

CLARIFICATIONS

ECMWF/RFI/2024/365 Request for Information for a High Performance Computing Facility (HPCF) for ECMWF

Clarifications issued 8 April 2024

This document will be updated as additional clarifications are published. See version history below.

Ref: ECMWF/RFI/2024/365	
ISSUED BY: ECMWF Computing Department	
First issue:	19 March 2024
Second issue:	28 March 2024
Third issue:	8 April 2024

We are pleased to provide the following clarification responses to the questions received:

First issue clarifications:

1 Ref: C1_RFI 365

Q: Would it be possible to have an extension to the timetable for submission of the RFI?

A: ECMWF needs time to prepare a business case for presentation to its governing bodies and their advisory committees and it is important that we have as much information to inform that case as possible.

A long extension is unfortunately not possible, but by using the contingency time we will extend the date for submission to **12 May 2024 23:59 UK local time**. Consequently the initial timeline in section 1.3.2 becomes:

5 March 2024	Issue of this RFI
6 March – 28 April 2024	Discussions between ECMWF and vendors to clarify the RFI specification
29 April 2024	Last date for submission of clarification questions for this RFI
12 May 2024 23:59 UK local time	Close of RFI

No further extensions will be possible. However this is a Request for Information, extra results and information may be submitted after the closing date of the RFI.

Second issue clarifications:

2 Ref: C2_RFI 365

Q: When the tender is officially released in a latter phase, will only the RFI participants be invited to the tender or it will be open to everyone?

A: The RFI is not a qualification step for the future tender. The future tender will be open to all eligible suppliers and participation in this RFI will not be an evaluation criterion.

Third issue clarifications:

3 Ref: C3_RFI 365

Q: Can an organisation respond to the RFI and then respond to the ITT as or through a different consortium?

A: The RFI and the future ITT will be completely independent. For the ITT you will be able to respond directly or as part of a bidding group. In cases where bids involve a number of organisations which may wish to work together to deliver the contract requirements, these organisations must identify a lead contractor who will sign the contract with ECMWF and who will be responsible for all contractual obligations. All other organisations which are part of the bidding group will be

considered subcontractors. ECMWF will not enter into multiple contracts with individual members of consortia.

4 **Ref: C4_RFI 365**

Q: In case of an onsite installation in Bologna, what is the available space and power budget? Can you briefly describe the actual cooling facility and if/how is it planned to change in the future? (max cooling power, what fraction of heat needs to be removed through water and at what water temperature)? May a change in the facility cooling design be considered for the new system?

A: We expect the future systems to be made of many building blocks, so the amount of space and power available is not a relevant factor for the individual building block. However, the ECMWF data centre in Bologna can currently provide around 10MWs of power and has two data halls each with around 500 m² of space for the HPCF. Most of the heat from the current system is removed via direct-liquid cooling.

Part of the purpose of the RFI is to help define what cooling may be needed by various technology solutions. Changes to the cooling system to support the HPCF are expected and are anticipated to be part of the ITT. These modifications will form part of the business case being prepared and factored into the overall budget projections being requested.

5 **Ref: C5_RFI 365**

Q: A building block is a unit of computational or storage performance: depending on the dimension of the defined building block, could this be the union of computational and storage components as well?

A: No, not for the RFI stage. The purpose of the RFI is to give us information about the performance, power and cost of a building block. The building blocks are not meant to be self-sufficient systems. A future system will be made up of many such building blocks as well as supporting infrastructure.

6 **Ref: C6_RFI 365**

Q: Would you consider any non-x86 platform for the CPU compute?

A: Yes, certainly, our codes do not require an x86 platform. For the ITT ECMWF will not mandate any architecture or technology but will set requirements and provide a suite of benchmarks. The evaluation will be based on achieving requirements and the performance delivered on the benchmarks.

7 **Ref: C7_RFI 365**

Q: Would you consider a parallel file system solution not based on Lustre or based on different disk technologies/system architectures?

A: As in clarification C6_RFI 365, ECMWF will not mandate an architecture or technology for the storage. The benchmarks will include representative I/O. We will consider any storage system that meets the technical requirements.

8 **Ref: C8_RFI 365**

Q: When mentioning the storage building block, and the overall target performance capabilities, and order of magnitude in terms of throughput, is there an additional information, beside the sheer

throughput itself, about the expected workload characteristics (e. g. workload profiles, block size and R/W ratio)?

A: This is a complex question and will be addressed in the ITT as requirements continue to evolve, especially with the growth of AI/ML.

We have many workflows, but the most important one is operational forecast production. In this ensemble workflow many individual members of the Integrated Forecast System, using IO-server ranks, concurrently write large, well-structured meteorological fields data out to the parallel filesystem, usually with file-per-rank O_APPEND writes of several MBs per system call. Concurrent with this, a host of “product generation tasks” read such fields from these field-containing files up to available/relevant offsets to create much smaller “products”.

The purpose of the RFI is not to delve into the detail at this point, but to provide ECMWF with base information on the cost, capabilities and requirements of a future system.

9 **Ref: C9_RFI 365**

Q: Do we need to provide all the requested benchmarks, or can we submit only the ones we were able to execute? Do we need to provide all the benchmarks outputs? For future technologies, can we submit an estimate without running the complete benchmarks set?

A: The RFI is an information gathering exercise, respondents should reply with whatever information they can. This is not a tender exercise we will not be relying on any of these projections provided in good faith as part of the response (i.e. these are non-binding figures/information).

Your extrapolations on future performance and associated methodology for arriving at these estimates will provide ECMWF with very useful information.

10 **Ref: C10_RFI 365**

Q: Required benchmarks – Predicted values are supposed to be described just providing an estimated wall time / power consumption or it is also necessary to provide all the criteria / projections used to achieve those figures?

A: Your extrapolations on future performance and your methodology for arriving at these estimates will provide ECMWF with very useful information

11 **Ref: C11_RFI 365**

Q: AIFS ML training benchmark and ECTrans benchmark must be run on the same accelerators, do the available versions of these benchmarks perform equally well on different chips / vendors or can we still expect that some optimizations before the ITT?

A: The ECMWF benchmark suite is always based on our current applications. It is anticipated by the release of the ITT, the benchmark will contain updated versions, extra components and I/O as well as a requirement for many more nodes.