

The Open Science Grid and High Energy Physics

Burt Holzman (Fermilab/CMS)
On behalf of the Open Science Grid

An aside: who am I?

- “Who are you and how did you get in here?”
 - “I’m a locksmith. And I’m a locksmith.”
- Not strictly from HEP -- dissertation and Postdoc in nuclear physics (heavy ions) at RHIC / BNL
- Currently working for CMS Grid Services @ Fermilab and manage several projects contributing to the Open Science Grid



What is the Grid?

“The Grid is an emerging infrastructure that will fundamentally change the way we think about – and use – computing. The word *Grid* is used by analogy with the electric power grid, which provides pervasive access to electricity”

[I. Foster & C. Kesselman, 1998]



What is OSG?

“Transform processing and data intensive science through a cross-domain self-managed national distributed cyber-infrastructure that brings together campus and community infrastructure and facilitating the needs of Virtual Organizations (VO) at all scales.”

[M. Livny, 2007]




What is OSG, really?

- More than 100 heterogeneous clusters of Linux machines
- Handful of different storage solutions with a common protocol
- Common interface to computing and storage
- Opportunistic use
- Interoperability with other grids




Birth of the OSG



OSG “born” about 2005



Birth of the OSG



OSG “born” about 2005



The Grid as Baby




Input



Output

Congratulations! It's a Grid!



- Babies are born knowing how to do 3 things: eat, sleep, and cry.
- OSG was born with
 - Globus
 - X.509 authentication
 - Condor
 - Virtual Organizations
- Allows management of collaboration users world-wide
- Distribution of global resources

OSG: Year 1

Crawling and Babbling



OSG: Year 1



- VDT packaging
 - Built on top of NMI
 - One stop shop for grid software

- LHC Tier 3s

- Accounting (“Gratia”)



- Is HEP getting what we pay for?

OSG: Year 2

Walking and Talking



OSG: Year 2



- Talking with others:
OSG-EGEE
Interoperability
- Mature information system, Resource Selection Service
- Wider reach of resources (WLCG)
- Matchmaking



OSG: Year 3

Complex Thoughts



Mine MINE MINE



OSG: Year 3



- Storage: dCache, BeSTMan + various backends (including Hadoop)
- Pilot-based Workload Management
- Well-defined storage (and publication of it!) enhances opportunistic computing
- Scalability, robustness, late-binding validation



What is OSG, really?


- More than 100 heterogeneous clusters of Linux machines
- Handful of different storage solutions with a common protocol
- **Common interface to computing and storage**
- **Opportunistic use**
- **Interoperability with other grids**



“Common Interfaces”




- CDF transitioned from distributed CDF Analysis Facilities (“CAF”) to the Grid
- Several grid access points (i.e. PacCAF, CNAFCAF, GlideCAF @ FNAL)






CDF on the Grid




GlideCAF at Fermilab




Asian GlideCAF - PacCAF

[D. Lucchesi CHEP09]



Opportunity




- DØ committed to opportunistic use of OSG
- Concentrated effort on both computing and storage



DØ on non-DØ resources

Production Last Year
By Segment




[J. Snow CHEP09]



Opportunity strikes again: **BooNE** BOOSTER NEUTRINO EXPERIMENT

- The MiniBooNE experiment is a 12m diameter sphere filled with 800 tons of mineral oil and instrumented with 1520 phototubes
- Goal: search the range of the LSND oscillation signal




- Computing challenge: reduce systematic errors from optical model



“It took us 6 months on OSG instead of 560 days” [C.Polly]

Cumulative Hours Spent on Jobs By Facility 1300 Days from Week 01 of 2006 to Week 29 of 2009




- | | | |
|------------------------------------|-------------------------|---------------------------|
| USCMS-FNAL-WC1-CE (1,214,563) | FNAL_CDFOSG_2 (465,898) | FNAL_GPFARM (465,722) |
| Fermilab General Purpose (252,601) | FNAL_GPGRID_1 (159,333) | FNAL_DZEROOSG_2 (128,613) |
| FNAL_DZEROOSG_1 (106,921) | FNAL_CDFOSG_1 (59,478) | FNAL_CDFOG_3 (49,931) |
| FNAL_GPGRID_3 (29,025) | FNAL_CDFOSG_4 (28,857) | Cdfosg1 (1,825) |
| Fermigrid1 (984.35) | FNAL_GPGRID_2 (831.80) | Cdfosg2 (603.57) |
| Test (0.00) | SDSS_TAM (0.00) | |

