Procurements 2017

National Supercomputer Centre at Linköping University

Niclas Andersson, Technical Director
nican@nsc.liu.se

Procurements 2017

Pre-announcement:

Supercomputer **AC2018** (LiU-2017-02204)

- General computing resource for Swedish academic research
- Funded by Swedish National Infrastructure of Computing (SNIC).
- Replacing current SNIC-resource Triolith (and possibly LiUs Gamma) at NSC
- Contract size: 90 – 100 MSEK
- Installation in two phases at NSC during 2018.

Supercomputer **NWP18** (LiU-2017-02205)

- Computing resource for daily numerical weather prediction for MetCoOp and climate research for MET Norway.
- Funded by Swedish Meteorological and Hydrological Institute (SMHI) and Norwegian Meteorological Institute (MET Norway).
- Replacing Frost(Bifrost), Alvin(Allvis), and Elvis(Allvis) at NSC.
- Contract size: 35 – 40 MSEK
- Installation of two systems at NSC (or NSC/SMHI) during 2018.

Pre-announcments, Benchmarks, Request for Information and an early draft of the Invitation to Tender are available at **www.nsc.liu.se/about/procurement**

Procurement Schedule

Procurement of AC2018 and NWP18 will be performed in parallel, just 1-2 weeks apart.

- 2017-06-15 – Pre-announcement
- 2017-07-10 – Invite to participate (30 days to answer)
- *2017-08-21* – Invite to tender (40-45 days to answer)
- *2017 Oct-Nov* – Negotiation
- *2017 Nov* – Contract

Evaluation of Tenders

Application performance / Cost

- Application Performace
  - Throughput performance of benchmarks
- Cost
  - Investment
  - Installation
  - Service & Support
  - Estimated power, cooling, and facility costs over four years.
  - Additional costs (e.g. complement cooling)

Performance and Cost will be validated at delivery.

AC2018

- Will become the largest and fastest HPC resource in Sweden
- General purpose computing resource for researchers in Sweden
  - Focus on performance of the most used applications
  - Long-tail of legacy applications
- Benchmarks (reflects >50% of current usage):
  - VASP, Gromacs, CP2K
- Deliveries
  - Delivery #1: Central equipment and capacity on the level of current resource Triolith (currently 1017 dual socket SandyBridge)
    - Completely delivered and taken into production before 1 July 2018
  - Delivery #2: Extending the capacity.
    - After Triolith has been decommissioned
  - Optional:
    - Additional capacity, additional resource

AC2018 Technical Details

- System Software: NSC Cluster System Environment (CSE) based on CentOS.
- Storage: Connect to existing storage
  - 4 x GSS26 (eight servers) from IBM/Lenovo running Spectrum Scale, Native RAID (GNR). Today connected to Triolith with IB FDR.
- Density: Minimum of 15 kW per (normal size) rack (peak).
- GPU: X% of the budget can be used for servers with GPUs if they achieve better performance on the VASP and Gromacs benchmarks.
  - A few GPUs for development.
- Large Mem: Approx. 5% of computing servers equipped with larger memory.
- Possible architectures:
  - CPU: Intel Xeon, Intel Xeon Phi, AMD Epyc, IBM Power9, ARM, …
  - Network: Mellanox FDR, EDR, Intel OPA, Eth RoCE ...
  - Cooling: Air, Water, Immersed

NWP18

- Main focus: Meteorological Co-operation on Operational Numerical Weather Prediction (MetCoOp)
  - Collaboration between SMHI, MET Norway and Finnish Meteorological Institutet (FMI).
  - Will be operative in late 2018.
  - NWP18 will be used together with a computing resource at FMI as a joint resource for HARMONIE-AROME Meteorological Ensamble Prediction System (MEPS) for the common regions of Sweden, Norway and Finland
- Benchmarks:
  - HARMONIE-AROME, NEMO
- Deliveries
  - System #1: Installation at NSC
  - System #2: Installation at NSC or SMHI.

NWP18 Technical Details

- Two systems (System #1 at NSC, System #2 at NSC or SMHI):
  - System #1, part A: Daily Numerical Weather Prediction (MetCoOp)
  - System #1, part B: Climate Research (MET Norway)
  - System #2: Daily Numerical Weather Prediction (MetCoOp)
- System Software: NSC Cluster System Environment (CSE) based on CentOS.
- Storage: Connect to existing storage
  - 1 x GSS26 (eight servers) from Lenovo running Spectrum Scale, Native RAID. Today connected to Triolith with IB FDR.
  - Additional storage will be procured separately
- Density: Minimum of 15 kW per (normal size) rack (peak).
- SMHI computer room has a upper limit of 125 kW.
  (system 1A and 2 has to be diveded accordingly).
- Possible architectures:
  - CPU: Intel Xeon, Intel Xeon Phi, AMD Epyc, IBM Power9, ...
  - Network: Mellanox FDR, EDR, Intel OPA, Eth RoCE
  - Cooling: Air, (Water)

Preliminary Requirements

(work in progress)

Service & Support

Requirements

- Installed equipment should be working and provide its indended service.
- Reduce and limit the time personnel at NSC spend on handling typical, recurring errors and failures (e.g. memory faults, PSUs failures) during the life-time of the system.
  - This include time spent on fault diagnostics, error reporting, ticket handling, logistics of parts (RMA handling), waiting for response, etc.
- The contractor is responsible for all advanced troubleshooting of equipment and components.
- Efficient escalation management to quickly resolve systematic errors

Service & Support

Service Levels

- Central or critical equipment
  - Spares on site and 4h or 8h response.
- The bulk of computing capacity
  - Minor faults affecting one server: replace within a week.
  - Major faults affecting half a rack or more: next business day

Deliveries (suggestion)

A Delivery

- Shipping and installation: 2 weeks
- Functional Tests: 2 – 4 weeks
  - acceptance testing
- Stability Test-period: 2 months
  - uptime, counting faults

Deliveries

- Pre-delivery of solution samples for testing and porting
- Tetralith #1: March – June
  - Core parts + replacing Triolith capacity
- Tetralith #2: July – October
  - Extending capacity
- NWP18 System #1: May – August
  - Installation at NSC (about 60-70% of NWP18)
- NWP18 System 2: June – September
  - Installation at NSC or SMHI