

Open Science Grid

Frank Würthwein
UCSD



“Airplane view” of the OSG

- ⇒ High Throughput Computing
 - Opportunistic scavenging on cheap hardware.
 - Owner controlled policies.
- ⇒ “Linux rules”: mostly RHEL3 on Intel/AMD
- ⇒ Heterogeneous Middleware stack
 - Minimal site requirements & optional services
 - Production grid allows coexistence of multiple OSG releases.
- ⇒ “open consortium”
 - Stakeholder projects & OSG project to provide cohesion and sustainability.
- ⇒ Grid of sites
 - Compute & storage (mostly) on private Gb/s LANs.
 - Some sites with (multiple) 10Gb/s WAN uplink.



OSG by numbers

- ⇒ 53 Compute Elements
- ⇒ 9 Storage Elements
(8 SRM/dCache & 1 SRM/DRM)
- ⇒ 23 active Virtual Organizations

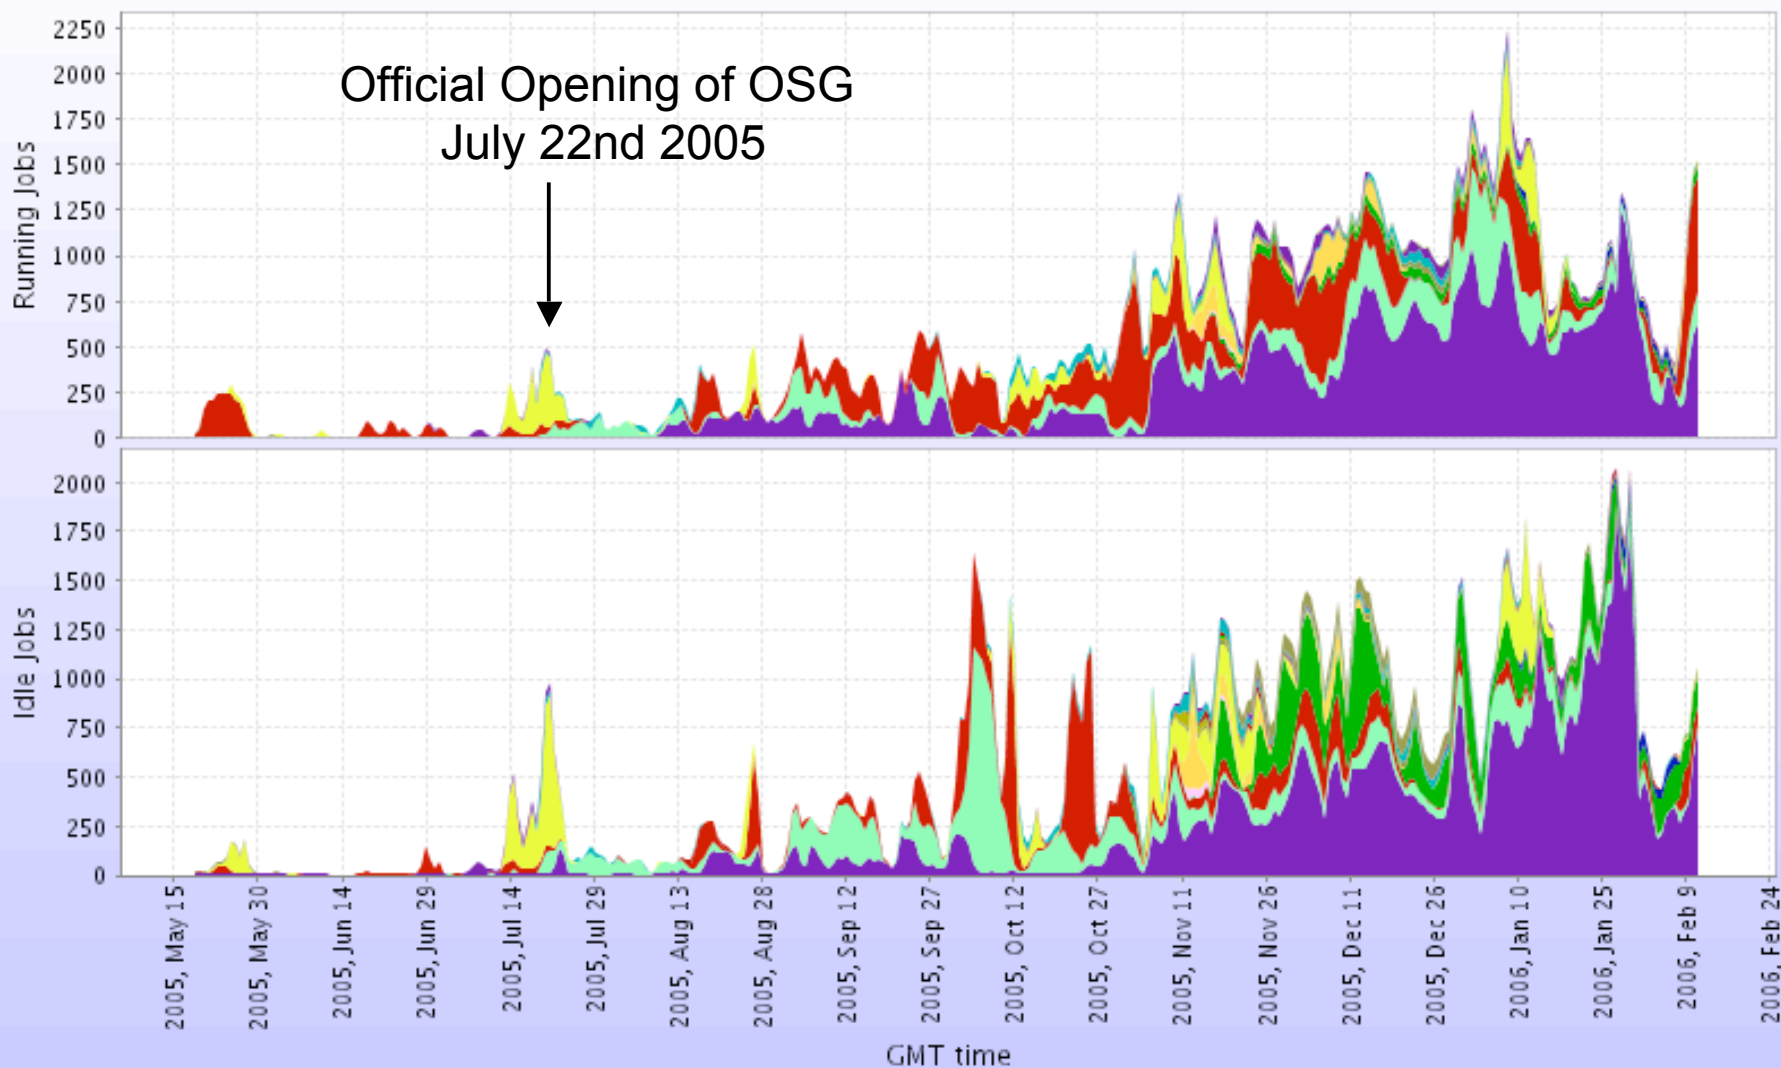
4 VOs with >750 jobs max.




