tag for Solution A

6.1.3 Solution B: Commercial service based on IP Transit services.

IP transit is a service allowing a customer to access and cross (transit) an IP network. Usually, it is required to give a smaller Internet Service Provider (ISP) access to another bigger ISP (upstream provider) and eventually the broader internet, which is a network of networks. An IP transit service enables this by using the Border Gateway Protocol (BGP[61]) protocol, which exchanges routing information between Autonomous Systems (AS[62]), thereby providing access to the rest of the internet. The service is very similar with IP peering between providers – ISPs, but IP Transit comes with a cost. It should be noted that, in the case of direct peering between two providers, they exchange only their own customer networks information for mutual benefit, and as a result, there is no service cost for this service, and no universal access to the Internet.

[61] https://en.wikipedia.org/wiki/Border_Gateway_Protocol
[62] https://en.wikipedia.org/wiki/Autonomous_system_(Internet)

Solution B is based on commercial providers and is similar to solution A, requiring the same list of HPC and data points and the list of Internet Exchange Points (IXPs) that needed for the commercial ISPs to interconnect locally (in each country) to all research and academic institutions, which are connected to the NRENs. Relevant services should also include last mile connections to and from the end points with the required access speeds. This solution follows the DestinE architecture which is based on a commercial provider and an IP transit service.

A major issue that this solution should overcome is multihoming. The end points are already connected to the NRENs networks and some of them are having IP address spaces provided by the NRENs. In this case a multihoming architecture should be followed that increases complexity and requires special treatment during incident and performance management.

Market based solutions are technologically feasible provided that the infrastructure does not suffer from congestion. Market based solutions should provide congestion monitoring tools as described in the basic services of the network. Security focused offerings are a natural fit for market-oriented solutions which are normally catered for enterprise requirements.

The strong points of this solution include:

- Established security solutions as a commodity on market offerings.
- Penalty based SLA solution are common on market-based offerings.
- An already operating streamlined operation, providing support and SLAs
- Flexibility in management through single point of contact and relevant management entity[63].
- Easier commercial access to the HPC centres.

The opportunities of this solution include:

- Resources could be used to enforce collaboration between network providers in Europe.
- Upgrades in backbone networks could also benefit other user groups

The weak points of this solution include:

- Commercial networks are designed and operated for providing services mostly towards home and business users.
- Not suitable for the HPC traffic patterns (overbooked commercial networks vs overprovisioned academic networks)
- Likely could not support Jumbo frames (MTU=9000 bytes) end to end.
- IP Addressing issues (many HPCs already use PA IP space provided by NRENs)
- Not able to use existing last mile circuits (typically IRU based)
- GÉANT’s infrastructure is not utilized.
- Increased network configuration and OAM complexity with each HPC centre having multiple upstream providers.

[63] Although a single commercial provider is unlikely to be able to service all the end points, a single management interface is expected.

- Many sites (EuroHPC and others) may refuse to connect to commercial providers (i.e. not being NRENs) (administrative)
- Interconnection Research networks outside Europe will be through commercial peerings.
- Might involve a long implementation period for Last miles.
- Providers are not used into introducing advanced architectures and technologies that might put the infrastructure and other clients into risk.

The threats of this solution include:

- Commercial providers might use the opportunity to lock on the EuroHPC clients.
- Education of the research community in operating such services will be reduced or even eliminated, leading into future capacity deficit in advanced networking for HPC applications.
- The number of end point connections, more 350 in total, within multiple administrative authorities will be very difficult to implement in the tight implementation timeframe introduced by the funding scheme.

Estimating the budget for this solution involves costing the necessary services, “last mile” and “IP transit” services, as well infrastructures and administration for provided advanced services. These costs were estimated through the costing methodology described previously. The same hypothesis was used for the relevant services – speeds for each point of interest. Based on this hypothesis, price tags were estimated for the first period between Q3 2025 and Q2 2029 and for the second period between Q3 2029 and Q4 2032. The cost for the main (mandatory) PoIs was estimated and the secondary (optional) ones. The relevant results are presented in the following table:

	Stage 1 (2025 to 2029 in M€)	Stage 2 (2029 to 2032 in M€)
Core network (main PoIs)	142,35 €	191,22 €
Access network (main PoIs)	31,25 €	41,98 €
Sub-total	**173,60 €**	**233,20 €**
Core network (secondary PoIs)	77,90 €	108,65 €
Access network (secondary PoIs)	17,10 €	23,85 €
Sub-total	**95,00 €**	**132,50 €**
Total	**268,60 €**	**365,70 €**

Table 3: Price tag for Solution B

If one compares the cost of Solution B with Solution A, for Stage 1, it is around 3 times (285%) more expensive for the main points, and more than 4 times (441%) more expensive for all the points (main and secondary).

If one compares the cost of Solution B with Solution A, for Stage 2, it is around 3 times (279%) more expensive for the main points, and again more than 4 times (438%) more expensive for all the points (main and secondary). The final numbers may be slightly updated based on the final list of PoIs that will be agreed with the EuroHPC JU.

6.1.4 Solution C: Implementing a dedicated hyperconnectivity network.

A solution also for EuroHPC JU is also to try to implement a separate hyperconnectivity network based on long term leases. This solution could result into duplication of the network infrastructure of GÉANT as the majority of the countries should be connected in order for just the EuroHPC infrastructure to be interconnected. Thus, the network architecture could follow the relevant architecture of GÉANT as presented in a previous chapter.

In this solution, the EuroHPC sites could be the core nodes of an extended wavelength (alien wavelengths) and dark fibre network on which national Data providers and HPC sites could be interconnected. This would require multiple redundant routes between the EuroHPC sites for resilience and traffic management reasons. Each core node should facilitate the necessary DWDM equipment that could allow multiple terabit connections to be implemented based on the capacity requirements.

This solution will definitely result into a very long implementation and a financially demanding project, at least for the phase of initial build of the project.

As the majority of the sites will remain connected through the NREN's networks there will be configuration and troubleshooting issues with multihoming, as described in scenario B.

The strong points of this solution include:

- Allows more control to JU
- Long term solution as infrastructure should be made available for at least 15 years.
- Fully compliance with the priorities of JU
- Easier commercial access to the HPC centres.

The opportunities of this solution include:

- Allow the duplication of the research network in Europe providing alternative solutions in the future.

The weak points of this solution include:

- It is a very demanding solution in resources, including time, administrative and engineering resources, and budget.

  - The only other relevant solution is GÉANT's network that took more than 20 years to implement and required the extended co-operation between the NRENs and research communities.
  - GÉANT's infrastructure is not utilized.
  - Increased network configuration and OAM complexity with each HPC centre having multiple upstream providers.
  - It is not clear how non EuroHPC sites could be included, i.e. may refuse to connect (administrative)
  - Interconnection Research networks outside Europe will be through commercial peerings.
  - Might involve a long implementation period for Last miles.
  - Who is going to operate such an infrastructure. Dependence to a new Service provider.
  - Duplication of existing infrastructure (can be considered waste of public funding).
- What happens after 2032?
- The least cost-efficient solution as there is no infrastructure re-use foreseen.

The threats of this solution include:

  - There significant risks for this solution related with the tight implementation timeframe introduced by the funding scheme.
  - The management scheme that will emanate from this solution will not have the long experience of the previous solutions.

The main disadvantages of this solution are its complexity and the extensive network that needs to be implemented. Based on the extensive size of the network to be implemented (that would require a separate small project to calculate it), the exact cost of this solution was not estimated, but it is expected that it would be significantly higher and more than an order of magnitude more expensive than the scenario A[64].

6.1.5 Solution D: A mixed solution expanding existing Infrastructure & introduce commercial services

A final solution examined is a hybrid solution expanding existing Infrastructure & introducing commercial services for segments of the network. This solution is similar with solution A, with the addition of some extra connectivity services, either at IP layer or optical layer. In this solution though, the upgrades of GÉANT/NRENs infrastructure are implemented through GÉANT/NRENs, while the other network segments are implemented through separate procurements by commercial providers, including both network and equipment providers.

Thus, this solution, will have similar characteristics with solution A, but increased complexity in managing and implementing it, as it will require a continuous co-operation between GÉANT, the

[64] This will depend on the actual number of points to be interconnected; still, basic (if not full-blown) NREN networks would be needed in each country and the pan-European backbone. The estimation has been informed by the budget numbers in the Compendium-2022-2023-IX.pdf (geant.org), taking into account both the numbers from NRENs and also GÉANT.

NRENs, EuroHPC JU and any commercial providers that will result through a separate bidding process. Furthermore, the cons of the commercial solutions (B and C) apply for the commercial network segments, including the increased cost. The cost of this solution is not analysed in detail, as it will depend on the portion of the network based on Solution A and on the portion of the network based on Solution B or C. For example, suppose half of the network is based on Solution A (GÉANT /NRENs) and half on Solution B (IP transit from commercial providers). In that case, the hybrid solution cost will be exactly in the middle between Solution A and B. Based on the analysis in Solutions A and B, this would mean the following indicative prices:

	Stage 1 (2025 to 2029 in M€)	Stage 2 (2029 to 2032 in M€)
Mixed approach (half network)	164,75 €	224,58 €

It is evident that this solution is much less cost-effective compared to Solution A (271% more expensive for Stage 1 and 269% more expensive for Stage 2). So, only limited implementation of network segments directly via commercial service providers should be considered and for special cases. A possible use of this solution would be that Solution A is used, with the exception of one or two other segments. For example, Destination Earth, as part of its data (cloud) architecture, are subcontracting connectivity to a commercial network provider. Since Destination Earth is/will be using heavily several EuroHPC hosting sites, a hybrid solution is envisaged for them (i.e. both GEANT/NRENs and commercial provider connectivity). As the DestinE current connectivity contract ends mid-2026, such a hybrid scenario can be considered in time for this period.

6.2. Assessment Results

Based on the Assessment methodology described above, each one of the solutions from the previous chapter was evaluated. The following table provides a descriptive overview of the results from this process. It highlights the strengths and weaknesses of each solution, as well as the relevant opportunities and threats, and assesses the performance of each solution in relation to each key element.

Criteria	Key Elements	Solution A
Fitness for purpose		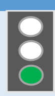
	Overprovisioning in the core network	Fully support relevant key element
	Support of Jumbo Frames	Fully support relevant key element
	Peering with international research networks and infrastructures	Fully support relevant key element

	Limiting implementation and timeline risks	Fully support relevant key element
Impact		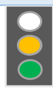
	Hyperconnectivity at terabit level	Fully support relevant key element
	Support of a Federated Supercomputing Infrastructure	Fully support relevant key element
	Promoting a World-Class European HPC and Quantum Computing Ecosystem	Fully support relevant key element
	Open and Inclusive Ecosystem	Partial support relevant key element, as there are specific requirements for the usage of the relevant resources (i.e. there is an Acceptable Usage Policy (AUP) for research and educational purposes)
Performance		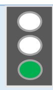
	Pro-active Based Measures & Guarantees	Fully support relevant key element
	Applicability of Efficient Performance Monitoring Mechanisms	Fully support relevant key element
	Well-defined Upgrade Rules for Congested Infrastructure	Fully support relevant key element
Security		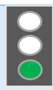
	Preventing Attacks on HPC Infrastructure	Fully support relevant key element
	Balance Between Performance and Security	Fully support relevant key element
	Proven DDoS Mitigation	Fully support relevant key element

	Facilitating Innovative Security Architectures	Fully support relevant key element
Support for Innovation		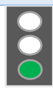
	support advanced networking initiatives	Fully support relevant key element
	support an active vibrant research community	Fully support relevant key element
	applying innovative technologies	Fully support relevant key element
	training and educational Initiatives	Fully support relevant key element
Sustainability		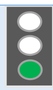
	Utilizing existing management infrastructure	Fully support relevant key element
	Reduced Complexity and OAM Cost	Fully support relevant key element
	Proven Service Provisioning	Fully support relevant key element
	Energy Efficiency	Fully support relevant key element

Table 4: Assessment of Solution A.

