

Attendees:

Local: Tim, Maarten, Alessandro, Ulrich, Stefan

Remote: Jan Just, Dirk, Victor, Andreas

Apologies: Manfred

News:

Manfred: We do have an SGE implementation for machinefeatures running at GRIDKA

Use Cases (ctd from last meeting)

LHCb (Stefan)

Use cases:

- MC Factory (aka job masonry), i.e. filling the "end of a queue" with useful jobs, e.g. monte carlo simulation of X number of events which fit until the queue length expires

- In multi-core environments discover the number of available cores in order to run any number of single- multi-processor using payloads that are deemed to be useful

Question on whether a running event processing application can be interrupted at any time? Gaudi (LHCb core software framework) is instrumented to catch signals and to interrupt the event loop at receiving a signal. The most common use case should be the usage of simulation jobs. It's probably not that easy to be used for "real data" processing, as the underlying bookkeeping might not be suited for this (e.g. double output file for a single input)

Tim: what is the processing time / event? Stefan: in the order of minutes.

General Discussion on next steps:

Definition: different resource systems that could possibly be covered by the taskforce:

- physical "bare metal" batch systems: ie. payload is executed on a hardware host without virtualisation (what most grid sites are doing now)

- "virtualized worker nodes": The worker node itself is a virtual image but used within a conventional batch system (e.g. slc6 at Cern, CNAF, ...)

- "private clouds": IaaS system running in the context of WLCG (e.g. openstack at CERN)

- "public clouds": IaaS system running outside the context of WLCG we cannot (or little) influence any changes (e.g. amazon EC2)

- "opportunistic computing": deployment of a virtual image on a remote host which is self managed (e.g. BOINC (LHC@Home))

- Stefan: Goal of this taskforce, in a first approach, shall be to come up with the minimal needed implementations. Later on we can go into enhancing the system with more details if decided to be useful

- Tim proposes to not rely on the current file implementation but instead to provide a data structure (e.g. json) which can be used by the VOs to discover the

features.

- in case of batch systems this will rely on the file system. In case of batch systems a simple tool needs to be developed that returns this data structure

- Question on how to communicate the feature information in a virtualized environment?

- Discussion on how to provide the information on IaaS systems in a dynamic way. Suggestion to use the "magic IP" to provide the information, but not clear if dynamic information can be conveyed by this mechanism. We need to test this.

- In case of virtualized WNs we need to understand how to get a meaningful number for the power

- Ulrich suggests that we can have a number provided as "minimum guaranteed" power that the site will provide on each WN, Stefan argues that this is not much more than we currently have with the BDII, can we do better than this?

- Stefan asks how many IaaS system types (e.g. openstack, opennebula, etc.) we would need to cover? Do sites already converge on some systems? Tim: no,

- Alessandro: was it possible with the system to retrieve the features of all jobslots at a given site at the same time? (e.g. talking to the nova controller?). Discussed but we may want to keep this feature for later.

- Dirk: Can we do the reverse? ie. push information to the resource provider about the payload (e.g. I'm finishing in x minutes). Discussed but decided to keep this as a possibility for later enhancement

Action Items:

- Define the structure of the features to be returned to the VO payload (e.g. json)

- Define and implement a tool which gathers the information on a bare metal batch system

- Deploy the system for physical batch systems on multiple grid sites and have it tested by VOs

- Understand the power number provided on virtualized WNs

- Test if we can provide dynamic information via "magic IPs" on IaaS systems

Aob:

- Next meeting Tue 24 Sept, 14.00 (agenda and venue to be announced on the mailing list)