

Workshop Discussion Summary



Denise Heagerty
26 June 2009

Overview

- **Requirements**
 - ATLAS, LHCb
- **Exploiting multi-core**
- **VMs: conserving IPv4 address space**
- **Security: trusting images**
- **Batch services: scheduling and accounting models**

Requirements: ATLAS

Extract from talk by Paolo Calafiura:

- ATLAS ready to transition multicore and virtualization prototypes to production**
 - WLCG should support multicore job submission**
 - CERN should add KSM to SLC5**
 - CernVM should transition from R&D project to official CERN platform**
 - Concerned about the number of VMM technologies**

Requirements: LHCb

Extracted from talk by Thomas Ruf:

- **Parallel processing should become fully supported for interactive and batch / Grid processing**
 - Only way to gain from processor developments
- **Virtualization seems to be the most efficient way to become independent of the local operating system**
 - Should be very attractive for the Grid, where changing from SLC3 to SLC4 already caused a lot of pain
 - From some LHCb institutes came the feedback that their system administrators don't want to know about it for shared machines.

Discussion: Requirements

■ Request for tool deployment

- Performance and other tools are now required by experiments
 - E.g. KSM, PERFMON, ...
- A test infrastructure is being prepared by IT Dept as a first step
- Experiments can request deployment

Discussion: Exploiting Multi-core

- **Could batch resources be dedicated for parallel jobs?**
 - Needs to be balanced with other requirements
 - Proposed to first test LSF functionality for mixed workload
 - E.g. reserving x cores on a single host
 - CPU should not be the only measure of batch resource utilisation
 - E.g. Jobs may be memory or I/O bound

Discussion: Security

- **What is needed for sites to trust VM images?**
 - Confidence that image has not been 'tweaked'
 - Traceability and integration with site logging
 - An agreed (security) update mechanism
 - Sandboxing?
 - No root access?
 - Off-site access only via proxies?
 - ...

Discussion: Conserving IP address space

- **Is private IP addressing feasible for worker nodes?**
(VMs significantly increase use of IP addresses)
 - Traffic needs to be understood
 - Proxy servers can provide a gateway for off-site traffic
 - Sites may be reluctant to run non-standard proxies (e.g. squid ok)
 - Past experience showed private addresses were unacceptable

Discussion: Batch Model

- **Do we need a paradigm shift?**
 - Move to a cloud-style reservation scheme?
 - Infrastructure as a Service?
 - Scheduling is still required to decide where to run a VM
 - Split resources between traditional batch and supplied VMs?
 - Run a batch scheduler inside a cluster of VMs
 - ...
- **Is CPU usage the right accounting model?**
 - Take account of memory, I/O, infrastructure (power etc), ...
 - Wall clock time + charge factor per CPU?
 - Slot time?