Arrhenius Procurement
Gert Svensson
Dep. Technical Director, NAISS

Arrhenius Procurement

- **Steering Group**
  - Björn Alling
  - Patrik Norman
  - Erik Lindahl
  - Niclas Andersson
  - Representative sensitive data

- Project leaders
  - Gert Svensson
  - Andreas Johansson

Procurement type

- Negotiated procedure with prior publication
- Two or several step
  - Invitation to demonstrate capacity to deliver such a system
  - Shortlist
  - RFP
  - Second step proposal submission
  - Possible negotiations
  - Possibly changes to the RFP
  - Third step proposal submission
  - Possible negotiations
  - New proposals
  - Contract

Evaluation

- Based on Total Cost of Ownership (TCO)
  - Purchase costs
  - Maintenance costs
  - Power and cooling costs

  - We give the max TCO for 5 years
  - Mandatory requirements
  - Optional requirements maybe
  - The vendor that can run the benchmarks fastest wins.

Tentative Time Plan

Funding and Agreements

2023

- Hosting Agreement
- Secure funding
- Start preparing site
- Procurement Agreement

2024

- Grant Agreement

Procurement

2024

March: Invitation to Participate

April: Selection of Vendors

May: Invitation to Tender

Aug: Tenders

Sep – Oct: Negotiation

Nov: Award + Contract

2025

Q1: Test system

Q2: Starting delivery

Benchmarks

Purpose

- Design, size and configure
- Evaluate
- Accept
- Measure performance and power

Preliminary Candidates

- Applications
  - CFD: Neko
  - MD: Gromacs
  - QC: VeloxChem
  - Climate: ICON
  - AI: Not yet decided
- Synthetic
  - HPCG
  - HPL (single node)
  - STREAM
  - IO-500

Trends in HPC Architecture

- More cores 128 → 256 cores/node
  - Less memory bandwidth per core
- Coupled CPU and GPU:s (cache coherent)
  - May need mix of CPU jobs and GPU jobs.
  - Fixed ratio between #cores in CPU:s and #GPU:s
- Everything with AI is highly expensive and possible with long delivery times.
- GPU:s less optimized for 64-bit floating point calculations (due to AI)

NAISS

General questions

- Size of partitions?
  - Proposal: 1000 Dual CPU nodes 250 GPU nodes (4 GPU:s/node)
- Do we anticipate application to span different partitions?
- Compatibility with x86, NVIDIA? (Third party applications)
- How much should we optimize for AI-applications vs simulation?

More general questions

- How much should we value fast delivery time?
- How about application where manycore is not optimal (data driven, bioinformatics, neuroscience etc.)?
- What is important for your science area?
- Different access methods
  - Batch
  - Cloud interfaces
  - Interactive use

Memory size

- Different RAM sizes?
  - New faster memory
  - More cores
  - Keep memory BW/core constant?

Storage

- Node vs network attached?
- How can we take care of increasing data sizes?
- Tiered data?
- Why isn't Swestore (dCache) more used?
- AI requirements?
- Should NAISS look into long term storage/archiving in the future?

NAISS

Sensitive data and Cloud

- Should sensitive data and cloud services applications run on the same hardware as the rest?
- On a fixed partition?
- Industry applications on the same partition?

Discussion