4 VOs with 100-750 max.

Farm	Running Jobs			
	Last value	Min	Avg	Max
ATLAS	618.6	0	232	2282
AUGER	0	0	0	5
CDF	171.9	0	74.22	752
CDMS	0	0	0.623	40
CMS	636.7	0	133.6	980
DOSAR	0	0	0.004	10
DZERO	56.69	0	17.57	77
FERMILAB	0	0	3.709	186
FMRI	0	0	0.459	46
GADU	0	0	10.71	301
GRASE	0.155	0	34.52	1045
GRIDEX	10.66	0	6.39	38
IVDGL	19.9	0	2.517	83
LIGO	0	0	0.681	82
MINIBOONE	0	0	0.001	14
MINOS	0	0	0	15
MIS	0.041	0	0.437	34
NANOHUB	0	0	0.002	1
OSG	0	0	0.095	14
PATRIOT	1	0	0.057	2
SDSS	3.421	0	11.59	244
STAR	0	0	11.92	114
USATLAS	0	0	2.166	16
Total	1519			

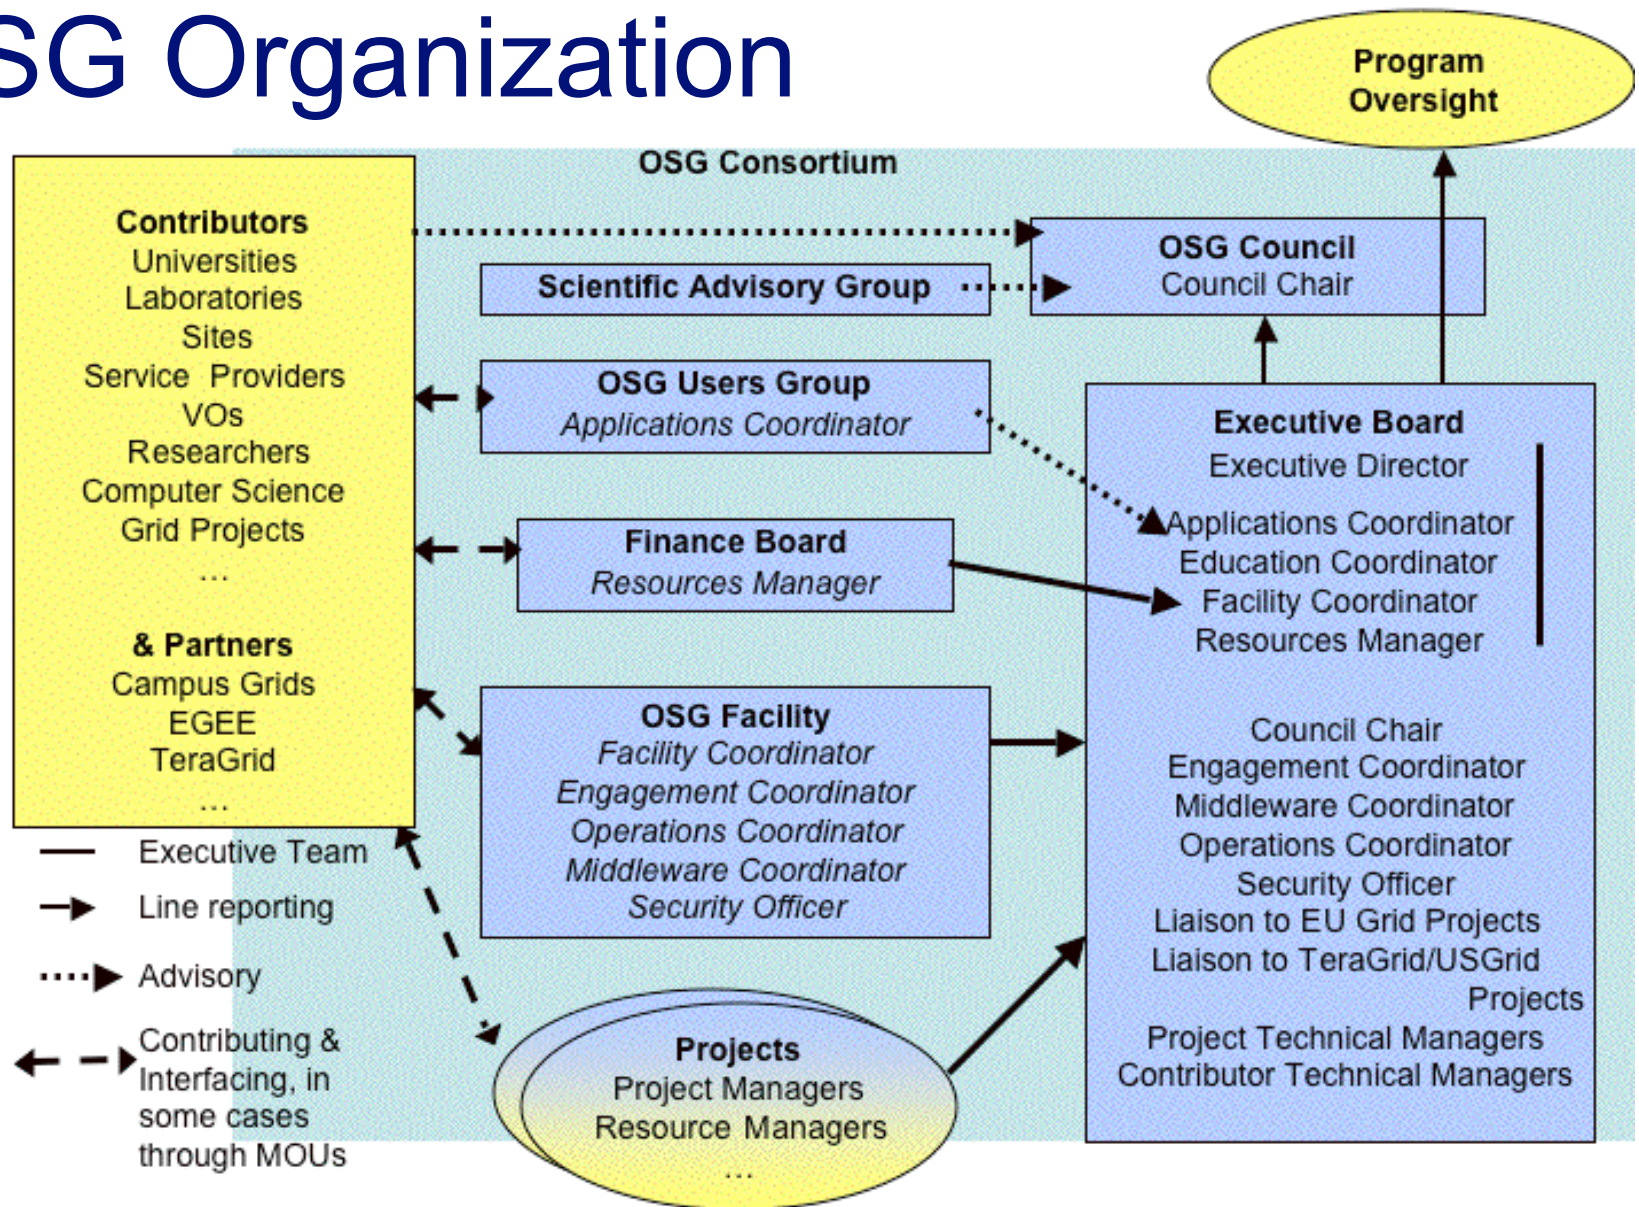


Total Jobs per VO





OSG Organization



OSG organization (explained)

- ⇒ OSG Consortium
 - Stakeholder organization with representative governance by OSG council.
- ⇒ OSG project
 - (To be) funded project to provide cohesion & sustainability
 - OSG Facility
 - “Keep the OSG running”
 - “Engagement of new communities”
 - OSG Applications Group
 - “keep existing user communities happy”
 - Work with middleware groups on extensions of software stack
 - Education & Outreach

OSG Management

Executive Director:	Ruth Pordes
Facility Coordinator:	Miron Livny
Application Coordinators:	Torre Wenaus & fkw
Resource Managers:	P. Avery & A. Lazzarini
Education Coordinator:	Mike Wilde
Council Chair:	Bill Kramer

The Grid “Scalability Challenge”

- ⇒ Minimize entry threshold for resource owners
 - Minimize software stack.
 - Minimize support load.
- ⇒ Minimize entry threshold for users
 - Feature rich software stack.
 - Excellent user support.

Resolve contradiction via “thick” Virtual Organization layer of services between users and the grid.