Criteria	Key Elements	Solution B
Fitness for purpose		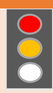
	Overprovisioning in the core network	Does not support relevant key element, commercial networks are overbooked networks best fitted for home and business users.

	Support of Jumbo Frames	Does not support relevant key element, end to end support of jumbo frames not possible.
	Peering with international research networks and infrastructures	Does not support relevant key element, peering will be implemented through traditional commercial peering agreements.
	Limiting implementation and timeline risks	Fully support relevant key element
Impact		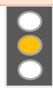
	Hyperconnectivity at terabit level	Fully support relevant key element
	Support of a Federated Supercomputing Infrastructure	Fully support relevant key element
	Promoting a World-Class European HPC and Quantum Computing Ecosystem	Does not support relevant key element, commercial networks are not open to new architectures and technologies.
	Open and Inclusive Ecosystem	Fully support relevant key element
Performance		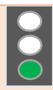
	Pro-active Based Measures & Guarantees	Fully support relevant key element
	Applicability of Efficient Performance Monitoring Mechanisms	Fully support relevant key element
	Well-defined Upgrade Rules for Congested Infrastructure	Fully support relevant key element
Security		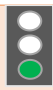
	Preventing Attacks on HPC Infrastructure	Fully support relevant key element

	Balance Between Performance and Security	Fully support relevant key element.
	Proven DDoS Mitigation	Fully support relevant key element.
	Facilitating Innovative Security Architectures	Fully support relevant key element.
Support for Innovation		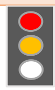
	support advanced networking initiatives	Fully support relevant key element.
	support an active vibrant research community	Does not support relevant key element, commercial providers focus on efficiently providing services to end users
	applying innovative technologies	Does not support relevant key element, commercial networks while applying innovative technologies are not focusing in deploying platforms for collaboration or HPC related initiatives.
	training and educational Initiatives	Does not support relevant key element, commercial providers are not providing actively platforms for experimentation and research.
Sustainability		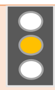
	Utilizing existing management infrastructure	Does not support relevant key element, while commercial providers have well established management teams, they mis the link with existing entities at HPC centres and research infrastructures.
	Reduced Complexity and OAM Cost	Fully support relevant key element.
	Proven Service Provisioning	Fully support relevant key element.

	Energy Efficiency	Fully support relevant key element.

Table 5: Assessment of Solution B.

Criteria	Key Elements	Solution C
Fitness for purpose		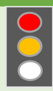
	Overprovisioning in the core network	Fully support relevant key element
	Support of Jumbo Frames	Fully support relevant key element
	Peering with international research networks and infrastructures	Fully support relevant key element
	Limiting implementation and timeline risks	Does not support relevant key element, implementing a dedicated DWDM based network for EuroHPC JU is a high-risk project in terms of both budget and timeframe.
Impact		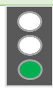
	Hyperconnectivity at terabit level	Fully support relevant key element
	Support of a Federated Supercomputing Infrastructure	Fully support relevant key element
	Promoting a World-Class European HPC and Quantum Computing Ecosystem	Fully support relevant key element
	Open and Inclusive Ecosystem	Fully support relevant key element
Performance		
	Pro-active Based Measures & Guarantees	Fully support relevant key element

	Applicability of Efficient Performance Monitoring Mechanisms	Fully support relevant key element
	Well-defined Upgrade Rules for Congested Infrastructure	Fully support relevant key element
Security		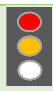
	Preventing Attacks on HPC Infrastructure	Fully support relevant key element
	Balance Between Performance and Security	Fully support relevant key element
	Proven DDoS Mitigation	Does not support relevant key element, implementing DDoS mitigation methods and technics requires experience and management methods that require time and experience.
	Facilitating Innovative Security Architectures	Fully support relevant key element
Support for Innovation		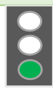
	support advanced networking initiatives	Fully support relevant key element
	support an active vibrant research community	Fully support relevant key element
	applying innovative technologies	Fully support relevant key element
	training and educational Initiatives	Fully support relevant key element
Sustainability		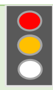

	Utilizing existing management infrastructure	Does not support relevant key element, as new infrastructures need to be deployed.
	Reduced Complexity and OAM Cost	Fully support relevant key element
	Proven Service Provisioning	Does not support relevant key element, as new service management and provisioning need to be implemented.
	Energy Efficiency	Fully support relevant key element

Table 6: Assessment of Solution C.

Criteria	Key Elements	Solution D
Fitness for purpose		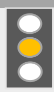
	Overprovisioning in the core network	Fully support relevant key element
	Support of Jumbo Frames	Fully support relevant key element
	Peering with international research networks and infrastructures	Fully support relevant key element
	Limiting implementation and timeline risks	Does not support relevant key element, based on the complex implementation and procurement scheme of the solution.
Impact		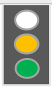
	Hyperconnectivity at terabit level	Fully support relevant key element
	Support of a Federated Supercomputing Infrastructure	Fully support relevant key element

	Promoting a World-Class European HPC and Quantum Computing Ecosystem	Fully support relevant key element
	Open and Inclusive Ecosystem	Partial support relevant key element, as there are specific requirements for the usage of the relevant resources (i.e. for research and educational purposes)
Performance		
	Pro-active Based Measures & Guarantees	Fully support relevant key element
	Applicability of Efficient Performance Monitoring Mechanisms	Fully support relevant key element
	Well-defined Upgrade Rules for Congested Infrastructure	Fully support relevant key element
Security		
	Preventing Attacks on HPC Infrastructure	Fully support relevant key element
	Balance Between Performance and Security	Fully support relevant key element
	Proven DDoS Mitigation	Fully support relevant key element
	Facilitating Innovative Security Architectures	Fully support relevant key element
Support for Innovation		
	support advanced networking initiatives	Fully support relevant key element
	support an active vibrant research community	Fully support relevant key element

	applying innovative technologies	Fully support relevant key element
	training and educational Initiatives	Fully support relevant key element
Sustainability		
	Utilizing existing management infrastructure	Fully support relevant key element
	Reduced Complexity and OAM Cost	Partial support relevant key element, as more contracts and peers will be involved in the process.
	Proven Service Provisioning	Fully support relevant key element
	Energy Efficiency	Fully support relevant key element

Table 7: Assessment of Solution D.

In terms of cost efficiencies solutions are ranked from more efficient to less efficient as follows:

- Solution A is the most efficient as utilization of GÉANT's and NREN infrastructure reduces significantly the relevant costs.
- Solution D is similar with Solution A, but introduces a significant extra administration and management overhead for all parties (i.e. GÉANT – NRENs, commercial network and equipment providers and EuroHPC JU, which has to deal with two sets of parties).
- Solution B follows with significant costs for implementing the services and minimal reuse of public funded infrastructure, while at the same time is not creating facilities capable to reduce the cost for the hyperconnectivity network in the long-run.
- Solution C is the least efficient as it replicates existing network infrastructures implemented within the last 20 years by the NRENs and GÉANT. On the other hand, this infrastructure could reduce the relevant costs, compared with Solution B, but not compared with Solution A or/and D.

The results of the assessment analysis for all four solutions are summarised in the following table.

Assessment Analysis Criteria	Solution A	Solution B	Solution C	Solution D
Fitness for purpose				

Impact	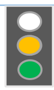	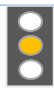	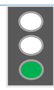	
Performance	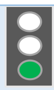		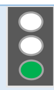	
Security	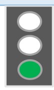		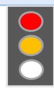	
Support for Innovation	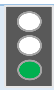		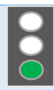	
Sustainability	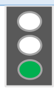	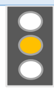	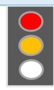	
Cost Efficiency	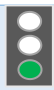		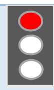	

- Concluding the results of the assessment analysis, Solution A: Expanding the current infrastructure-specifically leveraging the existing GÉANT and NRENs- emerges as the most suitable solution for implementing the Hyperconnectivity Network. GÉANT and the NRENs hold a technically unique position, being the best solution with a presence across all end points or Points of Interest (PoIs). In other words, NRENs interconnect the end points (HPC and Data Centres) directly, while commercial providers do not connect to the end points directly, rather they have to use Internet Exchange Points to reach these end points. They already provide services to all non-commercial PoIs, ensuring a broad and inclusive network. Additionally, through interconnected NRENs, GÉANT connects to approximately 10.000 institutions[65] across various scientific disciplines, which may be future HPC users.
- The infrastructure is based on publicly funded Indefeasible Rights of Use (IRU) for both core and access networks, offering a long-term and sustainable solution. Furthermore, several NRENs operate high-performance computing (HPC) and data infrastructures, underscoring their capability and reliability in managing complex and large-scale network operations.
- By choosing Solution A, the project can benefit from a fast implementation and deployment process. The existing infrastructure allows for seamless integration, minimising the time and resources required to get the project up and running. This approach is cost-efficient, leveraging established networks and reducing the need for significant new investments. It also ensures sustainability, as the infrastructure is designed for long-term use (15-20 years) and is supported by public funding. Operating costs are reduced due to economies of scale and shared resources within the existing network framework.
- Moreover, solution A supports innovation by utilising an established research network that fosters collaboration and development. It enables peering with major research networks worldwide, facilitating global connectivity and information exchange. The infrastructure also includes robust security measures and Denial of Service (DoS) mitigation services, ensuring

[65] About GÉANT - About GÉANT (geant.org)

network protection and reliability. Existing federated services, such as Trust and Identity Services, are already in place, providing a secure and efficient means of managing user identities and access across the network.

- Solution A also promotes bottom-up approaches, addressing a diversity of users' needs and corresponding data sites, and provides tailored expertise for various application domains concerned (e.g., Copernicus; e-Health, quantum). This requires a unique governance and management culture to maintain end-users trust. The GÉANT and NRENs network is inherently decentralized and scalable due to its technical design, with NRENs serving each country. Additionally, the network has a socio-economic driver perspective, aiming to bridge the core-periphery digital divide between potential users in central and western Europe with eastern and widening countries.
- It is important to note that for the hyperconnectivity network to provide services to an end site, an agreement with the end site's administration is required. This process can be lengthy and cumbersome. However, existing infrastructure, cooperation, and connectivity through the NREN's infrastructure at these sites are expected to minimise any implementation and timeframe risks.
- Many HPC Centres currently operate through Provider Aggregatable (PA) IP address space[66] provided by NRENs. Transitioning away from this IP space would be complex and could impact implementation in any alternative solution.
- Finally, the estimated cost of Solution B is approximately 3-4.5 times higher than Solution A, depending on the number of endpoints connected via IP transit services provided by commercial providers. The cost of Solution C is estimated to be more than an order of magnitude higher than Solution A, as solution A leverages long-term investments in dark fibre and equipment. Technical disadvantages are further analysed in Section 5.
- This comprehensive suite of services makes the expansion of the current infrastructure a strategic and advantageous choice for the project.

