



[EuroHyPerCon](#), funded by the [European High Performance Computing Joint Undertaking](#) (EuroHPC JU), is conducting a pivotal study on the hyper-connectivity of High-Performance Computing (HPC) resources. This study, in line with the strategic objectives of EuroHPC JU, aims to examine the connectivity requirements of international and national HPC systems across Europe. Our goal is to design a robust, future-proof connectivity service, including a detailed implementation roadmap.

The anticipated outcome of this study is a thorough analysis of the connectivity needs among EuroHPC systems, encompassing other significant European and national supercomputing and data infrastructures. We aim to understand the dynamics between these systems and their potential or future users, exploring the user landscape, available technologies, and service providers. Additionally, the study will offer insights into implementation strategies, detailing the service offerings, network architecture, implementation tools, and budget considerations. Ultimately, the study will present detailed specifications for the selected network design.

To achieve these objectives, we have developed three targeted questionnaires for different stakeholders: HPC users (the one you are now browsing), [HPC providers](#), and [network providers](#). We are also engaging in various other forms of data collection, such as interviews with major users, focus groups, and a series of workshops.

We greatly appreciate your time and valuable contributions to this survey. Your input is crucial in shaping the future of HPC connectivity in Europe.

For further information or any clarifications, please feel free to reach out to us at [surveys@eurohypercon.eu](mailto:surveys@eurohypercon.eu).

#### **GDPR Statement:**

other instruments, applying data protection rules in line with the General Data Protection Regulation (GDPR), as well as with corresponding national data protection regulations. In particular, EuroHyPerCon may collect non-sensitive personal information from the survey participants such as first and last name (optional), business-email, affiliation, role, country, and related projects, along with other information relevant to the provision or usage of services related to HPC sites. The information entered by the participants are stored in the LimeSurvey related electronic systems in Germany according to its related [Privacy policy](#). An agreement between LimeSurvey and EuroHyPerCon has been prepared following the rules for EU countries and clients.

The data collected will be processed by the EuroHyPerCon consortium 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). Participants who would like to remove their data can contact us at the information below. In general, no individual data will be published. Public presentations and a public report at the end of the study may be provided aggregating and summarising technical only related data, in particular for the design of the network and related services. In case of significant requirements from big users that may require special links and services, their consent will be requested for publishing their data. EuroHyPerCon may be asked to share some data also with its contracting authority (EuroHPC JU) (see [EuroHyPerCon Privacy Policy](#)).

Data Controller: [EuroHyperCon](#), [surveys@eurohypercon.eu](mailto:surveys@eurohypercon.eu). Contact persons: Fotis Karayannis, Innov-Acts, project coordinator and Dimitrios Kalogeras, Enomix.



## Section A: Contact Details

### A1. Please provide us with your contact details:

Name 1 (optional)	<input type="text" value="Maria Girone (MG)"/>
Name 2 (optional)	<input type="text" value="Simone Campana (SG)"/>
E-Mail	<input type="text" value="maria.girone@cern.ch"/> <input type="text" value="simone.campana@cern.ch"/>
Country	<input type="text"/>
Affiliation	<input type="text" value="CERN"/>
Sector (e.g. Public Sector, Health, ...)	<input type="text" value="Particle Physics"/>
Role	<input type="text" value="WLCG Leader (SC)"/> <input type="text" value="CERN/HPC integration Coordination (MG)"/>
Are you participating in a HPC-related project? If yes, please provide the relevant projects.	<input type="text" value="EC-funded projects such as CoE RAISE, interTwin, SPECTRUM"/>

### A2. Are you answering:

As an end user   
 On behalf of my organisation   
 On behalf of a project   
 Other

Other

Please note that this survey reflects answers on behalf of both CERN Tier 0 and World-wide LHC Computing Grid.

### A3. Could you please provide us with the name of the project? (if selected project above)

## Section B: HPC Usage Survey Section

In this initial section of the questionnaire, we delve into your engagement with High Performance Computing (HPC) resources. Our questions focus on the types of HPC resources you currently utilize, those you plan to use, and your anticipated future demands in terms of compute resources. This information is vital for us to gain a clearer understanding of both the present and future needs associated with HPC usage in Europe. By collecting this data, we aim to draw meaningful correlations between HPC utilization and network usage, enhancing our insight into the evolving landscape of HPC in Europe.