Me -- My friends -- The grid

Me: thin user layer

My friends:
VO services
VO infrastructure
VO admins

The Grid: anonymous sites & admins

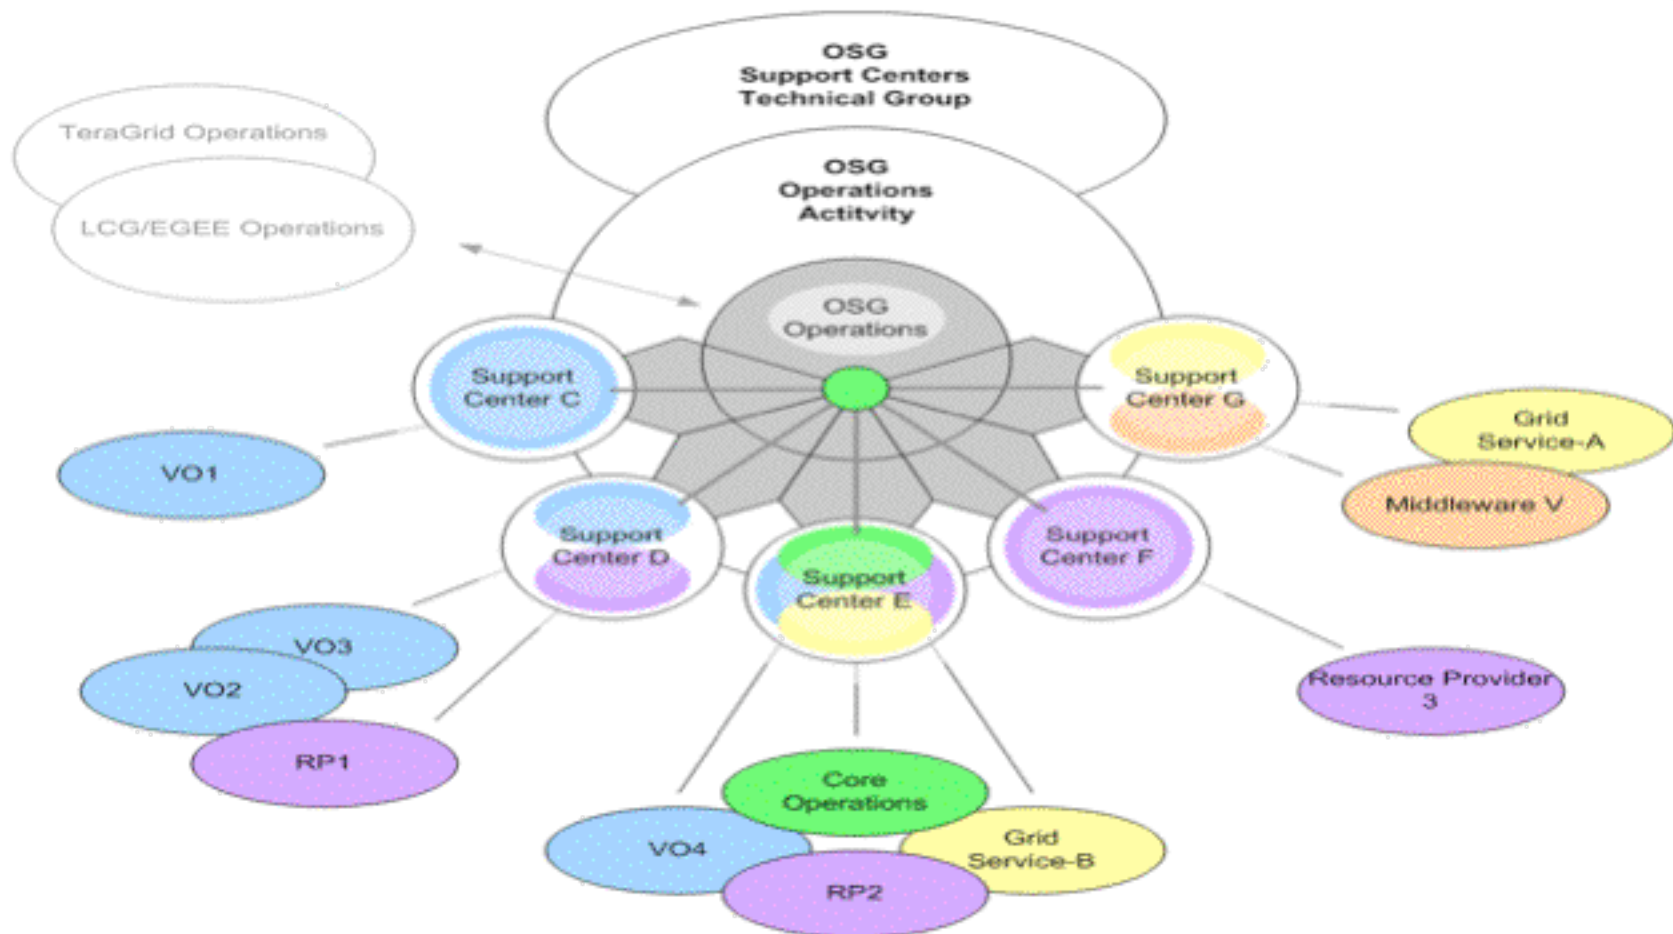
Me & My friends
are domain science
specific.

Common to all.




