



Recent Experience with the PanDA Workload Management Sysem

Torre Wenaus Brookhaven National Laboratory, USA

OSG Users Meeting Brookhaven National Laboratory June 16, 2008





Panda Overview



Workload management system for Production ANd Distributed Analysis Core team @ BNL, UT Arlington/CERN. 4 developers, 1 full time



- Launched 8/05 by US ATLAS to achieve scalable data-driven WMS
- Individual analysis as well as managed production
- Insulate users from distributed computing complexity
 - Lower entry threshold
- US ATLAS production since '05
- ATLAS-wide production since early '08
- OSG WMS program since 9/06
 - VO-neutrality
 - Condor integration
 - Security enhancements





PanDA Attributes



- Pilots for 'just in time' workload management
 - Efficient 'CPU harvesting' prior to workload job release
 - Insulation from grid submission latencies, failure modes and inhomogeneities
- Tight integration with data management and data flow
 - Designed/developed in concert with the ATLAS DDM system
- Highly automated, extensive monitoring, low ops manpower
- Based on well proven, highly scalable, robust web technologies
- Can use any job submission service (**CondorG**, local batch, EGEE, Condor glide-ins, ...) to deliver pilots
- Global central job queue and management
- Fast, fully controllable brokerage from the job queue
 - Based on data locality, resource availability, priority, quotas, ...
- Supports multiple system instances for regional partitioning





PanDA Pilots



- Prescheduled to worker nodes via grid (CondorG, gLite) or batch submission
- Upon launch, they validate environment and request a job from PanDA; exit immediately if no job
- Rate is dynamically regulated according to workload
- Payload job is monitored in parallel to execution; status reported periodically to PanDA server
- Job recovery capability
- Data handling for the local storage element
- Pilot+scheduler is the *principal insulation layer* for the rest of the system, and its users, from grid/site details
 - Encapsulates heterogeneity of the grid(s)





Security in PanDA



- Uses grid certificate based security for the PanDA server and its client communications (https)
- OSG workload management effort at BNL in collaboration with BNL Tier 1 and OSG very active in security extensions
- Currently integrating glExec to switch pilot identity to that of the downloaded user job
 - Prototype implemented and working; integration in production underway
 - Fermilab acting as US/OSG test site; finding an EGEE site is an ongoing challenge!
- Implementing limited proxies (pilot role) to restrict pilot capabilities
 - Prevent users from gaining access to extensive rights via the pilot proxy
 - Pilot proxy also concealed from the application code (removed from disk)
- Client<->Server validation, payload validation to come
 - Drawing on CMS/FNAL work
 - Will prevent hijacking PanDA pilots to run unauthorized jobs
- Authorizing pilot submitters and pilots themselves (beyond simply holding valid certificate) with time-limited token also to come







- 2006 2007: PanDA responsible for US ATLAS production and analysis, performed extremely well
 - ~2500 CPUs, ~35% of total ATLAS production, well in excess of US share of ~20% through efficient resource use
 - Many hundreds of different physics processes simulated
- 2008: PanDA takes over ATLAS-wide production
 - Roll-out has been smooth, no scaling issues seen
 - But scaling studies ongoing to keep ahead of the growth curve
 - Expansion includes a second full PanDA instance at CERN
 - For 'sociological' more than technical (scaling) reasons
- >500 analysis users, ~200 with substantial usage, 2.5M jobs in last 6 months
- 1 shifter, ~.1 FTE to manage PanDA ops itself



- PanDA deployment across ATLAS began fall '07
- Today PanDA production spans all 10 ATLAS regions
 - almost 200 sites/queues
 - ~16-18k concurrent jobs (when enough jobs are available!)

Torre Wenaus, BNL





- BNL has offered a production PanDA capability for OSG for the last ~1.5 years
 - In use by CHARMM (protein folding app) users for production for last ~1yr; yielded a publication
 - CHARMM effort ongoing, particularly in extending to MPI
- We stand ready to help and support new applications/ VOs interested to try out PanDA!
- Roll-out of new Condor glidein flavor underway
 - Scheduler (schedd) glidein allowing CondorG to deliver a persistent pilot submitter, Pilot Factory, within the perimeter of a site, under external grid control
 - Avoid the observed (by ATLAS) scaling limitations of the gatekeeper handling individual pilots

PanDA for OSG - CHARMM

Open Science Grid CHARMM protein folding application. Other users welcome!