7. Technical Specifications

The technical specifications of the relevant network and services is a set of specifications required to implement and provide the hyperconnectivity network. These specifications focus on a service-oriented network based on the outcome our analysis presented above.

The end sites identified and participating in the hyperconnectivity network are grouped in into the following categories and priorities:

a. Priority 1: This category includes the EuroHPC hosting sites, including the planned ones.
b. Priority 2: This category covers important national HPC centres and data providers from which HPC users retrieve their data, making them of vital importance.
c. Priority 3: This group includes important national HPC centres and data providers identified through the HPC users' questionnaire, with data transfer requirements exceeding 10 TB per day.

[66] What are Provider Aggregatable (PA) addresses and Provider Independent (PI) addresses? — RIPE Network Coordination Centre

d. Priority 5: This category includes other national centres and data providers identified through the questionnaires, interviews, and research conducted during Tasks 1 and 2.
e. Priority 6: This category includes national centres and data providers identified in the HPC users' questionnaire with data transfer requirements of less than 10 TB per day.
f. Priority 10[67]: This group comprises commercial end sites (mainly large industries or SMEs) that either presented HPC needs in the questionnaires or were identified during Tasks 1 and 2, or other end sites that are not relevant for the first hyperconnectivity tender. The commercial sites are expected to be interconnected via Internet Exchange Points (IXPs).

The specifications provide detailed descriptions and relevant information for the following components of the network:

- Creation of the interconnection framework: This involves defining the basic characteristics of the hyperconnectivity network.
- Interconnection of Major points belonging to the following categories:
  - ***Academia User base aggregation points***: These are the NREN (National Research and Education Networks) in each country.
  - ***Commercial User base aggregation points***: These include major Internet Exchange Points (IXPs) in every EuroHPC JU hosting country, as well significant European IXPs such as AMSIX-DECIX-LINX-MIX [68] [69] [70] [71]. At each of these points, the contractor must provide both IP connectivity and L2 interconnection ports to the hyperconnectivity network.
  - ***Major Public Cloud infrastructures:*** This includes platforms such as Amazon AWS[72], Microsoft Azure[73], Google GCP[74] and Oracle Cloud[75].
  - ***Major international Research Networks***: These include networks outside Europe, such as ESnet[76], Internet2[77].
- Interconnection of Priority 1 Points of Interest (PoIs): A total of 21 PoIs will be interconnected at the initial speeds specified for 2025 in Gbps.
- Interconnection of Priority 2 Points: A total of 135 PoIs will be interconnected at the initial speeds specified for 2025 in Gbps.
- Interconnection of Priority 3 Points: A total of 26 PoIs will be interconnected at the initial speeds specified for 2025 in Gbps.

---

67 The priorities are not arranged sequentially. During the initial development of the priorities, it was uncertain whether additional categories would be needed. In the end, the priorities were left as initially perceived, which explains the numbering gaps.
68 https://www.ams-ix.net/ams
69 https://www.de-cix.net/
70 https://www.linx.net/
71 https://www.mix-it.net/en/
72 Cloud Computing Services - Amazon Web Services (AWS)
73 Cloud Computing Services | Microsoft Azure
74 Compute Engine | Google Cloud
75 Cloud Compute | Oracle
76 US Energy Science Network - Home (es.net)
77 Home - Internet2

- Optional interconnection of Priority 5 and 6 points: A total of 151 PoIs may be interconnected, at the initial speeds specified for 2025 in Gbps.
- Upgrading interconnection speeds in 2028: The upgrade of interconnection speeds for the above points will be justified based on a specific upgrade algorithm (i.e. rule), subject to demonstrated needs.
- Backbone and aggregation layer upgrades: Upgrades for the backbone or aggregation layer are not directly addressed. However, it is required that the quality characteristics of the backbone network remain consistent, regardless of any upgrades at the access layer.

As part of the hyperconnectivity network, bespoke solutions for accessing the hyperconnectivity network may also be required. Candidates must include professional services in their proposals to design bespoke solutions tailored for special projects or infrastructures. If the implementation of an agreed bespoke solution involves the provisioning of new physical connectivity services -such as new access ports or dedicated point-to-point connectivity services- these services must be charged separately.

An initial set of Points of Interest (PoIs), i.e. end sites, has been identified, following the priorities outlined above. Most of these PoIs fall within the mandatory Priorities 1-3, and if feasible, from the optional Priorities 5-6. However, there is a small subset of optional PoIs (Priority 10) that will be connected via IXPs. In addition, a number of major cloud providers and Internet Exchange Points have been elevated to Priority 2 (mandatory). This change was made following stakeholder consultations, where potential users highlighted their reliance on cloud data services, from which data for HPC runs will be sourced, making it crucial for these providers to be interconnected to the hyperconnectivity network.

After the contract is signed, the EuroHPC JU will establish a validation process for the PoIs to ensure the following:

- Commitment of the end point to interconnect to the hyperconnectivity network.
- Eligibility of the end point, based on the funding requirements.

It must be noted that the EuroHPC JU reserves the right, during the implementation period, to replace one or more PoIs with alternatives not listed in the original plan, if such a change better serves its objectives. However, any replacement sites must be comparable in terms of location (country-wise or continental vs. offshore) and the required speed.

Candidates must also include a Performance Monitoring Platform (PMP) in their offer implemented as a service, to provide quantitative and qualitative metrics. This platform must deliver detailed and analytical measurements of the performance of the offered service. Relevant monitoring probes must be installed as follows[78]:

[78] For Priority 1 (of major importance), all sites must have a PMP point. For the other two priorities, which include a much larger set of points, only a fraction of PMP points is necessary to monitor service quality. The

- In all Priority 1 PoIs.
- In 10% of Priority 2 PoIs.
- In 5% of Priority 3 PoIs.
- In six (6) core nodes of the provider.
- In six (6) Major Internet eXchange Points (IXPs).

The exact Points of Interest (PoIs) falling under Priorities 2 and 3 will be determined at a later stage (e.g. prior the tender publication or during the implementation phase), considering the specific requirements and restrictions applicable to each PoI. Core Points of Presence (PoPs) and Internet Exchange Points (IXPs) must be proposed by the provider, based on available resources, and subsequently agreed upon the EuroHPC JU.

7.1. Funding Requirements

The estimated funding requirements for the proposed solution A -Expanding existing Infrastructure of GÉANT/NRENs- are presented in the following table.

	Stage 1 (2025 to 2029 in M€)	Stage 2 (2029 to 2032 in M€)
GÉANT	21,75	28,03
NREN – Backhauling (towards backbone network)	32,10	45,45
NREN Last mile (towards access network /end site)	7,05	9,98
Total	**60,9**	**83,46**

Table 8: Funding Requirements for Solution A

7.2. Funding Limitations

Funding for the hyperconnectivity network is secured through the CEF Digital (CEF2) funding program. The hyperconnectivity network falls under the Digital sector, which aims to enhance digital connectivity through high-capacity networks.

According to the eligibility criteria of the CEF2, all EU Member States are eligible to participate in the funding programs. This includes projects within their territories as well as those involving cross-border cooperation with other EU countries. Overseas Countries and Territories (OCTs), which are associated with EU Member States, are also eligible for CEF2 funding. Additionally, certain non-EU countries can participate in CEF2 projects if they are associated with the program. Notably, Moldova and Ukraine

lower the priority, the smaller the required fraction. A representative number of points has been selected for the backbone and IXPs.

have agreements allowing them to participate in CEF2 across all sub-programs, including transport, energy, and digital sectors [79,80,81,82].

It is important to note that, based on the CEF2 regulation for the Digital sector and discussions with EU personnel, entities from other countries that would typically be eligible under specific conditions are not eligible for the scope of the hyperconnectivity network[83] [84].

The funding for the network is secured in two (2) phases: one budget line of €60 million, which must be implemented by the end of 2029, and a separate budget line of €100 million, which must be implemented by the end of 2032, optionally extending the first implementation.

7.3. Analysis of alternative solutions for the hyperconnectivity procurement

As already mentioned, the following alternative solutions were analysed:

1. Expanding the existing infrastructure based on GÉANT & NREN networks
2. A service based on commercial IP transit solutions
3. Implementing a new dedicated hyperconnectivity network based on own fibre and IP/DWDM technologies.
4. A hybrid solution combining A and B (or C).

The alternative solutions were examined and assessed based on a set of techno-economic and impact criteria included in the tender specifications. After the scenarios comparative assessment, the best fit scenario was Scenario A – Expanding existing infrastructure (utilizing GÉANT and NRENs networks), as the existing infrastructure suitably expanded and adapted is fulfilling all the identified criteria elements.

These alternative solutions need to be further evaluated in light of the procurement options available to EuroHPC.

7.3.1 Overview of the various procurement procedures

Article 46 of the Financial Rules of the EuroHPC Joint Undertaking ("EuroHPC") provides that Title VII of Council Regulation (EU, Euratom) 2018/1046 shall apply to the procurement activities of EuroHPC.[85]

[79] https://cinea.ec.europa.eu/system/files/2022-10/corrected-eligibility_costs_under_cef.pdf
[80] https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/cef/wp-call/2023/call-fiche_cef-e-2023-pci_en.pdf
[81] https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/cef/guidance/list-3rd-country-participation_cef_en.pdf
[82] https://www.euro-access.eu/en/programs/60/Connecting-Europe-Facility-Digital
[83] https://cef-hellas.gr/images/CEFDIGITALArticles/2nd_Call_Digital_2022/12_call-fiche_cef-dig-2022-5gcorridors_en.pdf
[84] https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/cef/wp-call/2024/call-fiche_cef-t-2024-afifgen_en.pdf
[85] With some exceptions as specified in the same Article 43, which are not applicable to the case at hand, as well as any specific provisions of the constituent act or the basic act of the programme the implementation of which is entrusted to the EuroHPC JU.

Council Regulation (EU) 2018/1046 (the "Financial Regulation") is therefore the relevant legal framework that must be considered in order to analyse the procurement procedure options available for EuroHPC.

The following main options for the corresponding follow-up tender(s) have been identified, which needs to be procured starting from the first one this year, namely:

- use of the 'open' procedure;
- use of the 'restricted' procedure;
- use of the 'competitive procedure with negotiation' procedure;
- use of the 'competitive dialogue' procedure;
- use of the 'negotiated procedure without prior publication of a contract notice'.[86]

A. The 'open' procedure

The open procedure is the most standard procurement method used by European Institutions which are subject to the Financial Regulation. It enables them to acquire works, supplies, and services in a manner that is both straightforward, open, and competitive.

The open procedure is a 'one-step' procedure in which the contract opportunity is published in the Official Journal of the European Union ("OJEU") and submission of a tender response is made accessible electronically to all interested bidders. Potential bidders are required to submit both selection (qualification) information and final proposals simultaneously in response to the tender invitation. Only tenders that meet the selection criteria will be considered further.

