



EGEE09 Barcelona

ATLAS Distributed Data Management

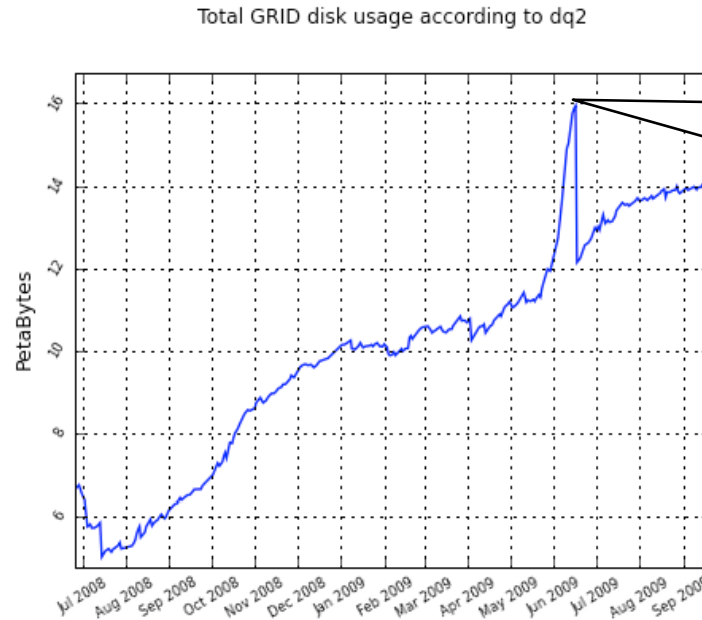
Fernando H. Barreiro Megino on behalf of the
ATLAS DDM team



1. Introduction
2. ATLAS DDM Architecture
3. ATLAS DDM and the users
4. Summary



- Presently



STEP09: Stress test involving all key elements from data-taking to analysis

- Expected during data taking

- ✧ Generation of RAW data

$$1.6 \text{ MB} \times 200 \text{ Hz} \times 50 \frac{\text{ksec}}{\text{day}} \times 100 \frac{\text{days}}{\text{year}} = 1.6 \frac{\text{PB}}{\text{year}}$$


- ✧ Processing and reprocessing

- ✧ Simulation production

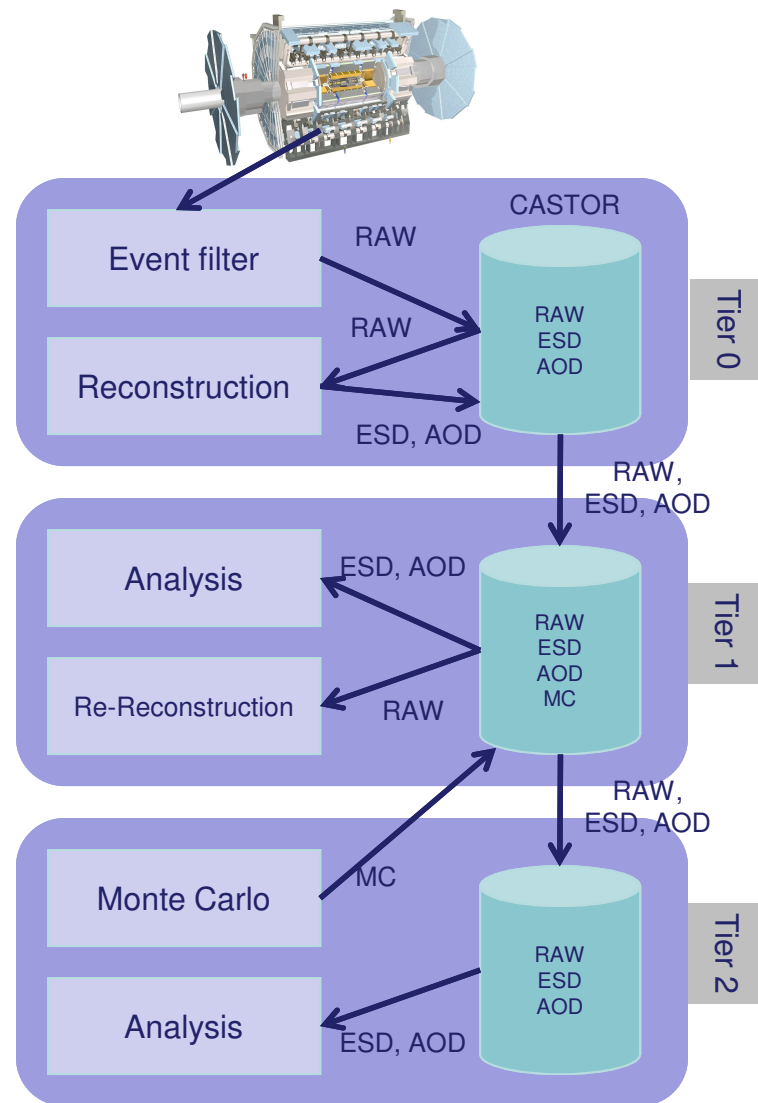
...