### B1. 1. Which kind of HPC resources are you currently using?

*Remark:* You may select more than one answer. Sites from category c. (i.e. institutional/regional) are outside of the scope of the study.

EuroHPC JU funded HPC centres



National HPC Centres

Regional (sub-region of a country) or Institutional Centres

**B2. 2. Can you please name the HPC Centers you are currently using?**

1.
2.
3.
4.
5.

**B3. 3. Which kind of HPC resources will you likely use by 2030 and beyond?**

- EuroHPC JU funded HPC centres
- National HPC Centres
- Regional (sub-region of a country) or institutional centres
- Cannot anticipate

**B4. 4a. How often do you use HPC resources?**

- Daily
- Weekly
- Monthly
- Quarterly
- Half Yearly
- Yearly

**B5. 4b. Please specify if there is no regularity in the HPC usage. For**

**instance, you may have intense usage for a month in the context of a study and no interaction for a long period of time:**

Currently, we use HPC centres for mainly two tasks, as detailed below.

- a) CERN Tier0: Use for application development through the EuroHPC JU benchmark and development access calls (as detailed in answer 5a); AI applications and training workflows; analysis workflows.
- b) CERN Tier1s: CERN experiments' production workflows, typically through National HPC centres.



**B6. 5a. How much compute resources (either CPU or GPU nodes) do you use?**

*Remark: This corresponds roughly with the access policies of the EuroHPC JU supercomputers. Development access is currently up to 15,000 node hours, regular access up to 80,000 node hours and extreme scale access more than 80,000 node hours.*

- less than 1,000 node hours per year
- 1,000 to 15,000 node hours per year
- 15,001 to 80,000 node hours per year
- 80,001 to 200,000 node hours per year
- 200,001 to 500,000 node hours per year
- more than 500,000 node hours per year
- I don't know

**B7. 5b. In case you have more than one applications/projects you may want to specify more details below:**

This figure in 5a above includes the development requests from CERN Tier0, the use of HPC resources from activities at CERN like AI analysis and ML training.

It also includes the use of HPC centres to perform processing tasks of LHC data, including Monte Carlo campaigns. Today, up to 10% of WLCG capacity comes from HPC systems, which amounts to approximately 100000 CPU cores.

**B8. 6a. Can you estimate the computational resources you will need by 2025?**

node hours

**B9. 6b. Can you estimate the computational resources you will need by 2028?**

node hours

**B10. 6c. Can you estimate the computational resources you will need by 2030 and beyond?**

node hours

**B11. Remark: If you cannot anticipate your requirements, please leave the boxes empty. You may also provide a free text answer in this box**

We estimate that HPC centers provide about 100000 CPU cores in 2023. Projections should be made accounting for an increased need of 20%/year.



## Section C: HPC Workflow Analysis

In this segment of the survey, our goal is to gain a comprehensive understanding of your HPC-related workflows, with a particular emphasis on data management practices. By exploring both your current and anticipated future workflows, we aim to acquire a deeper insight into the evolution of data requirements over time. This understanding is crucial for us to effectively cater to the changing data needs in the realm of High Performance Computing.

**C1. 7a. Does the following scenario describe your workflow?**

**Scenario 1 “Batch processing”:** Upload input data → run job(s)/simulation(s) → download output data

- No
- Partially
- Yes
- Don't know

**C2. 7b. Does the following scenario describe your workflow?**

**Scenario 2 “Interactive - Processing with feeding data during runtime”:** Upload input data → Loop: (run job → update data or upload new data) → download output data (Examples may include data assimilation, streaming, etc.)

- No
- Partially
- Yes
- Don't know

**C3. 7c. Please describe another scenario, if any:**

The "upload input data" phase can be substituted with a "read in continuous streaming from a remote location" .

**C4. 8a. Does your scenario require taking account of security and privacy?** For example, you need to upload personal health data, and thus need to follow specific rules of data sharing.

- Yes
- No



**C5. 8b. Please provide more details about your scenario. Please specify what type of sensitive data you need to transfer** (based on GDPR or other national/EU regulations, e.g. personal, health):

Nothing we do in HEP is strictly sensitive, but our infrastructure (AAI + all the services) is AAI ready, via industrial standards.