OSG Operations Model

Physical View

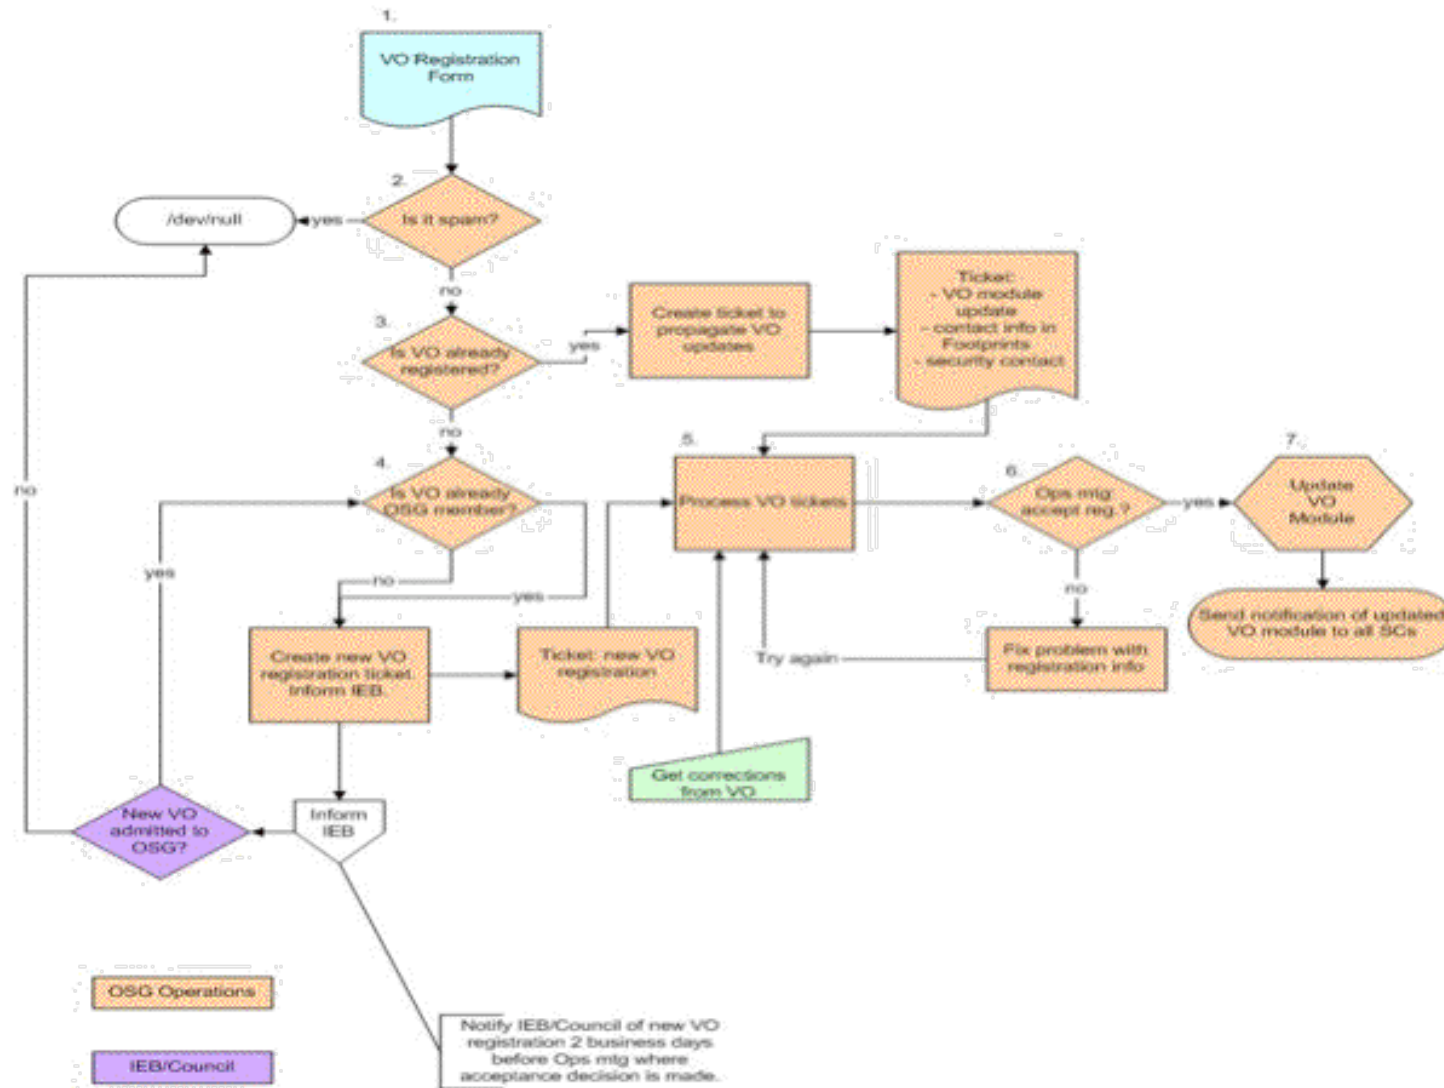



User Management

- ⇒ User registers with VO and is added to VOMS of VO.
 - VO responsible for registration of VO with OSG GOC.
 - VO responsible for users to sign AUP.
 - VO responsible for VOMS operations.
 - VOMS shared for ops on both EGEE & OSG by some VOs.
 - Default OSG VO exists for new communities.
- ⇒ Sites decide which VOs to support (striving for default admit)
 - Site populates GUMS from VOMSes of all VOs
 - Site chooses uid policy for each VO & role
 - Dynamic vs static vs group accounts
- ⇒ User uses whatever services the VO provides in support of users
 - VO may hide grid behind portal
- ⇒ Any and all support is responsibility of VO
 - Helping its users
 - Responding to complains from grid sites about its users.



Virtual Organization Registration



Compute & Storage Elements

- ⇒ Compute Element
 - GRAM to local batch system.
- ⇒ Storage Element
 - SRM interface to distributed storage system.
 - Continued legacy support: gsiftp to shared filesystem.