- Only EGEE sites currently are automatically discovered, configured and tracked by PanDA (via lcg-infosites)
- Currently extending this automation to OSG such that PanDA 'sees' all current OSG resources, properly configured
- Then extend further...
- PanDA's site information DB/cache can be extended with *any* info, then readily usable in site brokerage, VO/application matchmaking to sites, problem detection, ...
 - e.g. automatically maintain a CHARMM tag that includes the sites that pass validation tests today, support the application requirements of CHARMM, and support the CHARMM VO
 - CHARMM users merely need to send pilots to 'CHARMM'; PanDA handles the rest
 - No struggles by VO people to determine which sites work for them







- PanDA offers very extensive monitoring of the system at all levels
 - an important key to its ease of use and efficiency in terms of resource utilization and operations manpower
 - managed production jobs, individual analysis 'My Panda' pages, job diagnostics and logs, pilot and pilot scheduler monitoring, site attribute and status, error rates by type and site...
- All monitoring is available to OSG, including 'custom' VO-specific pages, but not well factored from ATLAS specifics
- Currently implementing a cleaner, OSG-user-friendly version of the monitoring interface

	BNL monitor	Production Clouds D	DM Pan	daMove	er AutoPilot	Sites A	nalvsis P	hysics data	Lisane I	Plote Prod	Dach DD	Dach				
					or ribior not		trangers 1	190100 0010	obage i	1013 11001	Jasii UDI	NDash				
17	3 min old Update															Not log
	Panda monitor Now in UTC	Panda Production Operations Dashboard														
	Shift log Wiki	Panda shift guide calendar mailing list ADCoS twiki calendar mailing list Mailing list for all questions related to ATLAS production: atlas-project-computing-operations@cern.ch														
	Jobs - search Recent running.	Servers: BNL:OK E	BNLdev:	OK CEF	RN:OK Logg	er:OK	Bamboo:	OK		Space	e available	e at sites:				
	activated, waiting,	Active tasks: CA:3 C	ERN:1	DE:6 ES	:2 FR:41 IT:4	NL:1 T	W: <u>8</u> UK:2	US:24		110	Sit	e	GB As	s of (UTC)		
	assigned, defined, finished, failed jobs	Bamboo submission	s, status	over las	at 12 hours					US	C2 conder		24115.06	16.07:52		
	Select analysis,	Jobs updated >12 h	rs ago:	activat	ted: <u>6741</u> run	nning:n	one			BULA	TI AS TI	20	73017 06	16 09:06		
	prod, install, test jobs	Jobs updated >36 h	rs ago:	transfe	erring:308					MWT	2 UC-pbs	3	22872 06	-16 09:34		
	Outob essent	Ganglia World Wide	Summa	iry						OU (OCHEP S	WT2-con	dor 7063 06	6-16 07:58		
	Job		World	Wide	- runnin	ig - d	lay		10	SLAC	CXRD-lsf		15969 06	6-16 07:59		
	Dataset	12 k							1001	SWT	2 CPB-pb	<u>is</u>	19044 06	6-16 09:33		
		11 k							/ 10	UC /	ATLAS_M	NT2	22876 06	6-16 09:29		
	Task request	10 k				-	-		0 10	UTA	SWT2		6572 06	6-16 09:33		
	Tack status	9 k						And I	NI TS	Other	r SEs repo	orting in la	st 3 days			
	lask status	8 k		-	1		-		22	FZK-	LCG2 MC	DISK	44673 06	5-16 09:13		
	File	10 7 K		_	-/					IN2P3	3-CC_MC	DISK	20328 06	6-16 09:13		
	Cummoriae	qo 6 k			~	-				NDG	E-T1 MCD	DISK	17502 06	-16 08:14		
	Blocks:	5 k								NIKH	EF-ELPR	OD MCD	ISK 9781 06	-16 08:14		
	days	4 k		-						PIC	MCDISK	00 1100	34267 06	-16 08:14		
	Errors:	3 k						100		RAL-	LCG2 MC	DISK	3676 06	6-16 08:14		
	days	2 k		_					8	TAIW	AN-LCG2	MCDISK	38198 06	6-16 08:15		
	Nodes:	1 K								TRIU	MF-LCG2	MCDISK	18371 06	6-16 09:13		
	Daily usage	0-														
	Tasks - search		DE	ES	FR U	00:00 K ■ T	rw 🗖 Ni	NDG								
	EvGen Task Reg	Generated by TRIUMF-LCG2 (times in UTC)														
	CTBsim Task Reg Task list															
	New Tag	Production job summ	mary, las	st 12 ho	urs (Details:	errors, r	nodes)									
	Bug Report Task browser	Cloud Information	Nodes	Jobs	Latest	Pilots (3hrs)	defined	assigned	waiting	activated	running	holding	transferring	finished	failed	tot trf other
	Datasets - search Dataset browser Aborted MC	Overall Production	6118	56289	06-16 09:40	7976	<u>0/0</u>	<u>494</u> / <u>0</u>	<u>300</u> /0	<u>12991</u> / <u>0</u>	<u>8706</u> / <u>0</u>	<u>2113/0</u>	<u>8837</u> / <u>308</u>	<u>11290 / 0</u>	<u>11861</u> / <u>0</u>	51% 1% 51%

