

CMS position on whole nodes, virtualization and Clouds

Claudio Grandi
(INFN Bologna)

- CMS is interested in a resource allocation mechanism that allows executing a process in **user space** on a **whole multi/many-core host**
 - Clear benefits in term of memory consumption
 - Still to be understood the implications for the local I/O
- We expect to have more answers by the work of the *Whole Node Task Force*

- CMS is not interested in virtualization *per-se*
 - It is mainly a site business: virtualization may help a site configuring resources in a more flexible way
 - CMS is happy with real nodes as far as there is agreement on the supported OS distributions
- CMS has nothing against sites using virtualization provided that:
 - the performances of the **virtual host match the CMS requirements** in terms of bandwidth, latency and aggregated throughput
 - monitoring of the application is not impacted

- Commercial clouds may be used by CMS e.g. for simulation but are currently too expensive
 - See e.g. CHEP'10 poster: A.Melo *Using Amazon's Elastic Compute Cloud to scale CMS' compute hardware dynamically*
- The use of a **Cloud interface** to access the WLCG resources is a possibility but depends on the implications
 - CMS prefers having **efficient** (local) **access to the site storage** in user space rather than having root access on a machine in the DMZ
 - Current **Grid interfaces** may be used to access whole nodes

- CMS is interested in accessing **whole nodes** where the processes may run in user space with **high efficiency** and **optimized access to the storage**
- Provided the above requirements are satisfied
 - CMS doesn't have strong requirements on the **interface** used to access the resources
 - CMS doesn't have strong requirements whether **virtualization** is used on the allocated resources
- Commercial clouds are usable for specific tasks but too expensive