

Integration Program Status

Rob Gardner

10-12-10

US ATLAS Facilities Meeting

SLAC

Evolving the Integration Program

- Starting now Phase 15 (FY11Q1) - our fourth year of Integration work!
- Core principles remain the same, tweaked to evolve with changing ATLAS requirements
 - Integrate all Facility components in the US Cloud (Tier 1, 2 3)
 - Phased program over 3 months
 - Includes fabric, grid, ATLAS services
 - Capacity, reliability, performance
- However focus has shifted to usage & access (both CPU and storage)
 - We have some well-established analysis usage patterns at the Tier 1,2 now (pathena,prun & fetching of job output)
 - With Tier 3 may come new modes requiring better and global access to storage

Phase 15

“We got some work to do”


- CREAM CE
- Analysis Triage
- Network & Throughput
- Tier 3 integration
- Global data access
- CPU & disk capacity upgrades

CREAM CE

- Anticipate CREAM CE update or significant integration testing during Phase 15
- CREAM-CE as a replacement to the LCG-CE, OSG-CE (Globus 2 based). Used in production >1 year
- No plans to deprecate OSG-CE, but all LCG-CE to be decommissioned by end of year
- Widely discussed in ATLAS ADC, see Graeme Stewart's talk, etc
 - <http://indico.cern.ch/conferenceDisplay.py?confId=107126>
 - Pilot factory submission issues have been addressed in a recent condor release 7.5.3
 - Sites are being encouraged to update
- Integration into OSG in progress - schedule later this week

Analysis Triage (Site Level)

- # analysis jobs in the facility continues to grow
- # analysis job types is unlimited and constantly changing
- Required storage access increasing
- # analysis job slots increasing



- Strategy
 - Sort failures by job set/user
 - Distinguish site-user errors
 - Contact users if necessary, cc DAST

Tier 3's and the Facility

- New Tier 3's coming online
 - Focus is on “T3g” type
 - See <https://twiki.cern.ch/twiki/bin/view/Atlas/Tier3gSetupGuide>
- As Tier 3 program evolves and requirements change we need to examine points of contact and required support from the Tier 1 and Tier 2 centers
- Immediate areas
 - Panda-Tier3 & worker-data affinity
 - Managed dataset transfers via dq2-FTS
 - LFC and local data management

Network & Throughput

- Shawn will describe in detail tomorrow
- Expect all Tier 2's to update to latest PerfSonar release (we'll have a demo from Jason)
- Bring the infrastructure into a routine “go-to” operational mode for periodic monitoring, troubleshooting, trend analysis, etc

Global Data Access

- In the US Facility (today) we have over 12 PB of disk deployed, over 5 PB deployed at the Tier 2s
- Access to this data comes is provided by a number of services, and clients for the various processing tasks supported (production, analysis, reprocessing, tag-skim, etc)
- User level access is a particular challenge
 - Local area versus wide-area
 - Direct (network reads) versus stage-in of whole files
 - User resource characteristics (storage caching) and modes
- WLCG demonstrator finding will be key in identifying directions forward

Summary

- Major accomplishment over past several months has been deployment storage capacity to meet the production and analysis needs of ATLAS this past summer
- Analysis needs met with good performance - though challenges remain
- In Phase 15 of Integration Program we likely see some major site-level upgrades
 - CPU & disk upgrades
 - CREAM CE (evaluation, +)
 - Network monitoring
- Lots of new work ahead - equipment, T3 integration, data access, ... 😊