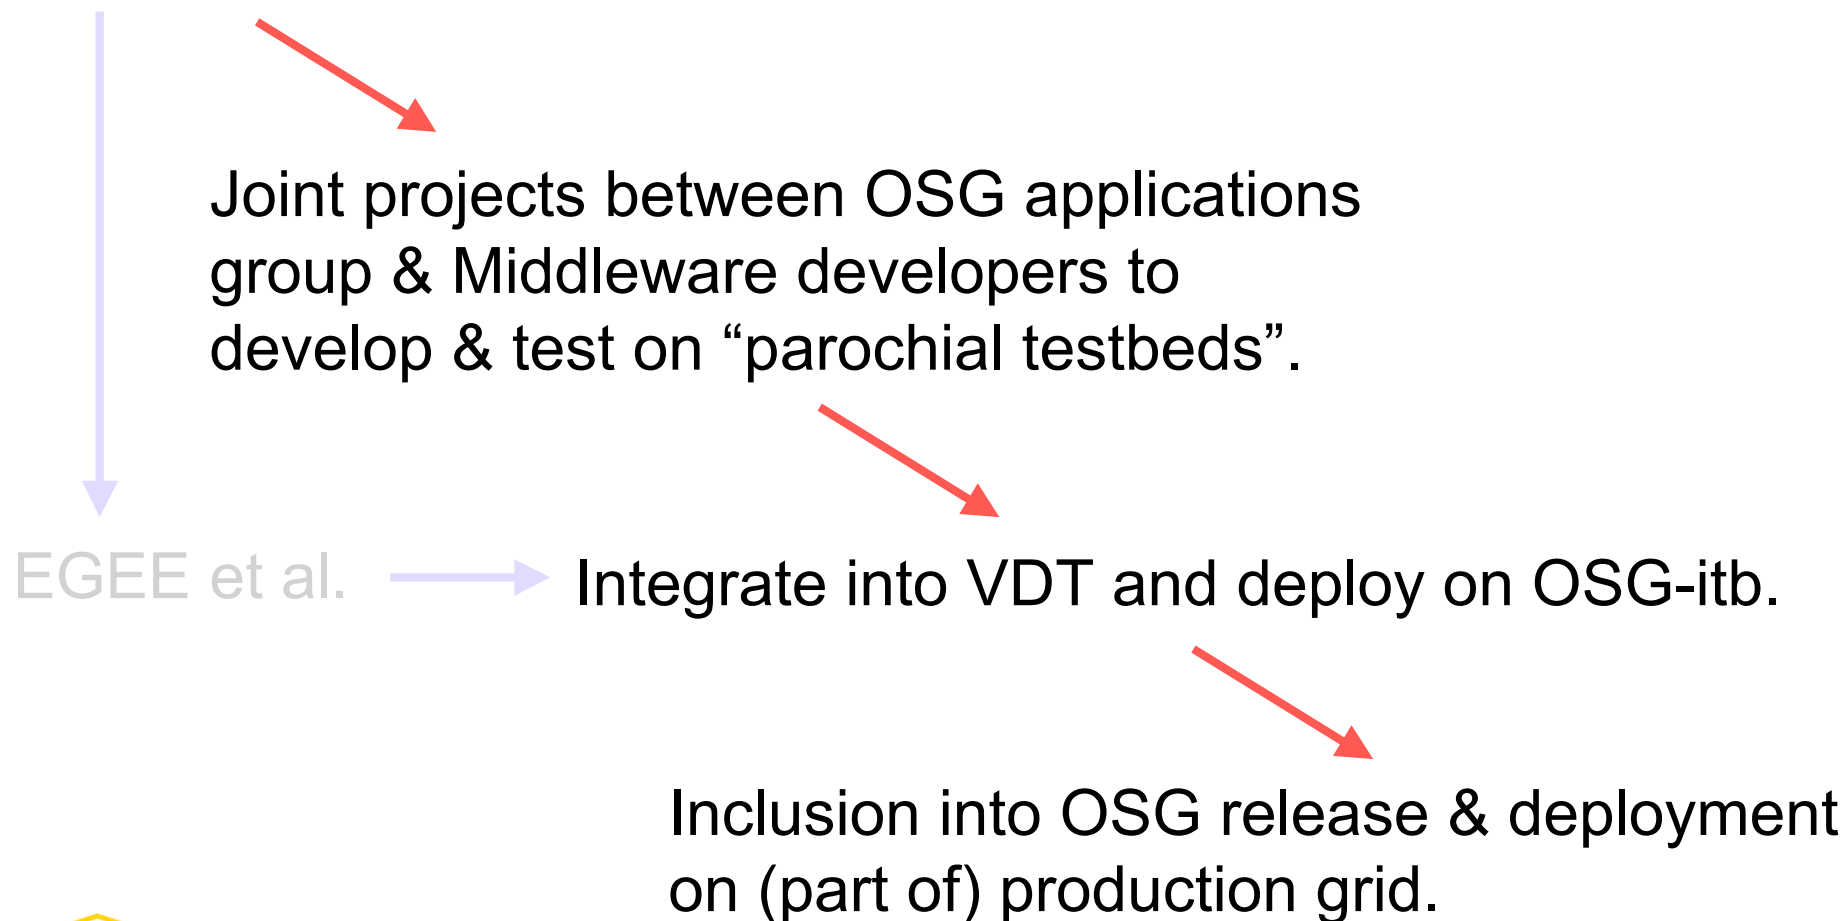


Disk areas in more detail:

- ⇒ Shared filesystem as applications area.
 - Read only from compute cluster.
 - Role based installation via GRAM.
- ⇒ Batch slot specific local work space.
 - No persistency beyond batch slot lease.
 - Not shared across batch slots.
 - Read & write access (of course).
- ⇒ SRM controlled data area.
 - Job related stage in/out.
 - “persistent” data store beyond job boundaries.
 - SRM v1.1 today.
 - SRM v2 expected in next major release (summer 2006).

Middleware lifecycle

Domain science requirements.



Challenges Today

⇒ Metrics & Policies

- How many resources are available?
- Which of these are available to me?

⇒ Reliability

- Understanding of failures.
- Recording of failure rates.
- Understanding relationship between failure and use.



Release Schedule

	Planned	Actual
OSG 0.2	Spring 2005	July 2005
OSG 0.4.0	December 2005	January 2006
OSG 0.4.1	April 2005	
OSG 0.6.0	July 2006	

Dates here mean “ready for deployment”.
Actual deployment schedules are chosen by each site,
resulting in heterogeneous grid at all times.

Summary

- ⇒ OSG facility is under steady use
 - ~20 VOs, ~1000-2000 jobs at all times
 - Mostly HEP but large Bio/Eng/Med occasionally
 - Moderate other physics (Astro/Nuclear)
- ⇒ OSG project
 - 5 year Proposal to DOE & NSF
 - Facility & Extensions & E&O
- ⇒ Aggressive release schedule for 2006
 - January 2006: 0.4.0
 - April 2006: 0.4.1
 - July 2006: 0.6.0