The minimum time limit for receipt of tenders is 37 days following the publication of the OJEU contract notice. However, the duration of the entire process, from preparation of the initial tender documents (including a draft contract) to contract award, is much longer and, depending on the availability of procurement professionals, it may last up to one year or even longer for more complex procurements such as the one for the hyperconnectivity.

In contrast to the other procurement procedures, the open procedure does not allow contracting authorities to limit the number of suitable candidates to be invited to participate in the procedure and all bidders are welcome to submit a complete tender in response to the OJEU contract notice (the call for competition).

The open procedure is often used when the required supplies, services or works are not particularly complex. The contracting authority simply advertises its specific requirements and invites applicants to bid.

[86] Relying on Article11.1 (l) "for the purchase of public communication networks and electronic communications services within the meaning of Directive 2002/21/EC of the European Parliament and of the Council". See ANNEX I of Financial Regulation 2018/1046.

The open procedure is mostly suitable for simple and straightforward requirements. Under the open procedure no negotiations with bidders are allowed,[87] implying that the contracting authority must be able to design detailed technical specifications upfront (to be defined in the tender documents). These specifications cannot be changed or finetuned during the procurement process or afterwards, and there is no possibility to develop more innovative solutions with bidders during the tender process.

B. The 'restricted' procedure

Alongside the open procedure, the restricted procedure represents another standard procurement method used under the Financial Regulation. Unlike the open procedure, the restricted procedure is a 'two-step' procedure, whereby the contracting authority publishes an OJEU contract notice, together with initial tender documents (a 'selection' document which includes exclusion and selection criteria) that set out minimum technical, economic, and financial criteria. Bidders submit requests to participate and must provide information in response to the selection document. The contracting authority then performs an initial evaluation to establish which of the bidders are qualified to perform the contract and those will be selected to be invited to submit a complete tender. The contracting authority is permitted to limit the number of bidders that it invites to tender which means that not all of the bidders that qualify will automatically be invited to tender (although it may be the case). Only those bidders that have been pre-selected will be issued with the full tender documents and invitation to submit complete and final tenders.

The minimum time limit for receipt of requests to participate is 32 days from the publication of the OJEU contract notice. The minimum time-limit for receipt of final tenders is 30 days from the sending of the invitation to tender to selected candidates. Like the Open Procedure, the duration of the entire process, from preparation phase to contract award, is much longer and, depending on the availability of procurement professionals, may last up to one year or even longer for more complex procurements such as the one for the hyperconnectivity.

Similar to the open procedure, no negotiations are permitted with bidders during this process, implying that EuroHPC must be able to design detailed technical specifications upfront (to be defined in the tender documents) with no flexibility to change or refine them during the procurement process and with no possibility to discuss any of the proposed solutions with bidders.

C. The 'competitive procedure with negotiation' procedure

The competitive procedure with negotiation is a 'two-step' process, similar to the restricted procedure. The process starts with the publication of an OJEU contract notice, and procurement documents are made accessible electronically. The first step involves assessing the exclusion and selection criteria (minimum technical, economic, and financial criteria). Any bidder may express an interest and request

[87] Contracting authorities may clarify aspects of the tender with bidders, but this possibility is very limited.

to participate, but only those shortlisted by the contracting authority are invited to the next phase in which they submit their proposals.

In the second step, at least three (3) shortlisted bidders are asked to prepare initial proposals. The contracting authority then engages in negotiations with these bidders based on their proposals, focusing on both technical and financial aspects of the tenders. The negotiation phase allows fine-tuning and adjustments to the initial proposals before bidders are invited to submit their final proposals for evaluation.

The minimum time limit for receipt of requests to participate is 32 days from publication of an OJEU contract notice. The minimum time limit for receipt of tenders is 30 days from sending the invitation to tender to selected candidates. In practice, the duration of the entire process, from preparation phase to contract award, is likely to take at least 12 months likely more than this for more complex procurements such as the one for the hyperconnectivity (depending on EuroHPC level of experience in conducting such a procurement procedure). In particular, the preparation of the tender process and the relevant documents (including the draft contract) for the competitive procedure with negotiation may take several months before the OJEU contract notice can be published.

The competitive procedure with negotiations can only be used in specific circumstances prescribed by the Financial Regulation, as follows:[88]

- Where only irregular or unacceptable tenders have been submitted in response to an open or restricted procedure.
- The needs of the contracting authority cannot be met without adaptation of a readily available solution.
- The purchase includes design or innovative solutions.
- The contract cannot be awarded without prior negotiations because of specific circumstances related to the nature, the complexity or the legal and financial make-up or the risks attached to the subject matter of the contract.
- The technical specifications cannot be established with sufficient precision by the contracting authority with reference to a standard.
- For concession contracts.
- For service contracts referred to in Annex XIV to Directive 2014/24/EU.
- For certain types of research and development services provided that the results do not belong exclusively to the contracting authority for use in its own affairs or where the service provided is not remunerated in full by the contracting authority.
- For service contracts concerning the purchase, development, production, or co-production of programmes for audiovisual media services as defined in Directive 2010/13/EU or radio media services or contracts for broadcasting time or programme provision.

[88] The same grounds apply for the use of the Competitive dialogue procedure, see further below.

The competitive procedure with negotiation allows for a more flexible approach by creating the possibility for an interactive negotiation with bidders to discuss their proposed approach before the contracting authority selects the tender offering the best price-quality ratio.

D. the ‘competitive dialogue’ procedure

The competitive dialogue procedure is also a ‘two-step’ process, similar to the ‘restricted’ procedure and to the ‘competitive procedure with negotiation’, whereby the contracting authority publishes an OJEU contract notice, and procurement documents are made accessible electronically. Similar to the restricted procedure, bidders submit requests to participate and must provide information in response to the selection document, which is used by the contracting authority to establish whether bidders are qualified to perform the contract and to select those that will be invited to participate in the ‘dialogue’ stage. The contracting authority is permitted to limit the number of bidders that it invites, which means that not all of the bidders that qualify must necessarily be invited to participate in the competitive dialogue phase. At least three (3) shortlisted bidders must be invited to the dialogue, unless there are not enough bidders meeting the exclusion and selection criteria.

The contracting authority may engage in a dialogue with bidders to identify the proper specifications and the best bidder. During this phase, all aspects of the project can be discussed with bidders and the number of solutions can be reduced as part of the dialogue process. Once the contracting authority is satisfied that it will receive proposals that will meet its requirements, it may declare the competitive dialogue phase closed and invite final bids from the remaining bidders.

The minimum time limit for receipt of requests to participate is 32 days from the publication of the OJEU contract notice. There is no minimum time-limit for receipt of final tenders (after the dialogue phase). Similar to the competitive procedure with negotiation, the duration of the entire process, from preparation phase to contract award, is likely to take at least 12 months and likely more than this for more complex procurements such as the one for the hyperconnectivity (depending on EuroHPC level of experience in conducting such a procurement procedure). In particular, the preparation of the tender process and the relevant documents (including the draft contract) for the competitive dialogue procedure may take several months before the OJEU contract notice can be published.

The competitive dialogue is designed for complex contracts, such as infrastructure construction, where the contracting authority may not initially know in sufficient level of detail what it wants to acquire. It can only be used under specific circumstances, which are identical to those used to justify recourse to the competitive procedure with negotiation.

In practical terms, the competitive dialogue requires significant resources to be managed, and therefore it is used only for exceptionally complex procurements where the contracting authority cannot objectively define the technical solutions that meet the needs or objectives or specify the project’s legal and financial framework. In such circumstances, the competitive dialogue provides flexibility through a dialogue stage aimed at exploring potential solutions, where all aspects of the project can be discussed and refined.

The choice between using competitive dialogue and the competitive procedure with negotiation is not always straight forward, as the grounds for use of both procedures are identical. It may depend on the contracting authority's objectives and its ability to draft detailed technical specifications from the start. If the contracting authority has a clear understanding of what it wants to purchase and can develop more precise specifications, but still require scope of negotiations, then the competitive procedure with negotiation would seem more appropriate. Conversely, if the uncertainty regarding the exact requirements is higher, competitive dialogue may offer better flexibility to refine and define the procurement needs during the dialogue process. On balance, it can be said that competitive dialogue should only be used where the purchase is really complex and if the contracting authority is not objectively able to define the technical means capable of satisfying the needs or objectives or not able to specify the legal or financial make-up of the project. In both cases, it is extremely important to have an experienced team of procurement officers that can manage these complex and highly regulated procedures.

E. The ‘negotiated procedure without prior publication of a contract notice'

This procedure can only be used in exceptional circumstances such as when only one company can complete the work (‘sole source’), or where publishing an OJEU contract notice is not possible due to specific reasons or in other specific circumstances, such as those specified in Article 11.1 of ANNEX I of the Financial Regulation[89].

One of these grounds is set out in Section (l) of Article 11.1 and allows recourse to this non-competitive procedure “for the purchase of public communication networks and electronic communications services within the meaning of Directive 2002/21/EC of the European Parliament and of the Council”.

Under the ‘negotiated procedure without prior publication of a contract notice', the contracting authority may conduct direct negotiations with one or more suppliers and enter into a contract without the need to publish the contract opportunity in advance.

7.3.2 Relevant considerations related to the service to be procured

A. Characteristics of the service

A very important finding of the requirement and gap analysis is that the services under discussion are <u>innovative advanced non-commodity services</u> in terms of:

- **network traffic patterns**:
    - As already presented, HPC research IP traffic patterns are much different than regular commercial IP traffic patterns. The latter commercial internet traffic consists of millions of small to medium (volumetric) flows (i.e. “mice flows”[90]), while in research HPC traffic patterns the number of users is far smaller, but each one of them may put excessive stress on the network infrastructure by requesting either long-lasting huge volumes of data

[89] Regulation - 2018/1046 - EN - EUR-Lex (europa.eu)
[90] https://www.sciencedirect.com/science/article/pii/S1389128605001118

(hundreds of TBs or even multiple PBs) to be transferred to a remote location (i.e. elephant flows[91]), or even by deploying low latency demanding applications. In essence, the HPC applications require non-overbooked – non-congested network, which is not the case with commercial IP networks, which are usually overbooked and may result in congestion when the demand for network resources exceeds the available capacity, causing high delays, high jitter, and even packet loss. Thus, research networks supporting "elephant flows" are rather innovative non-commodity services.

- Given the requirement of moving big data files in HPC applications, the end-to-end support of high Maximum Transmission Unit (MTU) size plays a significant role in efficient and fast data transfer. Research networks support end-to-end high MTU values (9000 bytes, also known as "jumbo" frames), while commercial IP providers usually use the standard 1500 bytes MTU, mainly for interoperability reasons. When comparing network overhead between different MTU sizes, such as 1500 bytes and 9000 bytes, the overhead (packet headers) is approximately 10%, while with an MTU size of 9000 bytes, it is only around 1%. Thus, 9000 MTU size networks are rather innovative non-commodity services. The nature of the data transfers between the HPC Ecosystem mandates the end-to-end use of Jumbo Frames (MTU size of 9000 bytes).

