

CLARIFICATIONS

ECMWF/RFI/2017/001 Request for Information for a High Performance Computing Facility (HPCF) for ECMWF

Clarifications issued 19 July 2017

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

Ref: ECMWF/RFI/2017/001
ISSUED BY: ECMWF Administration Department Procurement Section
First issue: 29 June 2017 Second issue: 4 July 2017 Third issue: 19 July 2017

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

First issue clarifications:

1 Ref: C1_RFI 001

Q: In Section 2.4 - Support Requirements, Para 5 ECMWF request a full-time application software support service. Could you please clarify which applications you require support on and if home grown code such as IFS what training will you provide and in what time frame? Do you need this service, along with the requirement for a full time software support resource requested in Para 3 costed within the spreadsheet?

A: The role of the “full-time application software support” service is not to support any particular application, but rather to assist developers with migration, optimisation and debugging of their code. A person involved in providing the support service is expected to be an experienced scientific software developer with knowledge of the vendor’s hardware and software who can work closely with an ECMWF or Member State developer to get the most out of their codes or diagnose a particular issue. Prior familiarity with IFS or Member States' applications is not necessary.

Training will not be provided. However, it is expected that an ECMWF or Member State developer will work closely with the application support person to understand the issue with the application.

Please include the costing for this support in the spreadsheet under the “Annual software support” column.

2 Ref: C2_RFI 001

Q: Are the storage performance metrics provided for the slowest of either Read or Write performance (i.e. both separate 100% Read and 100% Write performance tests must be > performance requirement)?

A: The figures in table 3 should be achievable as a minimum for the sustained aggregate bandwidth of concurrently executing IOR kernels in read and write mode, for all ratios of read:write bandwidths in the range from 2:1 to 1:2. Respondents can select suitable transfer sizes and alignments.

Second issue clarifications:

3 Ref: C3_RFI 001

Q: Does ECMWF have a completion date for the newly proposed Data Centre in Bologna and what mitigations are in place if this build slips?

A: As recently announced, the new ECMWF Data Centre will be located in Bologna, Italy. Contractual arrangements are currently under discussion, but with a view to having the Data Centre completed by late summer 2019.

The project plan contains some contingency for slippage. Respondents can assume that the Data Centre in Italy is available. Should anything happen to change this assumption, it will be dealt with in the Invitation to Tender.

4 Ref: C4_RFI 001

Q: Please could you provide details on how to access KRONOS?

A: All the RFI benchmarks, including Kronos are available for download from the ECMWF ftp site.

Third issue clarifications:

5 Ref: C5_RFI 001

Q: Could you clarify what counts as a "good" correctness check for the test-of-adjoint benchmark?

A: This test serves two purposes: to measure the accuracy of the compiler and math libraries and to estimate the cost of the minimisation process in the IFS 4D-Var data assimilation application. When the test completes, near the bottom of the output can be seen some text as follows:

```
TEST OF THE ADJOINT
                                12345678901234567890
< F(X) , Y > = -.13405644731607836206E+02
< X , F*(Y) > = -.13405644731674996706E+02
THE DIFFERENCE IS 22562.436      22562.4359033512      TIMES THE ZERO
OF THE MACHINE
```

If the difference in the zero of the machine is "excessive", meaning greater than 1000 for the moment, then examine the figures. Within a run the figures for $\langle F(X), Y \rangle$ and $\langle X, F^*(Y) \rangle$ should be exactly the same up to and, preferably, beyond 10 decimal places (the lower the resolution the more decimal places should be exact). If they are only accurate to nine decimal places, there is an issue somewhere (with the Compiler being most likely). If they are only accurate to eight decimal places, then there is a serious error somewhere and the run should be considered invalid.

For runs using the same inputs, number of nodes and cores/threads then the values should be consistent across multiple runs.