

Tier-0 Resources

Helge Meinhard, CERN-IT
WLCG Management Board
14-Jul-2015



Outline

- Hardware status
- VM performance
- ATLAS tier-0 issues

Hardware Status

- Disk pledges: fully installed and available beginning 2015
- Tape pledges: idem
- CPU pledges
 - 27'000 cores delivered rather late
 - After burnin, discovered that firmware upgrade was required – done during May
 - Availability uncomfortably close to start of LHC data taking
 - Now fully commissioned and in operation
 - Hardware pledges fully met

VM Performance (1)

- Most CPU pledges provided as virtual machines
- Compute cells: CPU pass-through, KVM caching set to write-back
 - KVM caching gains two orders of magnitude in IOPS
 - Precludes live migration
- Overhead from virtualisation larger than expected
 - See following results presented at a CERN-IT-internal meeting on 13-Jul-2015

VM Performance (2)


- Arbitrary hardware type with 32 cores, no further tuning (“type 1”)

HWDB HS06	VM Size (cores)	Per VM HS06	Total HS06	Overhead
357±16	4x 8	82.3±11	329	7.8%
	2x 16	150±5	300	16%
	1x 32	284±11	284	20.4%

VM Performance (3)

- Numerous handles to turn

HS06 - Percentage Overhead and HS06 ratings for full node VM (32x core Intel)



VM Performance (4)

- Tuning depends on exact hardware type

Type 1 optimised:

Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz

HWDB HS06	VM Size (cores)	Per VM HS06	Total HS06	Overhead	Overhead Reduction
357±16	4x 8	87±11	348	2.5%	68%
	2x 16	163.5±1	327	8.4%	52%
	1x 32	311±1	311	12.9%	37%

Type 2 optimised:

Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz

HWDB HS06	VM Size (cores)	Per VM HS06	Total HS06	Overhead	Overhead Reduction
344±1	4x 8	84±7	336	2.3%	65.7%
	2x 16	161±2	322	6.4%	22%
	1x 32	279±4	279	19%	13.6%

VM Performance (5)

- Optimised settings being rolled out across compute cells
- Given core distribution of VMs, overall performance hit is ~7...8%
- Reason for larger performance hit on large VMs (counter-intuitive!), and for hardware dependence, not understood yet
- Why use large VMs at all?
 - More flexible in context of multi-core (in particular 8-core) jobs
 - LSF limit of 5'000 worker nodes per instance

VM Performance (6)

- Even though hardware fully deployed to meet pledges, shown effect means some ~10% shortfall in usable capacity
- Some spare capacity being configured now, likely to get close to covering the shortfall
- More CPUs (some 10'000 cores) coming up for delivery in September, likely to become available in October
- Working on understanding the large performance hit, and on developing alternative scenarios in view of 2016

ATLAS Tier-0 issues (1)

- Separate LSF instance
 - Protection against unforeseen user behaviour on shared instance
- Started with 24-core workers with 480 GB of SSD disks
 - Very little swap space: jobs killed frequently
 - Tier-0 jobs need more than 2 GB in initialisation and termination phase
 - Memory leak in application, fixed meanwhile
 - Configured swap at expense of /pool: jobs killed frequently
 - Not enough space in /pool
 - Significantly improved pool cleaner, disabled core files

ATLAS Tier-0 issues (2)

- Instance beefed up by new 32-core machines deployed as 2 x 16-core VMs
 - Performance ranging from 60 to 155 HS06
 - Long (and painful!) investigations showed that CPU pinning introduced with OpenStack Juno was buggy – with 50% chance both VMs would be pinned to one and the same CPU, leaving the other idle. CPU pinning now deactivated
- VM swapping causing up to 30% I/O wait
 - See before for reasons and remedy
 - Multi-core jobs much better
- Hypervisor swapping, causing I/O wait
 - Lack of memory on HV with KSM off due to memory remnants of KSM switched on before, fixed by re-activating KSM

ATLAS Tier-0 issues (3)

- Followed up by very regular (sometimes daily) meetings and phone conferences
- Very good collaboration between ATLAS and the CERN-IT teams – big thanks to ATLAS
- Situation now stable and workable

Summary

- Hardware fully put in place – disk and tape in good time, CPU rather late
- Discovered unexpected significant performance hit on large VMs
 - Tuning gets it down
 - Capacity to compensate is on its way
 - Efforts continuing
- Long and painful (for ATLAS and CERN-IT) way to fully commission ATLAS tier-0 resources
- Situation now stable and workable

Questions?