- **network speeds and capacities:**
  - Under the hyperconnectivity tender multiple end points require services at terabit levels, thus, a large number of high-capacity end sites (HPC or data centres) are planned to receive services, exceeding 400Gbps, e.g. 800Gbps or even 1200Gbps in 2025 and more in 2028. In terms of standardisation, the 800GbE and 1.6TbE standards are being developed by the IEEE P802.3df Task Force. In February 2024, the 800GbE was approved and the first installation of 800Gbps equipment was announced in January 2024 [92] [93] [94] [95] [96] [97] [98] [99] [100] [101] [102]. In practice, there is no end user service over 400Gbps widely available in the market as a commodity. In fact, for such capacities bespoke solutions are required. In essence, many

[91] Elephant flow - Wikipedia
[92] https://en.wikipedia.org/wiki/Terabit_Ethernet
[93] http://www.ieee802.org/3/bs/NGOATH_3bs_Objectives_16_0122.pdf
[94] https://www.researchgate.net/profile/Mehrdad-Nourani/publication/3937931_System_Requirements_for_Super_Terabit_Routing/links/547c9fff0cf2cfe203c1f8ec/System-Requirements-for-Super-Terabit-Routing.pdf
[95] https://telecomtalk.info/lumen-launches-400g-wavelength-services-across-europe/686746/
[96] https://www.opensignal.com/2023/10/12/european-markets-show-significant-disparity-in-broadband-user-experience
[97] https://www.eenewseurope.com/en/toshiba-europe-shows-400gbit-s-qkd-quantum-network/
[98] https://neosnetworks.com/resources/press-releases/neos_networks_launches_national_400gbps_services/
[99] https://telecomtalk.info/zayo-group-upgrades-european-network-to-400g/879219/
[100] https://wholesale.orange.com/international/en/knowledge-hub/insights/is-400g-network-capacity-right-for-you.html
[101] https://www.lightwaveonline.com/data-center/article/14303686/nl-ix-deploys-a-range-of-nokia-service-routers
[102] https://telxius.com/en/telxius-delivers-400gb-s-ethernet-service-powered-by-ciena-and-infinera/

of the required capacities are rather innovative, non-commodity services.

- For the hyperconnectivity network to be deployed significant upgrades may need to be implemented by the end sites. New equipment may need to be acquired and installed. Such equipment, capable to provide and consume services over 400Gbps is currently not widely available in the market and significant delays can be expected. This situation effects the delivery of the service and the time plan introducing risks for the potential bidders. These risks need to be addressed and negotiated through the procurement procedure in order to find ways to mitigate them.

- **architectures and service provisioning:**
  - Connectivity services in the hyperconnectivity network need to support a set of innovative security and performance monitoring services allowing:
    - special security architectures to be implemented (e.g. Science DMZ[103]), to be able to control the hyperconnectivity traffic without needing a firewall for the whole traffic (as firewalls for traffic above 100Gbps are very expensive).
    - performance monitoring infrastructure to be installed throughout the network. This also involves the installation of specific hardware and software within the providers core and edge networks. Such an innovative requirement needs to meet specifications of the network providers, and may need to be negotiated during the procurement process.

B. Ability to draft detailed technical specifications upfront

Based on the Needs Assessment and Gap analysis performed, it appears that the hyperconnectivity network comes with a non-definite set of requirements, especially for the set of end points. For example:

- some of the important end-sites were still implementing or planning their infrastructure and/or site and were not able to provide concrete input,
- others did not have the resources to commit during the study and provided non- or limited input,
- in many cases upgrade plans of the end-users were affected by unforeseen circumstances that result into non-concrete input.

Some of the open issues identified include:

- ***The requirements from the AI factories are not fully known****;* although it can be assumed in most of the cases that the AI factories will be collocated with the EuroHPC hosting sites, the specific network needs cannot be easily predicted. Given also the fact that the procurement

[103] https://fasterdata.es.net/science-dmz/science-dmz-architecture/

of the AI factories will take some time to be concluded, it is foreseen that no major impact is expected in the network until year 2025 or even 2026. On the other hand, after that and towards 2028, the AI factories will be in full production and a change in the needs and traffic flows may be observed, depending also on the number of applications and usage of the systems.

- ***The requirements collected are not in all cases definite****;* For capturing requirements and assessment of the relevant input, a powerful set of instruments were utilised, which resulted in a long list of end points (~330) to be hyper-connected. Despite the multi-step validation process, including in some cases the end sites themselves (e.g. EuroHPC hosting sites), and the validation by the upstream providers of the end sites (NRENs), the concrete requirements in terms of dataset sizes that need to be moved for a very diverse set of scientific and industrial applications cannot be easily certified. The collected *requirements can be thus over- or under-estimated.* Furthermore, although a diverse set of channels was used for the dissemination of the study and its requirements collection tools, it is clear that *not all requirements have been collected. Requirements may be missing, or change or new requirements may appear over time. In particular, given the peer-review based award process of time in the most HPC systems, the actual applications that will be approved cannot be foreseen beforehand,* In addition, the revolution of the generative AI models in such a short period of time, clearly shows that new requirements can emerge within 1-2 years which cannot be foreseen and may require additional input through discussions with the potential bidder or bidders during the procurement process.

- ***The number of end-sites (HPC/Data centres) is not guaranteed;*** hyperconnectivity services involves a big number of end sites (HPC/data centres) which may require the installation of equipment and provision of services that could affect the operations of the internal infrastructures and their end-users. Furthermore, the end-sites are independent institutions with their own administrations and technical management, and thus, with independent and non-uniform priorities, policies, and plans. Agreement of each one of these entities is essential before providing hyperconnectivity network services to them. Such an agreement has not yet been achieved and constitutes a significant risk as it can affect the scale of the project and the relevant budget. This risk needs to be addressed and negotiated through the procurement procedure in order to find ways to address it. At the end, some end-sites may choose not to be connected; as a follow-up from the previous long list of end points (~330), and although in the first 60M tender, it is expected to be able to connect or upgrade around 170 end points (i.e. around half of the points which are mandatory), it is clear that some of them may not be connected or upgraded for different reasons (such as inability to persuade the legal signatory of the end site, considerable investment required from the end-site to upgrade its internal infrastructure in order to benefit from the hyperconnectivity or just bureaucratic reasons, e.g. delays in the decision-making process).

- ***The requirements from the EuroHPC Federation Project (EuroFP) are not yet known;*** the transition from a set of EuroHPC machines to a federated (and hyper-connected) EuroHPC ecosystem may have an impact on the network. As outlined in the Multi-Annual Strategic Programme 2021-2027 (version 2023)104, EuroHPC will be offering a "service that enables fast and reliable transfer of data between supercomputing resources and data sources/lakes, including data staging, transfer protocols, and data synchronisation". This will increase the interaction among the EuroHPC hosting sites. This is also described in the "Annex to the Commission Implementing Decision amending Implementing Decision C(2021) 9463 on the financing of the Connecting Europe Facility – Digital sector and the adoption of the multiannual work programme for 2021-2025"[105], where it is stated that "novel and future applications that depend on human intervention will require real-time and/or interactive computing, in which high-speed connectivity is essential for a fast access to the scattered big data and HPC resources and applications, in which time is critical to ultimately save lives or reduce material impacts". In the future, some paradigms of HPC-over-the-network may also come to play, or at least shifting or moving datasets on the fly across sites as this was stated in one of the interviews with Destination Earth, running critical applications and identifying the most suitable HPC site(s) with available resources to run some of its applications, such as the Extremes Digital Twin.

C. Time constraints

A significant constraint for the hyperconnectivity network is the timeframe imposed by the EuroHPC JU Work Programme and funding instruments (CEF). There are two (2) major constraints to that extent:

a. the process has to be initiated by the beginning of December 2024, and

b. the services must be implemented in a timely manner and be operational for a significant duration until 2029.

As already discussed, implementation of the services consists of upgrades to the providers network infrastructure that includes state-of-the-art equipment which in some cases involves lengthy times for delivery and commissioning by equipment providers. The same also applies for equipment upgrades needed in the end-sites' infrastructure. For example, 800Gbps capable routers and line cards can have more than 6 months for delivery, a massive order for such equipment could increase this timeframe even longer. Such risks could be mitigated if the existing plans of the possible providers and end-sites could be taken into account allowing the adaptation of the specifications accordingly. Such an approach could reduce time needed to implement the network giving the opportunity and the tools for a meaningful implementation in stages, significantly reducing risks for the timeplan.

Implementation risks and delivery times are also an issue that affects infrastructures within end-sites' premises. Consuming hyperconnectivity services may require upgrades from the side of end-sites too. Such upgrades are expected to consume significant resources and, in some cases, to even require re-design of parts of end-sites internal networking infrastructure. Such resources might not be available

[104] https://eurohpc-ju.europa.eu/document/download/ea5b58d8-0270-4247-ba5e-738173af4f1f_en?filename=Decision%2008_2023_%20Amendment%20MASP%202021-2027_0.pdf
[105] https://ec.europa.eu/newsroom/dae/redirection/document/95109

in time by end-sites affecting service delivery, project's time-plan, and resource engagement during implementation.

7.3.3 Evaluating Alternative Procedures for the Hyperconnectivity Procurement

A. The 'open' procedure

In light of the above analysis, the open procedure presents the following shortcomings if used for the procurement of hyperconnectivity:

- It requires detailed and complete technical specifications upfront (to be defined in the tender documents).
- It offers no possibility for negotiations with the selected bidder during the procurement process or before contract signature to finetune the requirements or to discuss the proposed solution and strategy for implementation of the contract.
- Due to their internal policy and constraints, GÉANT & NREN networks may not be willing to take part in a competitive tender procedure, which will place them in a position where they compete against their existing suppliers, and risk jeopardising long-standing collaborations.
- The estimated timeframe for the process may last 12 months or more (from procurement preparation to contract award) and will require a competent procurement function to ensure compliance with the applicable procurement rules in carrying out the procedure, or else risk a legal challenge.

B. The 'restricted' procedure

The restricted procedure suffers from the same disadvantages as the open procedure and may be slightly lengthier due to its two-step structure.

C. The 'competitive procedure with negotiation' procedure

The competitive procedure with negotiation could potentially be suitable for the Hyperconnectivity procurement, as it designed for complex contracts, and allows for flexibility to discuss and fine tune all aspects of the proposed approach offered by bidders thereby facilitating the adoption of creative solutions.

However, the main shortcomings of using the competitive procedure with negotiation are:

- It will need to be demonstrated by EuroHPC that one (or more) of the conditions for use of the competitive procedure with negotiations are met in the circumstances.

- It will require significant resources to plan and execute this procedure, including highly qualified procurement professionals in order to ensure that it yields the expected result and in order to avoid a potential legal challenge during the negotiation phase or after.

- It will require significant preparation time to develop the tender documents (including a draft contract) and ensure the design of the process to accommodate negotiations, while ensuring legal compliance and suitability for purpose.
- Due to their internal policy and constraints, GÉANT & NREN networks may not be willing to take part in a competitive tender procedure, which will place them in a position where they compete against their existing suppliers, and risk jeopardising long-standing collaborations.
- Due to the negotiation phase inherent in this procedure, it is likely to exceed the project time constraints, with estimated time frame of at least 12 months, likely even longer.

D. The 'competitive dialogue' procedure

The competitive dialogue procedure could potentially be suitable for the hyperconnectivity procurement, provided it can be argued that the EuroHPC is not objectively able to define the technical means capable of satisfying the needs or objectives or not able to specify the legal or financial make-up of the project.

