Site Survey

Michel Jouvin

January 24, 2007

WLCG Collaboration Workshop, CERN

Agenda

- Site survey goals and answers
- T2 size
- T2 administration
- T2 resources
- Conclusions

The Survey

- Follow-up for T2 survey made mid-June 2006
  - 45 answers, compared to 33 in June
  - Mainly T2s, (1 T1, 1 T3)
- Goal : get a picture of what T2s are exactly, their differences, their problems, their readiness…
  - Not formal : will not be used to see if matching MoU
- Broader country coverage
  - Mainly Europe
  - USA Canada
  - Israel
  - Asia/Pacific : Japan, China, India, Australia, Taiwan

VO coverage and FTE

- LHC : range from 1 to 4
  - All combinations, mainly 1 or 4
  - Alice is the less supported…
- Mostly support for non LHC VOs as well
  - Big differences in the number of non LHC VOs supported
    - A few sites dedicated to 1 or 2 LHC VOs
  - Generally lower priority in term of resource access
  - Other HEP, biomed, esr, national VOs, local VOs
- Number of VOs supported has generally an impact on site size and staff needed to run the T2
  - From 1 to 10 FTE, majority between 2 and 6 FTE
  - Not always related to T2 size (at first glance)
  - Federated T2s tend to have more FTEs

- Big differences between T2s in resource size and CPU/TB ratio
- Resources planned at LHC startup
  - CPU (kSI2K) : 400+ if 1 VO, 800+ if 4 Vos
    - 3 T2s plan 2500+ kSI2K
  - Disk : 50 to 800 TB !!! (not related to number of VOs)
    - Some T2s probably devoted to MC
  - Network (external) : mainly 1 Gb/s today but several pla 10 Gb/s
    - Evolution compared to previous survey
    - A few T2s with < 0.5 Gb/s (e.g. Australia)
  - Generally significant ramp up planned in the coming yea
    - Some T2s at 10% or less of their final level, probably new ones
- (Almost) No MSS planned
  - Some exceptions (Spain), but generally comparable to d space

- Number of sites making the T2 : 1 to 8 !
  - Site : geographical
  - 1 site : ½ of (answering) T2s
  - Multi-site : majority between 2 and 5
  - Number of sites seen by the MW : generally 1 / site
    - A few exception
- Largest T2s are federations
  - Italy : all T2s are 1 site and support (mainly) 1 VO
  - Several countries have only one T2 made of federations
  - Related to local/national configuration : lot of small labs vs. large universities/institutes
- Impact on experiments not clear
  - Fragmented SE space (not always very large, <100TB) for analysis

OS / MW versions

- Much more homogeneous…
- OS : SL(C)3 32-bit mainly, asking for SL4 64-bit
  - Vast majority using SL/SLC, also Debian/CentOS (2)
    - But lot of sites still using <= SL 3.0.5
  - Interest in SL4 64-bit expressed by a large number of sites
    - Several sites already running 64-bit for testing or productions
    - Several asking 64-bit support for SE (kernel 2.6 perf improvements
- MW : gLite 3.0 everywhere (except US)
  - LCG part + some tests with gLite CE/WMS
  - 2 sites still running LCG 2.7
  - US using OSG 0.4
  - Sites feel confortable with continuous release process
  - General feeling is that MW is more stable than 6 months ago
    - Not so many (new) requirement requests…!!!

T2 Administration

- Mainly “distributed administration” = each site independently
  - Not always a unique (consolidated) answer to the survey
  - Often a technical coordinator able to act at each site
  - A few sites thinking about cross site logins : ssh, gsissh, sudo… for better support coverage (e.g. holidays)
  - Sometimes, vendor tools used (mainly installation)
- Deployment : site independence mainly
  - Sometimes agreement of minimum set of tools
  - A few exception : deployment managed by Quattor from unique repository
  - Mainly YAIM (+KS), a few Quattor or local/vendor tools
    - Quattor usage/interest increasing, in particular w/ QWG templates
  - Generally, but not necessarily, same batch scheduler or product

- Most common configuration = 1 CE / site
  - No CE spanning sites (some expression of interest : 3)
  - Sometimes several per site, e.g. 1 CE / VO
    - 1 site with 1 CE per type of HW
  - Generally not seen as problem : let MW / experiment SW deal with the situation
- LRMS : mainly Torque/PBS w/ or without MAUI
  - Several Torque v2 already (S. Traylen distribution)
    - Default in Quattor QWG (support for MPI)
  - LCG : several SGE, 1 Condor, 1LSF
    - Condor : better integration with MW generally requested
  - OSG (US) : Condor
  - Fairshare used at a large number of sites but not all
    - Why ?
  - Job priorities inside VO : ½ sites say they don't enforce them…

T2 SE

- Only ½ answered questions about SE : difficult to interpret
  - Answers : 2/3 using DPM, 1/3 dCache, Castor in Spain
    - A few (3-5) sites moved from DPM to dCache but not a trend
    - Assessing DPM future (not only support) remains critical…
  - Not always consistency inside a federated T2
  - Still a few Classic SEs (~5)
- 1 SE / site everywhere (almost)
  - 1 T2 with 1 SE / VO
  - Somes sites with 2 SEs, generally because of migration
  - A few sites planning for a unified SE across a federated T2
- A few sites have not yet made their final decision fo SE
  - Particularly those still running Castor1

T2 Helpdesk and Support

- Helpdesk : diversity
  - 1/3 have or plan a help desk, 1/3 rely/participate to national helpdesk, 1/3 with no grid support
  - Sometimes rely on GGUS (or OSG)
  - Can be not very formal (best effort)
- Support : from 0.5 to 3 FTE
  - Coverage not mentionned
  - Some T2s : participation to national helpdesk

Conclusions

- Picture is less complex than 6 months before
- Majority of T2s participated to SC4 and related activities
  - T2s who didn't participate are generally small and new sites and thus require some attention
- Several T2s have a significant ramp up to achieve i the coming year(s)
- Sharing management tasks between sites is considered by an increasing number of T2s
  - In particular inside federation
- MW is seen as more mature and more stable
  - Main request : better integration w/ monitoring tools like Lemon and Nagios