**C6. 8c. Are there any other types of requirements in terms of security?**

Yes

No

Don't know

**C7. 8d. Please specify:**

**C8. 9. At which phase of your process do you transfer the biggest volume of data?**

*You can select more than one if e.g. upload and download are around of the same volume*

Before executing the workflow on HPC (Upload)

During the runtime of your workflow on HPC (Runtime)

After executing your workflow on HPC (Download)

**C9-14. 10a. What takes most of your time when running an HPC job: data preparation, data transfer to the HPC center, executing the application, downloading the results or analysing the results data? Please provide a rough estimation using percentages.**

*This will give a rough idea of how much time the user requires for each of the steps. But also of how the different times can be further optimized in the future. Example: 10% data preparation, 5% data transfer, 50% execution, 10% downloading, 25% analysis. Please give percentages for each stage.*

*[Percentages should add up to 100%.](#)*

**Data Preparation: 0%**

**Data Transfer: 10%**

**Runtime on HPC: 80%**

**Data Download: 10%**

**Further Data Analysis:**

**Do not know:**



**C15. 10b. If there is another challenge, in terms of time spent to execute an HPC job, please elaborate:**

**C16. 10c. In case you believe that in the future some of these phases can be optimised, please describe shortly below.**

As stated, data download could be substitutes with "access data via streaming protocols from a remote location" .

**C17. 11a. Have you encountered any challenges or bottlenecks related to data transfer to or from HPC centres?**

Yes

No

Don't know

**C18. 11b. Please specify:**

Bandwidth is generally scarce (our domain data centers are 100-1000 Gbps) routing is not optimal (some locations cannot ne reached from the compute nodes due to firewall rules) .

## Section D: Data Size

In this section, we cover aspects related to the actual data size that you need to transfer to and from HPC centres. This, together with information on your typical scenarios helps to identify (future) bottlenecks to optimize the network design.

**D1-4. 12. Following question 10, could you please estimate how much percentage of data for your entire workflow is transferred during the uploading phase (input), during loop updates (optional) and download phase (output)?**

	Percentage
Uploading Phase (input):	50
Loop (optional):	30
Downloading Phase (output):	20
Do not know:	



**D5. 13a. How much data (on average) do you need to transfer each time you execute your workflow on a HPC system?**

- less than 100MB
- 100MB to 1GB
- 1GB to 10GB
- 10 to 100 GB
- 100GB to 1TB
- 1TB to 10TB
- 10 to 100TB
- 100TB to 1PB
- more than 1PB

**D6. 13b. Please provide more details on your job scenario and execution workflow, including your scientific discipline/sub-domain:**

In the SIMULATION workflow, small generator level events are uploaded (or read via streaming) and processed.

In the reconstruction workflow, much larger data from the detector is uploaded / read via streaming. Processed events are usually smaller than the input data, and are saved to a remote data center.





**D7. 14a. Do you expect that the data size will change over time? Please try to estimate how much data you would need to transfer in 2025:**

*Please specify a percentage compared to 2023 (e.g. 150% compared to 2023).*

*Remark: Your estimations can be based on experience related to the current growth from previous years (e.g. 150% by 2025 compared to 2023).*

*Leave it blank if you cannot estimate.*

Percentage

**D8. 14b. Do you expect that the data size will change over time? Please try to estimate how much data you would need to transfer in 2028**

*Please specify a percentage compared to 2023 (e.g. 150% compared to 2023).*

*Remark: Your estimations can be based on experience related to the current growth from previous years (e.g. 150% by 2025 compared to 2023).*

*Leave it blank if you cannot estimate.*

Percentage

**D9. 14c. Do you expect that the data size will change over time? Please try to estimate how much data you would need to transfer in 2030 and beyond**

*Please specify a percentage compared to 2023 (e.g. 150% compared to 2023).*

*Remark: Your estimations can be based on experience related to the current growth from previous years (e.g. 150% by 2025 compared to 2023).*

*Leave it blank if you cannot estimate.*

Percentage

**D10. 15a. How long does the current data transfer approximately take? (for the average size of data declared in question 12). You may select more than one answer if the answer differs (e.g. depending on the network traffic).**

- less than 30 minutes
- between 30 minutes and 1h
- between 1h and 5h
- between 5h and 10h
- between 10h and 24h
- more than 24h

**D11. 15b. Please provide some additional information:**

**We use protocols like webdav, xrootd and gsiftp. The bandwidth of our data centres is (for the largest) 100-1000Gbps, so transfers can happen fast potentially, provided the HPC centre has bandwidth.**



**D12. 16a. Are you satisfied with the current data transfer time?**

- Very satisfied
- Satisfied
- OK (Neutral)
- Dissatisfied
- Very Dissatisfied

**D13. 16b. Please explain what the issue may be (if you know).** E.g. what network seems to be the bottleneck (e.g. institutional, backbone, HPC centre's network access link) or if there are differences depending on the time of the day, and state the amount of data you want to transfer in a given time:

The bandwidth of HPC centres should increase and reach the Tbps like in our domain data centres, within years.