However, similar to the case of the competitive procedure with negotiation, the main shortcomings of using the competitive dialogue are:

- It will need to be demonstrated by EuroHPC that one (or more) of the conditions for use of the competitive procedure with negotiations are met in the circumstances.
- It will require significant resources to plan and execute this procedure, including highly qualified procurement professionals in order to ensure that it yields the expected result and in order to avoid a potential legal challenge.
- Due to their internal policy and constraints, GÉANT & NREN networks may not be willing to take part in a competitive tender procedure, which will place them in a position where they compete against their existing suppliers, and risk jeopardising long-standing collaborations.
- It will require significant preparation time to develop the tender documents and ensure the design of the process to accommodate negotiations, while ensuring legal compliance and suitability for purpose.
- Due to the dialogue phase inherent in this procedure, it is likely to exceed the project time constraints, with estimated time frame of at least 12 months, more likely longer.

E. The 'negotiated procedure without prior publication of a contract notice'

Article 11.1(l) of ANNEX I of the Financial Regulation (hereinafter "**Article 11.1(l)**") allows recourse to the negotiated procedure without prior publication of a contract notice "for the purchase of public communication networks and electronic communications services within the meaning of Directive 2002/21/EC of the European Parliament and of the Council". As will be further discussed below, this exemption could potentially apply in the circumstances to justify the use of this procedure for the hyperconnectivity procurement.[106]

The advantages of utilising this procedure for the hyperconnectivity procurement are evident. It provides the necessary flexibility to thoroughly discuss all aspects of the project with GEANT and NRENs (as per Scenario A – Expanding existing infrastructure by utilising GÉANT and NRENs networks) and to tailor the best solution collaboratively with GEANT. This level of flexibility exceeds that of the competitive procedure with negotiation or the competitive dialogue, as it is not subject to any legal constraints designed to ensure 'fair negotiations' among bidders participating in the competitive procedure with negotiation or the competitive dialogue. Additionally, this method does not impose minimum timeframes, and does not require an extremely lengthy preparation time, thus allowing for a relatively rapid execution to meet the previously mentioned time constraints.

7.3.4 Arguments in support of utilising Article 11.1(l) for the Negotiated Procedure Without Prior Publication of a Contract Notice[107]

A. The scope of Article 11.1(l)

Article 11.1(l) allows recourse to the negotiated procedure without prior publication of a contract notice "for the purchase of **public communication networks** and **electronic communications services** within the meaning of Directive 2002/21/EC of the European Parliament and of the Council".

Directive 2002/21/EC, no longer in force, has been replaced by directive 2018/1972/EC. Preamble (15) of directive 2018/1972/EC provides that "the definition of electronic communications services should therefore contain three (3) types of services which may partly overlap, that is to say <u>internet access services</u> as defined in point (2) of Article 2 of Regulation (EU) 2015/2120 of the European Parliament and of the Council, interpersonal communications services as defined in this Directive, and <u>services consisting wholly or mainly in the conveyance of signals</u>. The definition of electronic communications service should eliminate ambiguities observed in the implementation of the definition as it existed prior to the adoption of this Directive and allow a calibrated provision-by-provision application of the specific rights and obligations contained in the framework to the different types of services."

Article 2 of directive 2018/1972/EC also includes the following definitions which are relevant to determine the applicability of Article 11.1(l):

- '**electronic communications service**' [Article 2(4)] means "a service normally provided for remuneration via electronic communications networks, which encompasses, with the exception

[106] Supporting arguments for use of this procedure are presented below. However, EuroHPC should conduct its own independent legal review and analysis to ensure compliance with the Financial Regulation or any other applicable regulations.

[107] This Section aims to offer potential legal justification for using the negotiated procedure without prior publication of a contract notice, based on Article 11.1(l). However, it is not intended to substitute for independent legal advice. EuroHPC should conduct its own thorough legal analysis to ensure compliance with the Financial Regulation and any other applicable laws and regulations.

of services providing, or exercising editorial control over, content transmitted using electronic communications networks and services, the following types of services:

(a) ‘internet access service’ as defined in point (2) of the second paragraph of Article 2 of Regulation (EU) 2015/2120;[108]

(b) interpersonal communications service; and

(c) services consisting wholly or mainly in the conveyance of signals such as transmission services used for the provision of machine-to-machine services and for broadcasting”;

- **‘electronic communications network’** [Article 2(1)] means “transmission systems, whether or not based on a permanent infrastructure or centralised administration capacity, and, where applicable, switching or routing equipment and other resources, including network elements which are not active, which permit the conveyance of signals by wire, radio, optical or other electromagnetic means, including satellite networks, fixed (circuit- and packet-switched, including internet) and mobile networks, electricity cable systems, to the extent that they are used for the purpose of transmitting signals, networks used for radio and television broadcasting, and cable television networks, irrespective of the type of information conveyed”.

- ‘**provision of an electronic communication network**’ [Article 2(16)] means “the establishment, operation, control or making available of such a network”.

- **‘public electronic communications network’** [(Article 2(8)] means “an electronic communications network used wholly or mainly for the provision of publicly available electronic communications services which support the transfer of information between network termination points”.

B. The purpose of the EuroHPC hyperconnectivity procurement

To legally rely on the exemption set out in Article 11.1(l), EuroHPC must be able to demonstrate that the principal purpose of the procurement is the purchase of **public communication networks** or **electronic communications services,** as per the definitions explained above. It is acceptable to include ancillary services (such as consulting or professional services) if they support the principal purpose of providing public communication networks or electronic communications services.

The EuroHPC hyperconnectivity service procurement aims at laying out and achieving a secure, and hyper-connected European HPC and data infrastructure. After analysing the relevant options based on the discussions with EuroHPC JU and the end-users, as well through market analysis of the relevant available services in Europe, the solution was designed as a service-based network providing to end-users advanced internet access and connectivity. The services, include:

[108] Article 2(2) of Regulation (EU) 2015/2120 provides the following: ‘internet access service’ means a publicly available electronic communications service that provides access to the internet, and thereby connectivity to virtually all end points of the internet, irrespective of the network technology and terminal equipment used”.

- Advanced “Internet access service”, in-line with Article 2 of directive 2018/1972/EC, which includes:
    - IP Connectivity at high-bandwidth (up to more than 1,2Tbps), low latency, and high reliability to meet HPC requirements;
    - Support of IP Jumbo Frames: packets larger than the standard 1500 bytes (as in commercial networks), typically around 9000 bytes;
    - physical access high availability: highly available physical access to support maintenance, troubleshooting, and upgrades
    - connectivity with Public Cloud Providers (e.g., AWS, Azure, GCP, OVH)
    - rapid provision of direct logical circuits to Public Cloud Providers
    - high bandwidth peering in Major Internet Exchange Points (IXPs)
    - end-User Support/Helpdesk (Level 1, 2, 3)
    - network-level security: DDoS Protection Service
    - specific automated Network Upgrade Process (i.e. when traffic consistently exceeds 60% of a link capacity)
    - direct high bandwidth peerings with NREN Networks (if this is not the case)
- Performance measurement infrastructure and services: allowing end-user to actively monitor performance, availability, and reliability of network services (e.g. uptime, latency, packet loss, jitter);
- Bespoke implementations, e.g. virtual Ethernet Private Lines (EPLs), virtual Private LANs reserved or guaranteed bandwidth, between end-users of the network;
- Professional Services: consulting, project management & technical expertise to ensure the successful consumption of internet access services (e.g. upgrading end-user’s internal infrastructure) and implementation of additional service (i.e. reserved or guaranteed bandwidth services).

C. The services that are provided by GÉANT

In support of the explanation provided above, which argues that the principal purpose of the procurement is the purchase of electronic communications services,[109] it is also useful to consider the type of services provided by GÉANT. As the potential service provider, GÉANT and NRENs networks have been identified as the best fit scenario (Scenario A – Expanding existing infrastructure utilising GÉANT and NRENs networks).

The GÉANT Association is a collaboration of European National Research and Education Networks (NRENs). Together (i.e. GÉANT and NRENs), they deliver an information ecosystem of infrastructures

[109] The analysis presented in this document is not intended to replace independent legal advice. EuroHPC should conduct its own thorough legal analysis to ensure compliance with all applicable laws and regulations and to satisfy itself that the circumstances justifying the use of Article 11.1(l) are indeed met.

and electronic communication services to advance research and educational institutes and innovation on a global scale. They are interconnecting and thus providing internet access service to around 10.000 institutions. The GÉANT Association has national members, representative members (NORDUnet) and associates (such as CERN and ESA).

GÉANT is providing electronic communication services and public communication network, which is divided into three (3) parts:

A. the Infinera DWDM open line system;

B. the Infinera DWDM transponders layer and Juniper-based Internet protocol / multiprotocol label switching;

C. (IP/MPLS) network [Infinera; Juniper].

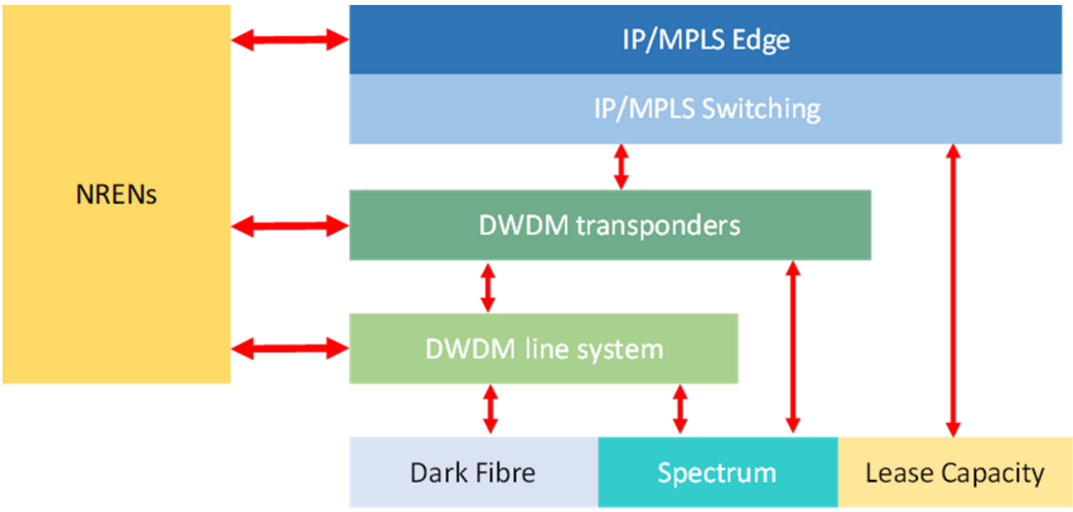

Figure 16: GEANTs layered network implementation

The Infinera DWD open line system runs on top of dark fibre and spectrum links and provides management of the optical, DWDM signals generated by the transponders layer, that is, amplification, multiplexing and power balancing. The "open" nature of the Infinera DWD open line system means that this layer can accept signals generated by different sources without restriction to a specific vendor or a specific technology. Spectrum services for NRENs are delivered directly by this layer.