<u>0</u>

0

<u>71</u>

57

<u>0</u>

0

3159

0

1473

31

<u>20</u>

1

<u>683</u>/0

0/0

2092

7

<u>110</u> 5% 0% 5%

19

73% 0% 73%

561 7608 06-16 09:40 155

66 115 06-16 09:39 478

datasets Panda subscriptions

Datasets

CA





- MySQL performance, load sharing, failover, partitioning
- Apache performance, tuning, redundancy, multi-host
- PanDA service tuning in light of stress tests
- Security -- as discussed
- Code organization, configuration management, packaging
 - Housecleaning after 2.5 years! Maintainability, ease of component and service installation, easier pickup and usage by other VOs
- OSG WMS
 - PanDA usage support (CHARMM, CHARMM/MPI, ...)
 - Validation of Condor schedd glide-in based Pilot Factory
 - OSG PanDA WMS service on BNL facility infrastructure
 - Other PanDA services for OSG: OSG- and VO-level site information/ validation; resource usage tracking/management; ...



Summary



- PanDA has delivered its objectives well
 - Efficient resource utilization (pilots and data management)
 - Ease of use and operation (automation, monitoring, global queue/broker)
 - Excellent scaling behavior
 - Makes its users happy, in ATLAS and OSG
- ATLAS-wide deployment has been effort-intensive but smooth and successful
 - Only NorduGrid still incomplete; hybrid approach taking longer
- Data management/handling always the greatest challenge
 - PanDA now has its own data handling mechanism, could be used by OSG for data-driven workflow
- Leveraging OSG effort and expertise, especially in security
- On track to provide stable and robust service for LHC physics datataking
 - Ready in terms of functionality, performance and scaling behavior
- Ready in terms of capability & available manpower to expand OSG applications
 - As a proven, hardened, user-friendly tool for utilizing the OSG and/or implementing OSG-wide services

Torre Wenaus, BNL





More Information

- PanDA wiki
 - <u>http://twiki.cern.ch/twiki//bin/view/Atlas/PanDA</u>
- PanDA monitor
 - <u>http://panda.atlascomp.org</u>
- The US ATLAS supported PanDA Team (1-2 pending hires)
 - Kaushik De, project co-lead (UT Arlington)
 - Tadashi Maeno, lead developer (BNL)
 - Paul Nilsson, developer (UT Arlington/CERN)
 - Aaron Thor, developer (UT Arlington/BNL)
 - Torre Wenaus, project co-lead/developer (BNL)
- BNL Facility PanDA Support
 - John Hover, Yuri Smirnov, Dantong Yu et al.
- BNL OSG WMS Team
 - Jose Caballero, Barnett Chiu (UTA/BNL), Maxim Potekhin (lead)



Supplementary Slides



Torre Wenaus, BNL







Production vs. Analysis

- Use the same PanDA and underlying grid infrastructure
 - Same software, monitoring system, facilities
 - No duplicated operations manpower
 - Analysis benefits from problems found/resolved by shifters
- Separate computing resources
 - Different queues mapping to different clusters
 - Relative usages can be controlled without competition
- Different policies for data management
 - Data delivery to sites is in production workflow, but not in analysis
 - Analysis jobs go to where the data currently resides
 - User can specify the site, or let the system decide
 - Analysis outputs reside on the local SE until retrieved by user
- PanDA user-level accounting/quota system supports analysis; being extended to production (at physics working group and project levels)
 - Priorities also supported; important basis for brokerage decisions

