

USATLAS Facilities

Tier2 Model, Target Capacities and Architecture

Overview

- Computing Model:
 - Tier2 Functions
 - Tier2 Capacity commitments and targets
- Tier2 Services.

Tier2 Functions

- Provide computing resources to ATLAS collaborators for Monte-Carlo simulations and data analysis (mostly AOD based)
 - Include data storage and data management services for the above functions.
 - Include network resource management.
 - All the activity is performed non-interactively (using PanDA).
- Provide *flexible* means for role/individual/group/project authentication, supporting multiple VOs.
- Provide controls for *flexible and efficient* resource allocation to projects and individuals.
- Provide tech support to (US)ATLAS and OSG VOs.
- Provide mechanisms for monitoring and usage accounting.
- No interactive user work, no direct user support. (These functions are delegated to T1 and T3s)

T2 Functions: CPU resource provision

- Access to the resource:
 - Support for Grid based access (OSG Gatekeeper) and submission from local service (PanDA pilot generator).

T2 Functions: CPU Resource Provision

- Control of resource utilization – Allocation policy service:
 - Ctrl of throughput – services that enforce policies based on cumulative and instantaneous metrics (max # of jobs, max total CPU or wall time, decay constants).
 - Ctrl of latencies – services that enforce access rules (preemption mechanisms)
 - Priorities based on VO/project/group/user/utilizationHistory
 - Service agreement policies (LSF)
- The solutions are implementation specific. Does it make sense to not standardize?

T2 Functions: Storage Provision

- 2 types:
 - Central, flexible, high-performance (NFS, IBRIX, ...?)
 - Used for applications, scratch space and communication areas (e.g. proxy transfers for grid access).
 - Distributed, low-cost (dCache)
 - Large volume data storage for bulk data processing.
- Control and allocation policy implementation:
 - FLEXIBLE implementations of: AccessCtrl, Quotas, Leases, CleanupPolicies
- Interfaces: SRM, Gridftp, ...?
- Data Management:
 - Catalogs, DQ2 site services, ...?
- Information/Monitoring

Tier2 Functions: Access Control

- Provide authentication and authorization:
 - Multi CA authentication
 - Service specific authorization
 - Roles
- Provide means to reconstruct an audit trail:
 - Log session history and archive tracking/accountMapping information.
 - Account pool model preferable to the more simple/convenient group accounts scheme.

Tier2 Functions: Monitoring and Accounting

- Provide info regarding:
 - Systems status, alarms, events.
 - Utilization, loads and historical info.
- Centralize status information collection if possible, to eventually leverage monitoring staff from other centers (USATLAS or OSG GOCs).
- Archive utilization info to assess efficiency, detect problems and prepare reports.

Tier2 Functions: Tech Support

- Provide communication channels and procedures for tech support specific to the site.
- Since there will be no user support per se, the effort is concentrated at supporting T2 systems and batch production.

Capacity Profile

- Where are we relative to these targets?
- Do we achieve optimal utilization?
- We've under-provisioned storage!
 - Can we adjust by increasing network utilization?
 - Will the Tier3s change this conclusion?

US ATLAS Tier 2 Centers (Federations)

4-Jan-06

| Tier 2 Center | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 |
|------------------------------|------|-------|-------|-------|--------|------|
| Boston/Harvard | | | | | | |
| <i>CPU (kS2k)</i> | - | 350 | 730 | 1,090 | 1,600 | |
| <i>Disk (TB)</i> | - | 170 | 370 | 480 | 630 | |
| Southwest | | | | | | |
| <i>CPU (kS2k)</i> | 500 | 900 | 1,500 | 1,700 | 2,100 | |
| <i>Disk (TB)</i> | 60 | 200 | 380 | 540 | 700 | |
| Midwest | | | | | | |
| <i>CPU (kS2k)</i> | 360 | 510 | 900 | 1,100 | 1,300 | |
| <i>Disk (TB)</i> | 50 | 130 | 260 | 465 | 790 | |
| Western | | | | | | |
| <i>CPU (kS2k)</i> | | 880 | 1772 | 3204 | 5269 | 8437 |
| <i>Disk (TB)</i> | | 64 | 244 | 498 | 858 | 1430 |
| AGL | | | | | | |
| <i>CPU (kS2k)</i> | | 387 | 726 | 1152 | 1850 | 2774 |
| <i>Disk (TB)</i> | | 156 | 304 | 471 | 770 | 1055 |
| Total for 5 Tier 2's | | | | | | |
| <i>CPU (kS2k)</i> | - | 3,027 | 5,628 | 8,246 | 12,119 | |
| <i>Disk (TB)</i> | - | 720 | 1,558 | 2,454 | 3,748 | |
| Target for 5 Tier 2's | | | | | | |
| <i>CPU (kS2k)</i> | - | - | 2,255 | 6,514 | 9,864 | |
| <i>Disk (TB)</i> | - | - | 928 | 2,919 | 4,908 | |

Tier2 Services: CPU and Storage access

- OSG Gatekeeper(s):
 - Recognizing multiple (OSG) CAs.
 - Role based authentication support (Prima)
 - Logging session and mappings (Audit).
- OSG Gridftp:
 - Use similar user/group mappings with the above. (if possible currently)
- SRM:
 - Which implementation?
 - Similar user/group mapping, again.

Tier2 Services – CPU Allocation

- Queuing system with:
 - Restrictions based on #jobs, SumCPU/Wall.
 - Base priority function of group/user
 - Preemption (suspension/requeuing)
 - Service Agreements or Hard/Soft Allocations
- Provides:
 - Guaranteed allocation to production and DA.
 - Low priority access to regular users.
 - Preemption
 - Fair-share policy using utilHistory + basePriority

Tier2 Services: Storage and Data Management

- Provide “global” FS (NFS, IBRIX,...?) and low-cost mass storage (dCache, ...?).
- (Are we on schedule with capacity delivery?)
- Implement SRM and Gridftp on top.
- Provide: access ctrl (file protection), (user/group) quotas, leases, automatic maintenance. Depending of storage implementation, some features are not available.
- Implement Data Management specific services:
 - DDM site services, LRC, ...?
- Monitoring status and usage, and publish info. Create the necessary tools.
- Accounting... write or adopt tools...

Tier2 Services: Authentication and Authorization

- Deploy GUMS with role based authentication.
- Recognize all OSG VOs and import from multiple VOMS.
- Preferably run dynamic account mapping (pool of accounts) for audit purposes.
- If not, implement at least the 4 USATLAS accounts scheme (debug, production, regular user, ATLAS) and separate the other VOs in individual accounts.
- Log and archive mappings and sessions (w/o details).
- Followup the gPlazma developments.

Tier2 Services: Monitoring and Accounting

- Deploy Ganglia and eventually push info centrally.
- Consider using Nagios (shall we standardize?) and centralize info collection.
- MonAlisa? Others?
- Followup the accounting developments in LCG and OSG
- Run whatever accounting you can, in the meantime.

Tier2 Services: Support

- Define responsibility chain and call up procedures.
- Appoint contacts (and create comm channels)
- Become accustomed with RT for ticket processing.
- Write documentation and procedures (😊)

Tier2: What did I forget?

- ...?