The next layer up the stack is the Infinera DWDM transponders layer, constructed utilising Infinera Data Centre Interconnect (DCI) transponders. This layer is responsible for activating point-to-point capacity over the line system routes. Transponders terminate short-reach Ethernet signals of 10, 100 and 400 Gbps capacity, and generate high- capacity DWDM signals for transmission over the DWDM line system. These high-capacity point-to-point connections are used for carrying the trunks between the IP/MPLS routers of GÉANT. In some cases, transponders are directly connected to spectrum supplied by a provider; this is the case when either the amount of spectrum is very limited (therefore not justifying the investment in the line system deployment), or there are limitations preventing the deployment of the full line system. “Managed Wavelength” services, or point-to-point high-capacity guaranteed Ethernet services to NRENs are also delivered directly by this layer.

The upper layer of the stack is the IP/MPLS layer, which is today managed by Juniper MX series routers; this layer is undergoing change as part of the IP/MPLS refresh where Juniper MXs will be replaced by Nokia 7750 SR-7s and SR-2se devices. This layer is responsible for transmission of IP packets and Ethernet frames between connectors across the GÉANT network. All connectivity services not delivered by the other layers are delivered by the IP/MPLS layer.

This network is operated, maintained and available to the research institutes and educational institutions throughout Europe. The ‘public’ nature of the communication network provided by GEANT is therefore evident by the fact that members and associates of GEANT use the electronic communications network of the GÉANT Association wholly or mainly for the provision of publicly available electronic communications services which support the transfer of information between network termination points. See in this regard the definition of public electronic communications network’ referred to above.

Based on the analysis presented above, it can be concluded[110] that:

- The best fit scenario is Scenario A – Expanding existing infrastructure (utilising GÉANT and NRENs networks)
- The most advantageous procurement alternative for EuroHPC (both in terms of flexibility and in terms of timeframe) would be to utilise the negotiated procedure without prior publication of a contract notice and engage in direct negotiations with GÉANT.
- Subject to further independent legal analysis to be conducted by EuroHPC, it can be argued that the principal purpose of the hyperconnectivity procurement is the purchase of electronic communications services within the meaning of Directive 2018/1972/EC. Therefore, it can be argued that the circumstances outlined in Article 11.1(l) of ANNEX I of the Financial Regulation are met, allowing EuroHPC to utilise “the negotiated procedure without prior publication of a contract notice” for direct negotiations and contract award to GÉANT and NRENs networks.

8. Implementation Roadmap

8.1. Roadmap & Evolution Plan

The EuroHPC hyperconnectivity network, will be supported by two main tenders. The first hyperconnectivity tender, with an indicative budget of 60M EUR, is expected to be launched by December 2024 at the latest. In fact, the tender must be published before the 8th of December to secure the allocated funding. The relevant requirements and technical specifications have been prepared for the first tender. The indicative timeline for the tender service spans 2025 to 2029, depending on the specific tender process determined by the EuroHPC Governing Board (GB). This

[110] Disclaimer: the information and guidance provided in this document should not be considered legal advice. This document is not intended to replace professional legal advice.

means that the service duration of the first tender will be approximately four (4) years. Assuming that the awarded contractor requires a lead time of around nine (9) to twelve (12) months to initiate the hyperconnectivity service, the likely timeline would be from the second half of 2025 to the second half of 2029.

The second hyperconnectivity tender, with a budget of 100M EUR, is expected to be launched towards the end of 2027 and must be contracted by 2028. According to the EuroHPC JU Work Programmes and related documents, funding for the second tender will be released gradually, starting with 10M EUR in 2025, followed by 30M EUR in 2026, and finally 60M EUR in 2027.

Two (2) main scenarios are being outlined for the transition between the two tenders. Scenario 1, where the second tender term begins after the first ends, and Scenario 2, where there is an overlap in service offerings between the two tenders. The decision between these scenarios will depend on the legal and financial frameworks associated with CEF funding, the specific tender call(s), and the preferences of the EuroHPC JU and its Governance Board. Another important factor is whether the second tender will be awarded to the same service provider.

- The first scenario is straightforward. The handover, either from the first to the second contract with the same service provider (Scenario 1a) or from the first to a new service provider (Scenario 1b), is expected to occur in the last months of the first service term, likely in 2029. If the same provider is retained, the handover will be simpler and may take only a few months. However, if a different service provider is selected, the process will require at least 5-6 months.
- The second scenario is more complex. The first consideration is the same as in the first scenario: whether there will be a change of provider. If there is no change, it is referred to as Scenario 2a, while a change of provider is referred to as Scenario 2b. The second consideration involves whether the second tender will only extend the service period (Scenario 2a1 and 2b1) or extend both the service period and expand coverage to more end sites (Scenario 2a2 and 2b2). In the latter case - extension and expansion of service, which is the most likely outcome - the expansion to new sites can start earlier (in 2028), while the handover to the second tender contract/provider will take place again in the last months of the first contract service term (in 2029). In the less likely case (extension of service without expanding to new end sites), and if a new provider is selected, the new provider would need to connect the end site with a second connection (dual-homed mode), and both connections would need to operate in parallel for a short period. This may lead to technical issues, such as the capability of end sites' border routers to handle Full Internet Routing Tables (FIRT), which is critical for a smooth transition. In such cases, the handover may require a longer period and more technical effort, potentially extending over a year (from the second half of 2028 to the first half of 2029). If the same service provider is retained, and given the larger budget for the second tender, it would be impractical to simply extend the service period without also expanding coverage. As a result, Scenario 2a1 may not be sufficient to absorb the budget, making Scenario 2a2 (expansion to more sites) necessary. In summary, the feasible scenarios are as follows:

    - Scenario 1a: Tender 2 starts after the end of Tender 1 with the same provider.

- Scenario 2a1: Tender 2 partially overlaps with Tender 1, with the same provider and the same sites.
- Scenario 2a2: Tender 2 partially overlaps with Tender 1 with the same provider and expansion to new sites
- Scenario 1b: Tender 2 starts after the end of Tender 1 with a new provider.
- Scenario 2b1: Tender 2 partially overlaps with Tender 1, with a new provider and the same sites.
- Scenario 2b2: Tender 2 partially overlaps with Tender 1, with a new provider and expansion to new sites.

A series of technical and non-technical aspects, related to the implementation roadmap and evolution scenarios, have been analysed, including the strategy for upgrades (i.e. backbone-first, edge-first or mixed), cost estimation for the entire period (until 2032), new technologies that may influence the network evolution, and administrative considerations, such as the requirement that each end site agrees to hyperconnectivity in order for the service to terminate within their premises.

For the upgrade strategy, a mixed strategy is preferred, with particular emphasis on the backbone to ensure readiness for serving and monitoring edge sites, including the Performance Monitoring Points (PMPs).

Regarding the cost/pricing aspects a sensitivity analysis has been conducted, identifying three (3) main scenarios: the “base” scenario, the “accelerated” scenario (i.e., upgrades occurring sooner than expected) and the “delayed” scenario (upgrades occurring later than expected). As anticipated, the delayed scenario is the most cost-efficient, as telecom circuit/equipment prices tend to decrease over time[111]. The accelerated scenario is the most expensive, while the base scenario falls in the middle. It is expected that the upgrade of end sites will be a combination of these three (3) scenarios. However, it is difficult to predict which will dominate, as several variables remain uncertain, such as the implementation and demand for AI factories. Only the accelerated scenario is projected to slightly exceed the indicative 100M EUR budget for the second tender. For the base and delayed scenarios, a surplus of approximately 10M and 15M EUR, respectively, is estimated. This way it is expected that a significant amount of funds will be available for other purposes. These may include the purchase or upgrade of end sites’ border routers/switches[112] and related components (e.g. cards) for end sites (HPC/Data centres) that were unable to connect during the first tender due to challenges (e.g., border routers unable to support hyperconnectivity speeds). Other potential uses, depending on CEF and respective call rules, could support the implementation of Science DMZs[113], which can be combined with firewalls.

[111] State-of-the-network-2023.pdf – Pricing, page 5 (telegeography.com) - TeleGeography reports 100 Gbps wavelength prices decreasing at an average of 12% annually

[112] Border routers/switches of HPC/Data sites are the network elements that interconnect the sites with the broader internet. The hyperconnectivity service will terminate at these border elements. However, upgrading these elements is not within the scope of the tender. This means that if the sites are unable to afford upgrades to their border elements, they may not be able to connect to the hyperconnectivity service.

[113] The Science DMZ Architecture and Security

In the event that significant new sites emerge and are not included in the list of Points of Interest (PoIs), they should also be incorporated and covered by the remaining budget. Such end sites may include AI factories, although it is expected that most of these will be collocated with existing EuroHPC hosting sites, which will already be hyperconnected.

New technologies that could impact the hyperconnectivity network are analysed, primarily focusing on the new 800GbE standard and the forthcoming 1600GbE standard, expected around 2026. The 800GbE technology is anticipated to be integrated into the hyperconnectivity network in the initial years of service (2025-2026), while the 1600GbE standard may become more relevant towards the end of the first tender and during the second tender.

8.2. Integration of new sites

The integration of new sites into the hyperconnectivity network is essential due to the dynamic nature of the network, evolving socio-economic and technical factors, and the development of new EU level projects and initiatives. Relevant examples include "AI factories[114]" (if their hosting sites are not already connected to the hyperconnectivity network), new Common EU Data Spaces[115] or new Research Infrastructures that function as data providers with HPC needs, which go live and require hyperconnectivity. The following paragraphs provide guidelines for managing network changes, especially when new sites need to be connected to the hyperconnectivity network. The network is a public resource, specifically designed to support Europe's High-Performance Computing (HPC) ecosystem. As such, only sites directly related to this ecosystem are eligible for connection, while users external to this ecosystem can access HPC resources through commercial networks and established peerings/Internet Exchange Points (IXPs) with the hyperconnectivity network.

Prospective sites wishing to connect to the hyperconnectivity network must submit a formal ***connection request***, documenting their relationship with the European HPC ecosystem and detailing their specific relation. The application must also describe any existing network flows with already connected HPC sites. The hyperconnectivity network offers a streamlined application process through a web-based form, and requests must be submitted by an authorised administrative representative. Responses to connection requests are guaranteed within ten (10) working days[116], with the possibility of the request being forwarded to a local partner for local connection management.

The technical ***assessment of connection requests*** is vital to ensure that the proposed site's networking needs are fully understood. This assessment includes evaluating the required capacity, future traffic projections, and any additional requirements, such as low latency or minimal jitter. Since the hyperconnectivity network is designed for large-scale "elephant-flow[117]" traffic, it is essential to estimate capacity needs in advance. If an incoming connection risks overloading the access network, upgrades to the backbone or secondary links may be necessary to prevent congestion. Additionally, reliable connection setups are supported through redundant links or nodes to ensure network

[114] AI Factories | Shaping Europe's digital future
[115] Common European Data Spaces | Shaping Europe's digital future (europa.eu)
[116] This period is based on common practice and can be adjusted during the related negotiated procedure
[117] Elephant flow

resilience in case of failure. Security is also of paramount importance, with new sites required to comply with their providers' Acceptable Use Policies (AUP)[118] and implement robust security measures to prevent malicious activity.

The process of integrating a new site must follow a structured **timeline**, beginning with a preparation phase in which the network or local partner conducts a comprehensive needs assessment to identify the site's specific requirements. Following this, a detailed project plan is developed, outlining objectives, milestones, and necessary resources. This phase typically takes about one month.

Next the design phase, during which the network architecture is carefully crafted to ensure scalability, redundancy, and security. The appropriate hardware is selected to ensure seamless integration with the existing infrastructure, and data migration plans are prepared. This phase is expected to last an additional month.