## Section E: Data Locality

**E1-5. 17a. Please provide some more information regarding the workflows/jobs you typically execute on HPC (you may omit the question if you don't know the answer).**

*User scenarios types are derived from question 7.*

*Scenario 1 "Batch processing": Upload input data → run job(s)/simulation(s) → download output data*  
*Scenario 2 "Interactive - Processing with feeding data during runtime": Upload input data → Loop: (run job → update data or upload new data) → download output data (Examples may include data assimilation, streaming, etc.) or any other scenario you have previously specified.*

Examples:

Workflow Name	Execution Frequency	User Scenarios type	Data to be Transferred	Data Transfer recurring window
Workflow 1: Weather forecast, deterministic	4/day	1 (batch) / 2 (interactive) 3 (Other)	15TB per run	6h
Workflow 2: Weather forecast, ensembles	2/day	1 (batch) / 2 (interactive) 3 (Other)	300TB per run	12h

Workflow Name	Execution Frequency	User Scenarios type	Data to be Transferred	Data Transfer recurring window
Workflow 1: Simulation	1000s/day	batch	1-10 GB	24h
Workflow 2: Reconstruction	1000s/day	Batch	5-100 GB	24h
Workflow 3:				
Workflow 4:				
Workflow 5:				



E6.

17b. Comments (if any):

**E7-35. 18a-i. From where to where are you transferring data (add sites locations/names)**

Mark with an x (you may omit the question If you don't know the answer). Extend table if needed.

Workflow Name 1 Simulation	User / Data Site A CERN	User / Data Site A FNAL	User / Data Site A KIT, CNAF, BNL, RAL, LYON
HPC Site A: CINECA	X	X	X
HPC Site B: BSC	X	X	X
HPC Site C: KIT	X	X	X

Workflow Name 2 Reconstruction	User / Data Site A CERN	User / Data Site A FNAL	User / Data Site A KIT, CNAF, BNL, RAL, LYON
HPC Site A: CINECA	X	X	X
HPC Site B: BSC	X	X	X
HPC Site C: KIT	X	X	X

Workflow Name 3	User / Data Site A	User / Data Site A	User / Data Site A
HPC Site A:		X	
HPC Site B:			X
HPC Site C:	X		



**E36. 18j. Comments (if any):**

The production workflows would read/write from the ~200 centers of WLCG ... only a few listed in the table. Some national large data centres host the vast majority of the data, typically CERN Tier0 and the WLCG Tier1s (order of 20 sites).

**Section F: Network service specific**

**F1. 19. Which is the current network provider(s) from your user/data site(s) towards the HPC centres?**

*Please specify the first network leaving the user/data site if it is a single site, or multiple networks if it is multiple user/data sites. In case you know more information about the whole span of network providers to reach the HPC centre(s), please provide these.*

Relevant NREN (e.g. DFN, GARR, GRNET, SURFnet)

Don't know

Other/commercial, please specify:

Other/commercial, please specify:

**F2. 20. Are you using the same provider for accessing the Internet?**

Yes

Don't know

No, please specify:

No, please specify:



**F3. 21. In the case the data that need to be transferred to the HPC centre are using the same provider both for Internet and HPC services access, are you using some type of Virtual Private Network technology for accessing HPC services?**

- No
- Don't know
- Yes, please specify:

Yes, please specify:

## Section G: Other comments

**G1. 22a. Do you have any further comments?**

**G2. 22b. Please provide any attachments (documents, diagrams, etc...):**

**We would like to thank you for your participation!**

You are encouraged to **reach out to other relevant stakeholders** regarding EuroHyPerCon, forwarding our questionnaires, including other users.

EuroHyPerCon will present some initial results at the EuroHPC Summit in March 2024 and also organise a validation workshop later in 2024. In case you would like to keep in touch, please subscribe at [EuroHyPerCon Study Stakeholder Registration - Forms](#).