

RFI Q&As

1. What is the duration of the transition plan?
A: NCEP estimates that the transition from the legacy system to the new system will be completed by January 2017.
2. What is the new system's Operational Readiness review (ORR) process and authority to assume full operational availability?
A: The new system must pass all acceptance tests and run in a parallel mode with the new system for approximately 30 days. Acceptance tests shall be developed jointly by the vendor and NCEP.
3. How will transition costs be incorporated into the evaluation?
A. All costs must be part of the vendor's proposal.
4. Will the incumbent be allowed to use existing assets to run benchmarks?
A: No
5. Regarding requirement 3.2.7 NPS on time production, given the requirements for OUT and DUT, what additional benefit does this requirement provide to NCEP?
A. On time production is a separate metric from system availability and resource contention and other system issues can cause runtime variability when the system is up. This must be avoided hence the separate requirement that the system be capable of running production with consistent times.
6. Given that product delivery requirements stated in sections 2.1.3, 2.2.1, and 3.3.3 involve both Contractor and NCEP resources and is a collaborative effort, how does NCEP determine the cause of interruptions to product delivery?
A. NCEP monitors all of the products that are produced by WCOSS and their distribution. If a product is not produced by WCOSS on schedule or delivered on time then an investigation is launched to determine the root cause of the failure. Depending on the type of failure the investigation can involve both Government and Contractor personnel.
7. Does NCEP consider the current WCOSS Phase I system (with Sandy Bridge nodes) and the Phase II system (Ivy Bridge nodes) to be a heterogeneous configuration?
A: Yes
8. Is it correct that determining the "best possible" technically compliant solution implies a cost trade-off so that a heterogeneous solution that offers consider cost savings could be considered better than a more expensive homogeneous solution?

A: Yes

9. Regarding requirement 3.3.5.3 concerning fine grain types of processors:

a. What does it mean to state there may be a requirement in the future?

A: NCEP is providing scope within the contract to accommodate recent and future technologies.

b. Does the contractor have to provide a specific statement of how this capability can be added to the proposed architecture?

A: No, but any architecture proposed must be flexible and extensible to accommodate these types of processors.

10. Is it required that the backup recovery system be based on removable storage media, given that the total required volume is relatively small?

A: The Backup /recovery System (BRS) requirements can be found in section 3.9.1 of the RFI.

11. Regarding requirement 3.10.5 Latency, What are the expectations for contractors and this requirement?

A: Section 3.10.5 provides information of the current user WCOSS experience.

12. Regarding requirement 3.19.11.1 concerning Mechanical and Plumbing drawings, what is the value and purpose of absorbing the additional cost of providing this documentation to the Government?

A: NCEP will re-evaluate this requirement prior to issuing the RFP.

13. If a WCOSS has 1,000 nodes and 400 are needed to run the NPS, is it the case that the ODT is the number of minutes when less than 4000 nodes are available?

A: No, ODT occurs when NPS cannot be run reliably within its windows. This can occur even when a large fraction or all of the nodes are available if a subsystem such as I/O is degraded.

14. For disk storage is it correct that if there is a portion of the disk storage not necessary to the running of the NPS, that resource is not considered computing the NPS?

A: Yes, that is correct.

15. Appendix C “High Watermark Chart”:

a. What is the meaning of the data in the line “38.2% Production, 21% Parallel down from 77.7% early June 2014” and how does it relate to the statement that production utilization is 59.3%?

A: In June the production machine was 77% utilized with a combination of the NCEP Production Suite jobs and parallel test jobs.

The graph shown was for the month of July and NCEP was pointing out that in July the production machine was 59% utilized as compared to June.

- b. What is meant by production and parallels given that different utilization is given for these (38.2% Production, 21% Parallel)?

A: Production represents the NCEP Production Suite jobs and parallels are test jobs run on the production machine.

- c. What codes are included in both the top and bottom parts of the chart (list of codes and the assigned colors)? Are NAM, HRRR, NLDAS, and HYSPLIT all part of the chart bottom.

A: See RFI site for updated charts for Tide and Gyre. Updated Spreadsheets with charts and data are provided.

- d. What does it mean that CPC Hurricane Seasonal Forecasts are “White Space”?

A: In the chart referenced, the graphs emanating from the bottom represent the NCEP Production Suite (NPS) jobs and the graphs emanating from the top represent parallel jobs running on the production system. The white space referenced refers to the white space on the chart which shows unused cycles. NCEP’s goal is to make more use of these unused cycles. The reference to CPC Hurricane Seasonal Forecast is an internal recommendation to run those jobs in the white space portion of the machine.

16. How does NCEP approach scheduling of nodes — specifically, are multiple parallel codes scheduled on a single Do parallel and non-parallel codes ever share a node?