The implementation phase involves procuring and installing the necessary hardware and software, configuring network settings, such as IP addresses and VLANs, and performing initial testing to verify functionality. This phase takes approximately one to two months.

The testing and optimisation phase includes rigorous stress tests and security assessments to ensure the network's resilience and performance under heavy loads. Network configurations are fine-tuned based on test results to optimise performance, and optional staff training is offered to ensure smooth network management. This phase is expected to be completed within five months from the initial starting date.

It is important to note that the available budget is a key factor in the integration of new sites. As presented in the implementation plan, there are two related hyperconnectivity tenders. The service period for first one is expected to run from 2025-2029, and the second one from either 2028 or 2029 to 2032 (or possibly beyond).

For new sites interested in applying for hyperconnected in the future, an associated budget must be available to cover interconnection costs (for the end site) and, if necessary, the access and backbone networks. The first hyperconnectivity tender includes a list of end points. If not all of these sites are hyperconnected, there may be a budget surplus that could allow for the connection of new sites. Otherwise, it will be up to the contractor(s) to determine whether they can connect the new site.

Most identified scenarios suggest the second tender will have a budget surplus, making it easier to accommodate the interconnection of new sites (provided that the projections of the implementation roadmap prove accurate).

Finally, in order to integrate new sites during the course of the two (2) hyperconnectivity tenders, flexibility in the tendering process will be required. An open tender cannot accommodate the addition of new sites, as the list of sites must remain fixed after the tender is published. Therefore, a negotiated procedure is necessary.

118 National Research and Education Network policies – GEANT Compendium online version - compendium.geant.org/policy

8.3. Service Handover

The handover procedure, in case there is such a necessity, should be designed and implemented based on a specific framework, fit for the hyperconnectivity network services. This handover includes a plan with instructions and associated artifacts to facilitate the transfer of the hyperconnectivity service. It covers technical and service details, identifies the parties involved, and outlines the capacities of experts who can advise the EuroHPC JU in the event of a required transition. The handover may be necessary due to a change in the electronics/network service provider and/or a change in the technology provider/supplier (e.g. of IP or optical equipment). One conceivable, though unlikely, scenario is the transition from the service provider selected for the first hyperconnectivity tender (60M EUR) to a new provider for the second hyperconnectivity tender (100M EUR).

As the hyperconnectivity network involves hundreds of end sites, referred to as Points of Interest or Points of Interconnection (PoIs), as well as numerous different administrative entities, the handover plan for such services is expected to be a complex and resource-intensive task. Although the specifications of the hyperconnectivity network have been documented, the exact implementation details are not yet fully known. Technical details are crucial for the development and implementation of a concrete handover plan. Nevertheless, based on the current specifications, a handover framework is presented.

A key objective of the handover plan is to ensure uninterrupted hyperconnectivity service during the transition for all sites. The handover plan aims to facilitate the comprehensive transfer of knowledge, documentation, and operational responsibilities from the outgoing provider to the incoming one. This includes providing an updated list of connected end sites, detailed network diagrams, configuration files, operational procedures, and troubleshooting guides.

The handover plan template is further detailed, highlighting the main elements that need to be identified and finalised at the time of the handover. These elements include the parties involved, timeline, communications, documentation, validation, risk management, handover execution and post-handover support.

- In line with the service specifications, the parties involved include the relevant EuroHPC JU team, the selected service provider consortium (outgoing) and the corresponding incoming service provider consortium, along with their respective teams and third-party vendors for both the optical and internet layers. The end points (i.e. end sites) and their teams, are also key stakeholders in the hyperconnectivity ecosystem. These include HPC providers, data providers, major cloud providers, Internet Exchange Points (IXPs) and NRENs, in cases where the selected consortium from the first tender does not include NRENs.
- The handover timeline depends on the overall hyperconnectivity implementation timeline, as outlined in the paragraphs above. Depending on the specific scenario implemented regarding the two hyperconnectivity tenders, it is expected that a period of at least 5-6 months will be needed for the actual transition.
- A robust communication plan is essential for the successful handover of any network or project. In the case of the hyperconnectivity network services, where a large number of stakeholders are involved, the communication plan becomes a core element for the

handover's success. It ensures that all parties have access to relevant information, deadlines, assignments, and progress updates, thereby enhancing productivity and keeping everyone informed.

- Technical documentation is a crucial factor for the successful handover of hyperconnectivity network services, as it provides a comprehensive reference for the incoming team. This documentation should cover several key areas to ensure continuity and effective management, including an updated list of hyperconnected end sites with their connection speeds, network diagrams, IP addressing schemes and VLAN configurations, routing tables and policies, firewall rules and policies, bandwidth allocations, performance monitoring setup and documentation, as well as bespoke implementation details.
- Testing and validation of the services from the outgoing to the incoming provider are key to a successful handover. Before the actual handover takes place, a series of tests must be performed, including end-to-end connectivity tests for all sites, performance baseline measurements (e.g., latency, packet loss, jitter), failover and redundancy tests, security assessments and vulnerability scans, application performance testing and capacity and scalability tests.
- The handover of hyperconnectivity network services, which involves complex and extensive infrastructures, presents several risks that must be meticulously managed to ensure a successful transition. These risks, if not properly addressed, can lead to significant operational challenges and disruptions.
    - One of the most critical risks during the handover is *network downtime*. Changes made during the transition can inadvertently lead to unplanned outages or performance issues, resulting in lost productivity and revenue. To mitigate this risk, it is essential to implement a phased transition approach, allowing for incremental changes that can be tested and validated before full deployment. This approach helps ensure that any issues are identified and resolved without impacting the entire network.
    - The transition period can introduce *security vulnerabilities* due to improper configurations or outdated policies. These vulnerabilities may expose the network to threats such as unauthorised access or data breaches. Enhancing security monitoring during the handover and conducting thorough security assessments can help identify and address potential risks, ensuring the network remains secure throughout the transition.
    - Poor handover processes can lead to *non-compliance with regulatory requirements*, risking penalties, and legal challenges. To prevent this, it is crucial to develop a comprehensive change management and communication plan that includes compliance checks and ensures all regulatory obligations are met during the transition.
    - *Insufficient transfer* of system and domain *knowledge* from the outgoing team to the incoming team can impair effectively infrastructure management. Establishing a shadowing period and conducting comprehensive knowledge transfer sessions are vital to bridge these gaps. This ensures that the new team is equipped with the necessary understanding and skills to maintain and operate the network efficiently.

- *Inadequate communication* between teams during the handover can result in misunderstandings and errors. Maintaining clear communication channels and providing regular updates between development and operations teams is essential to align expectations and facilitate a smooth transition.
- *A Lack of proper documentation* can make it difficult for the new team to understand and maintain the infrastructure. Ensuring comprehensive documentation of network configurations, processes, and operational procedures is critical for providing the incoming team with a clear understanding of the network environment.
- *Difficulties in tracking changes* made during and after the handover can lead to configuration drift and troubleshooting challenges. Leveraging automation tools and maintaining detailed change logs can help track modifications accurately, ensuring consistency and facilitating easier troubleshooting.
- *Manual configuration* changes during the handover introduce the risk risks of errors that can impact network stability and security. Utilising automation tools to manage configurations and changes can reduce the likelihood of human errors and enhance the reliability of the transition.
- Inconsistent network changes during the handover can lead to *integration issues* between systems. Conducting thorough compatibility testing before full deployment helps identify and resolve potential integration challenges, ensuring that all systems work harmoniously post-transition.
- *Delays on delivery* of parts of the new network, or *interoperability problems* between existing and new infrastructure may also disrupt the handover process resulting into delays or/and service degradation.
- By addressing these risks with *standardized handover processes*, leveraging *automation tools*, ensuring comprehensive *knowledge transfer and documentation*, conducting *thorough testing*, and *maintaining clear communication channels*, EuroHPC JU can effectively manage the complexities of a hyperconnectivity network services handover. Proper planning, training, and monitoring are critical to achieving a successful and seamless transition.

- Handover execution should include formal handover meeting(s), as it is a critical step in the transition process of hyperconnectivity network services. These meetings serves as structured forums to ensure that all aspects of the handover have been addressed and that the new team is fully prepared to assume operational control.
- After a completed and successful handover, the EuroHPC JU should plan for a period of post-handover support. Establishing a robust post-handover support structure is crucial for ensuring the continuity and stability of network services after the transition. This structure helps address any issues that arise during the initial phase of the new team's operations and ensures that they have the resources and support needed to manage the network effectively. Finally, for a successful handover of the hyperconnectivity network services, the EuroHPC JU should ensure that the appropriate parties' teams are fully engaged, with the necessary capacities and expertise to ensure a seamless transition and continued operation.

9. Conclusions and recommendations

Following the analysis conducted in the EuroHyPerCon study, the study team recommends Solution A as the most viable and efficient pathway for establishing a hyperconnected European HPC infrastructure. By leveraging the GÉANT and NREN networks, EuroHPC can implement a scalable, resilient, and cost-effective solution that aligns with its strategic objectives.

The EuroHyPerCon study underscores the critical importance of a timely and coordinated effort to develop a robust and sustainable network infrastructure tailored to the HPC needs. This presents an excellent opportunity for HPC providers, data providers, HPC users and network providers to acquire advanced networking services and build a hyperconnectivity network optimised for the HPC ecosystem. Through this initiative, the EuroHPC JU not only addresses immediate connectivity challenges, but also strengthens Europe's position as a global leader in scientific research and innovation for decades to come.

The recommended next steps include initiating the tender process, finalising the list of prioritised sites, and progressing towards the implementation deadline of 2025. Additional recommendations for the tenders are provided below:

- **End Sites and Ecosystem Integration:**

  The number of end sites to be connected is a crucial aspect of the two hyperconnectivity tenders. While the EuroHPC JU prioritises the EuroHPC Hosting Entities, it is essential to recognise that the ecosystem also includes a range of other European and national HPC and data centres.

  Key flagship initiatives, such as Destination Earth, depend on the efficient transfer of data across this ecosystem and to/from the EuroHPC sites. This may also involve commercial cloud providers, necessitating their direct interconnection alongside a sufficient number of Internet Exchange Points.

  As more Research Infrastructures and Data Spaces become operational, they will feed data into HPC centres for simulations and calculations. Therefore, integrating as many end sites as possible is critical for the long-term success of the hyperconnectivity network, supporting the EuroHPC strategy to create a federated and hyperconnected HPC ecosystem.

- **Risk Monitoring and KPIs:**

  Several risks and KPIs must be proactively monitored to enforce mitigation measures and anticipate the network's evolution. Monitoring will provide valuable insights, such as the number of connected sites and challenges faced by sites unable to connect.

  This data will help determine if surplus budget allocations might be available in subsequent tender(s), enabling the accommodation of additional services for end sites, such as the purchase of border routers/cards or additional equipment/configurations (e.g. Science DMZs).

  It is particularly important to note that, in the first tender, the hyperconnectivity demarcation point was set at the end site's border router, with the cost of upgrades beyond this point falling outside the scope of the first tender. This creates an opportunity for the second tender to be optimally designed to complement the first by including additional resources, especially for EuroHPC sites, maximising the network's end-to-end usage and the overall impact of the hyperconnectivity effort.