Total: 2,965,192 Hours, Average Rate: 0.03 Hours/s



Opportunity keeps knocking

Opportunistic CPU Hours by Site
197 Days from Week 00 of 2009 to Week 28 of 2009



- | | | | | |
|---------------------|-------------------|-----------------|--------------------|-----------------|
| ■ USCMS-FNAL-WC1-CE | ■ Other | ■ Nebraska | ■ FNAL_GPGRID_1 | ■ FNAL_GPGRID_3 |
| ■ MIT_CMS | ■ FNAL_FERMIGRID | ■ MSU-OSG | ■ FNAL_GPGRID_2 | ■ UFlorida-HPC |
| ■ OU_OSCER_ATLAS | ■ OU_OSCER_CONDOR | ■ OU_OCHEP_SWT2 | ■ Purdue-RCAC | ■ UFlorida-PG |
| ■ Purdue-Steele | ■ MWT2_IU | ■ UCSDT2 | ■ RENCI-Engagement | ■ SPRACE |


Maximum: 234,043 , Minimum: 43,495 , Average: 114,310 , Current: 43,495



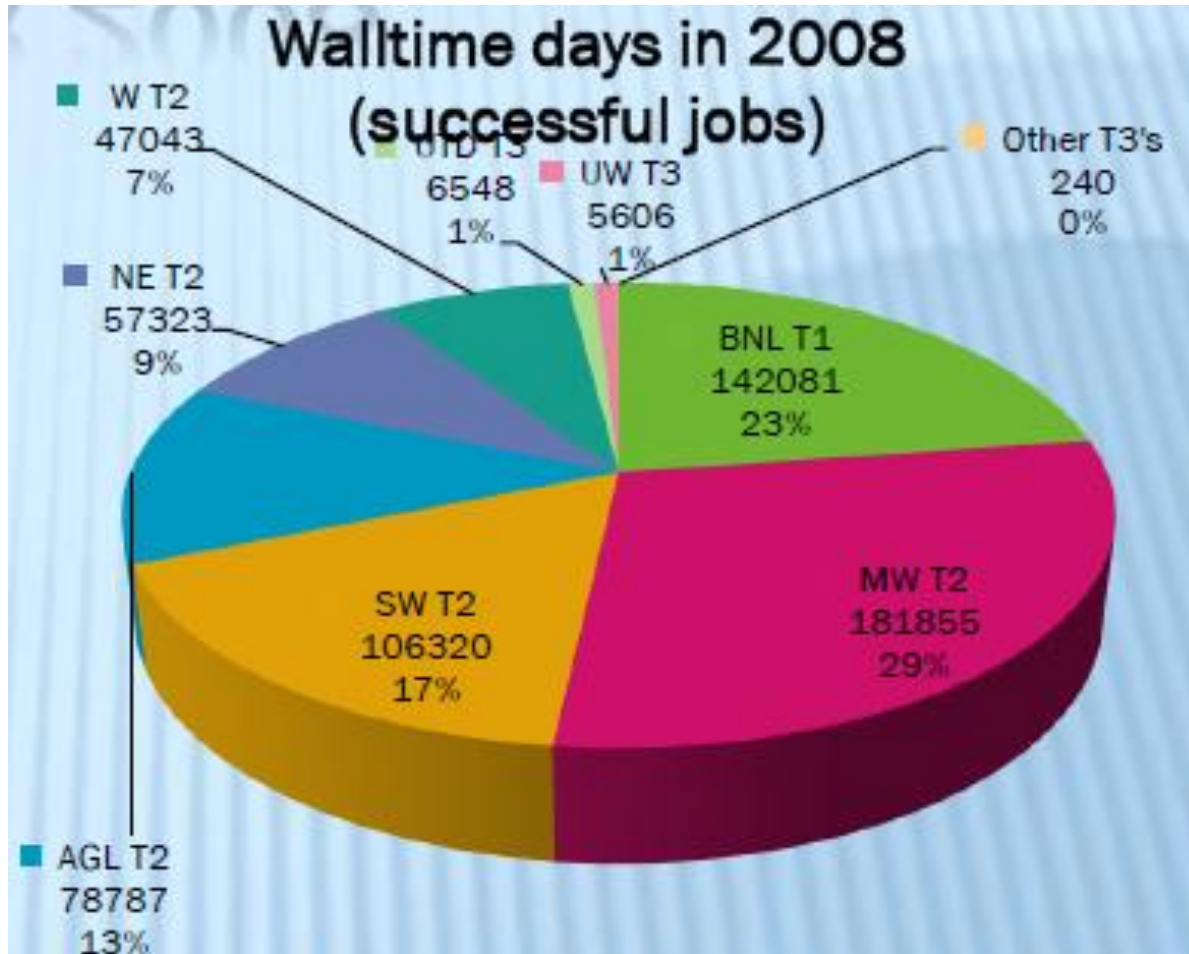
Success with Pilots



- ATLAS's PanDA software turns OSG into big virtual pool!
- Abstracts the entire grid for ATLAS and analysis use
- Job slots are pre-validated before jobs begin



ATLAS – PanDA in 08




Total of 15 million hours in 2008

[P.Nilsson ACAT08]



ATLAS: OSG totals

Cumulative Hours Spent on Jobs By Facility
1302 Days from Week 01 of 2006 to Week 30 of 2009



BNL_ATLAS_1 (22,042,771)	MWT2_UC (8,661,856)	AGLT2 (6,944,888)	WT2 (5,900,751)
MWT2_IU (5,865,460)	UTA_SWT2 (4,691,966)	BU_ATLAS_Tier2 (4,689,004)	SWT2_CPB (3,809,727)
OU_OCHEP_SWT2 (2,276,272)	UTA_DPCC (1,018,287)	IU_OSG (968,224)	UC_ATLAS_MWT2 (817,824)
OU_OSCER_ATLAS (722,537)	Other (328,345)	Nebraska (256,978)	IllinoisHEP (200,440)
GLOW-ATLAS (187,581)	UC_Teraport (137,745)	HU_ATLAS_Tier2 (111,135)	MCGILL_HEP (98,135)


Total: 69,729,937 Hours, Average Rate: 0.62 Hours/s




Interoperability



- Interoperability is crucial for the CMS experiment across OSG and EGEE
