Information Day
Technical Aspects

7 September 2016

Helge Meinhard

CERN IT department

Helge.Meinhard@cern.ch

All information contained herein are for discussion purposes only and shall not be considered a commitment on the part of CERN or the Buyers group

Helix Nebula – The Science Cloud with Grant Agreement 687614 is a Pre-Commercial Procurement Action funded by H2020 Framework Programme

Outline

- R&D aspects
- Prototype and Pilot installations
- Virtual Machines
- Shared storage
- Networking

Highlights of questions asked rather than a complete overview of the technical aspects of the project

R & D

- Workload uses **POSIX file I/O subset:** open, close, read, write, seek

- **Transparent data access:**

  - Such workload can run without modifications on contractors' as well as buyers' resources
  - Workload to "see" global view of files and data sets
  - Streaming as well as pre-fetching to be supported
  - Data movement to be done automatically and without user intervention (no "managed storage pool")

Prototype and Pilot Installations (1)

- **Purpose:** test PCP concepts with test suites and real workload

- Intended **scale** as an order of magnitude**:**
  - Prototype: 3’500 cores, 350 TB shared storage, 10 Gbps per supplier over three months flat
  - Pilot: 10’000 cores, 1 PB shared storage, 40 Gbps per supplier over eight months flat
  - **No pay-per-use**
  - Expected to be realistic

- Pilot is intended to represent ~5% of production facility, hence solutions should **be scalable** by factor ~20

Prototype and Pilot Installations (2)

- **Out of scope** of HNSciCloud:
  - Potential resource requirements beyond pilot
  - Process optimisation, sharing, common interfaces across contractors and/or buyers
  - Installations at buyers' sites
  - GPUs and other accelerators

- **Workload:** primarily batch processing; some HPC
  - List of workloads: http://www.hnscicloud.eu/hnscicloud-user-groups

Prototype and Pilot Installations (3)

- **Redundancy and High Availability:**
  - Requirements on infrastructure, storage and VMs to be clarified in call-offs for prototype and pilot phases
  - No special requirements on services the buyers will establish on contractors' resources

Virtual Machines (1)

- **Hyperthreading / symmetric multi-threading** can be enabled on hypervisor
  - Beyond that, no over-commitment of CPU and RAM
  - Call-offs for prototype and pilot phase expected to specify minimum performance

- **RAM:** Most VMs expected to be requested with 4 or 8 vCPUs and >=1.875 GiB or >=3.75 GiB per vCPU
  - These are minimum values – larger is compliant, smaller is not

- **Local VM storage:** 30 GB per vCPU of several performance classes

Virtual Machines (2)

- **HPCaaS** expected to represent less than 5% of resources
- **OS:** Images for CentOS 6 and 7, Debian 8, Windows server 2008R2 and 2012 (with licences) to be provided by contractor; support for currently supported versions of Linux and Windows
  - Windows expected to run on less than 5% of resources
- **Configuration:** Responsibility of Buyers' group; contextualisation methods required at least at cloud-init level
- **Services and applications:** Responsibility of Buyers' group
  - Contractor responsible of Container Orchestration Engines
  - Application licences, if any, will be provided by buyers

Shared Storage (1)

- In addition to local VM storage
- To support **streaming** as well as **pre-staged** access of data sets
- **Data sizes:**
  - Average: O(10 GB) / file, O(10’000) files / data set
  - Up to ~300 GB files
- **Throughput:** up to 2 GB/s per stream in pilot phase
- **Latency:** requirements to be clarified in call-offs for prototype and pilot phases

Shared Storage (2)

- **Lifecycle:**
  - To be instantiated at beginning of phase
  - Available throughout the phase
  - Data persistency beyond project phase not required
- **Data sharing** across OS (between Linux and Windows) not required
- **Head nodes,** if any, to be provided in addition to compute capacity

Shared Storage (3)

- **Data types:**
  - Bulk of data is non-SQL
  - Mostly proprietary formats such as ROOT files, images
  - XML files
- **Metadata** largely workload-dependent
  - WLCG: Oracle

Networking

- Contractor responsible for **connection to GÉANT**
  - See https://www.peeringdb.com/net/3386
  - 10 Gbps aggregate between contractor and GÉANT for prototype
  - 40 Gbps aggregate between contractor and GÉANT for pilot, possibly over multiple 10 Gbps circuits
  - Pipes expected to be mostly filled
- **Data sources and destinations:** mostly, but not exclusively buyers' in-house data centres
- **Public IPv4 addresses** as required by specifications to be provided by contractor

HELIX NEBULA
THESCIENCECLOUD

Thank you
Questions?

CERN