A: Most jobs are scheduled on parallel nodes and have exclusive use of the full node. A small fraction of parallel jobs share nodes. These jobs use a small fraction of a node. Jobs that use more than one node do not share nodes. Serial jobs share nodes and many serial jobs can run on the same node. Serial jobs also have mechanisms to reserve a large fraction or all of a node for memory, threading, or I/O bandwidth reasons.

17. Appendix E: WCOSS Physical Description (per system). Please provide the following additional information:

- a. Is the disk space usable or raw?

A: Usable

- b. What are the Service Nodes used for?

A: Logins, and network transfers by design but also serial and small (fraction of node) parallel jobs.

c. What are the configurations of the Service Nodes?

A: Login Nodes:

6 IBM x3650 M4 servers

Two 8-core Intel sandy Bridge, 2.6Ghz

Memory: 128 GB

One 10 Gig E NIC

One Mellanox FDR HCA

OnboardGB Ethernet NICs connected to the cluster management VLAN

Redundant power supplies

Two 500 GB SFF NL-SATA drives (RAID1)

Management Nodes:

2 IBM x3650 M4 servers

Two 8-core Intel sandy Bridge, 2.6Ghz

Memory: 32 GB

6 PCIe Slots

One 4x FDR IB

OnboardGB Ethernet NICs connected to the cluster management VLAN

Redundant power supplies

Two 500 GB SFF NL-SATA drives (RAID1)

Service Nodes:

2 IBM x3650 M4 servers

Two 8-core Intel sandy Bridge, 2.6Ghz

Memory: 32 GB

One dual-port Mellanox 10 GigE NIC

One 4x FDR IB

OnboardGB Ethernet NICs connected to the cluster management VLAN

Redundant power supplies

Two 500 GB SFF NL-SATA drives (RAID1)

d. Is it correct that the “Disk Storage” is managed using GPFS?

A: *Yes*

18. What is the hardware and software configuration of the current shared storage system?

A: *Hardware (per site):*

3 IBM X3650M4 servers acting as GPFS NSD servers with 8 Gb fiber channel adapters with 10 Gb Ethernet ports

SW: IBM Active Cloud Engine

Disk Capacity: 250 TB

19. Appendix F: WCOSS Filesystem I/O Measurements:

- a. On the first chart (TP1), it appears the write rate is between 0.5 and 1.5 GBps and reads are 0.5 to 11 GBps. It appears the disk utilization is 50% – 100% busy, and total data read and written in 24 hours is 249.29 TB. Is this correct?

A: Yes, NCEP will work to provide the source data for these charts and publish them on the RFI website.

- b. Please define GPFS TP1, GPFS TD1, GPFS TMV, etc.

A: These are separate file systems with separate I/O resource devoted to them.

- c. Please explain the purpose for the different file systems.

A: One purpose is to provide some usability if a filesystem is corrupted. Filesystem repair after corruption can be very lengthy. A second purpose is to provide a guaranteed subset of I/O performance for subsets of the WCOSS work, isolating this work at the disk channel level from other WCOSS work.

- d. If a unit is not provided, can one assume the unit is in bytes?

A: Yes

- e. Do any of the file systems have SSD?

A: No

20. Who is responsible for running the NPS?

A: Government Personnel.

21. Who is responsible for the approach to scheduling?

A: Vendors can advise but the Government makes final decisions on policies

22. Section 3.3.7 mentions a minimum of 2 GB memory per core for compute nodes. What are the expectations for other memory sub-systems such as head nodes and nonscalable (graphical) workloads?

A: The amount of memory can vary on those sub-systems.

23. Are RHEL and CentOS sufficiently POSIX compliant for purposes of the WCOSS project?

A: Yes

24. What are the security requirements, if any, for the data center in which the equipment is provisioned?

A: See section 3.19.7 in the RFI

25. If a future NCEP directed upgrade causes the system to outgrow its facility (e.g., power, cooling, space, etc.), then will the facility upgrade costs be included in the cost proposals for refresh?

A: All anticipated facility costs should be included in the vendor's proposal.

26. What is the planned contractual arrangement for WCOSS II?

A: NOAA's planned contracting approach is to award a firm-fixed-price contract, with firm-fixed-price options, for delivery of the initial Phase I system and the Phase II upgrade along with operation of the system for approximately four years after initial Phase I system delivery. The Phase III and Phase IV upgrades, along with contemporaneous operation of the system, will, at the time of initial contract award, be included in the contract as unpriced options.

27. Why did the draft statement of work attached to the RFI not include requirements for the size of the WCOSS II system?

A: The RFI describes the expected budget profile for the life of the contract. Offerors are invited to propose the most capable system that can be provided for the funding available. The capability of the systems proposed by offerors will be evaluated by NOAA and will be a major consideration in deciding to which offeror to make award.