- OSG and EGEE information systems coherent
- “Which Grid” is transparent to end-user



CMS: analysis during STEP09




CRAB: CMS's analysis application
User picks dataset, application does the rest
(querying dataset bookkeeping service, information systems, handling jobs, etc.)

Not just HEP on OSG

Monthly Wall Hours per VO for non-HEP VOs

53 Weeks from Week 28 of 2008 to Week 29 of 2009




- engage
- glow
- nysgrid
- ligo
- aceace
- nanohub
- johns
- sbgrid
- osg
- icbr
- fgstore
- cigi
- chen
- mis
- jis
- Other
- cyt701
- plukashev
- wu
- osgedu

Maximum: 1,357,159 Hours, Minimum: 0.00 Hours, Average: 722,188 Hours, Current: 878,054 Hours



Over 330 million hours served

Cumulative Hours Spent on Jobs By Facility
1142 Days from Week 22 of 2006 to Week 28 of 2009



Other (67,094,922)	USCMS-FNAL-WC1-CE (58,776,529)	FNAL_CDFOSG_2 (28,215,154)	FNAL_DZEROOSG_2 (24,897,240)
BNL_ATLAS_1 (22,613,731)	Nebraska (18,158,247)	FNAL_CDFOSG_1 (16,884,245)	MIT_CMS (13,694,374)
GLOW (12,718,239)	FNAL_CDFOSG_3 (9,705,305)	Purdue-RCAC (8,778,795)	MWT2_UC (8,416,406)
UCSDT2 (8,267,199)	MWT2_IU (6,574,995)	CIT_CMS_T2 (6,488,964)	OU_OSCER_ATLAS (6,252,062)
WT2 (5,777,663)	OU_OSCER_CONDOR (4,877,122)	UFlorida-PG (4,866,720)	UTA_SWT2 (4,676,835)

Total: 337,734,758 Hours, Average Rate: 3.42 Hours/s



Conclusions

- Open Science Grid provides
 - Common interfaces
 - Opportunistic Computing
 - Interoperability with other Cyberinfrastructures
- What will Year 36 look like?



Regrets & Thanks

- Regrets to many, many OSG virtual organizations I left out (STAR, LIGO, and many many many more)
- Thanks to organizers of course!
- Thanks to Joel Snow, Doug Benjamin, Chris Polly