- Central link between WLCG and ATLAS analysis components
- Manage the experiment's data:
 - Data movement between associated sites
 - Bookkeeping & accounting
 - Data access to
 - Production systems
 - Physics meta-data systems
 - Analysis systems
 - End users

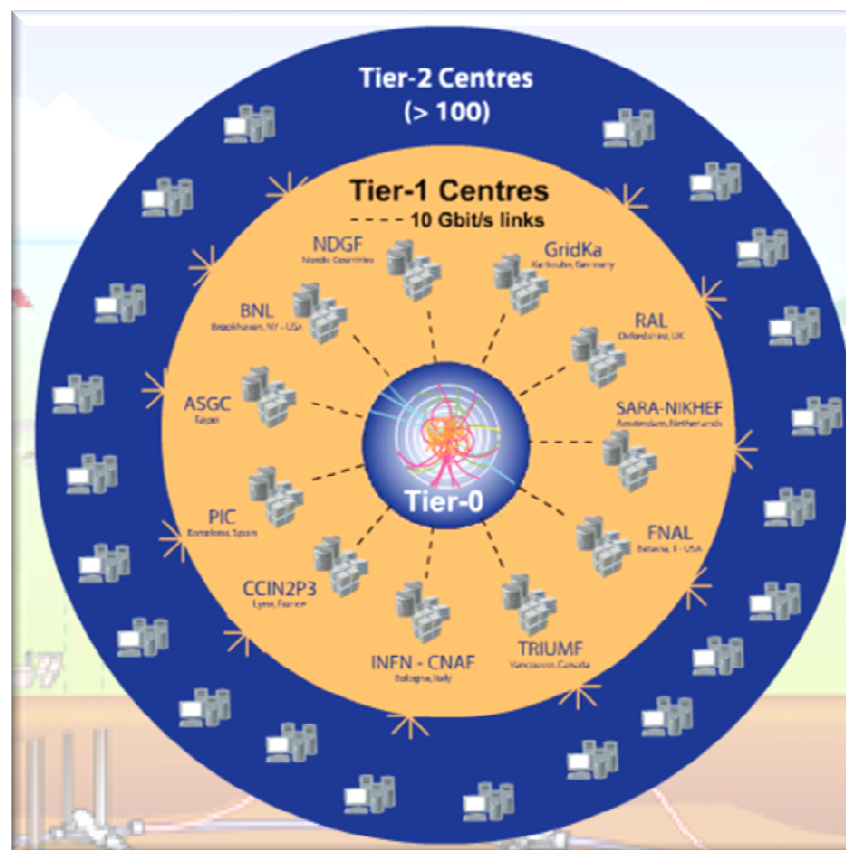
...following ATLAS' Computing model

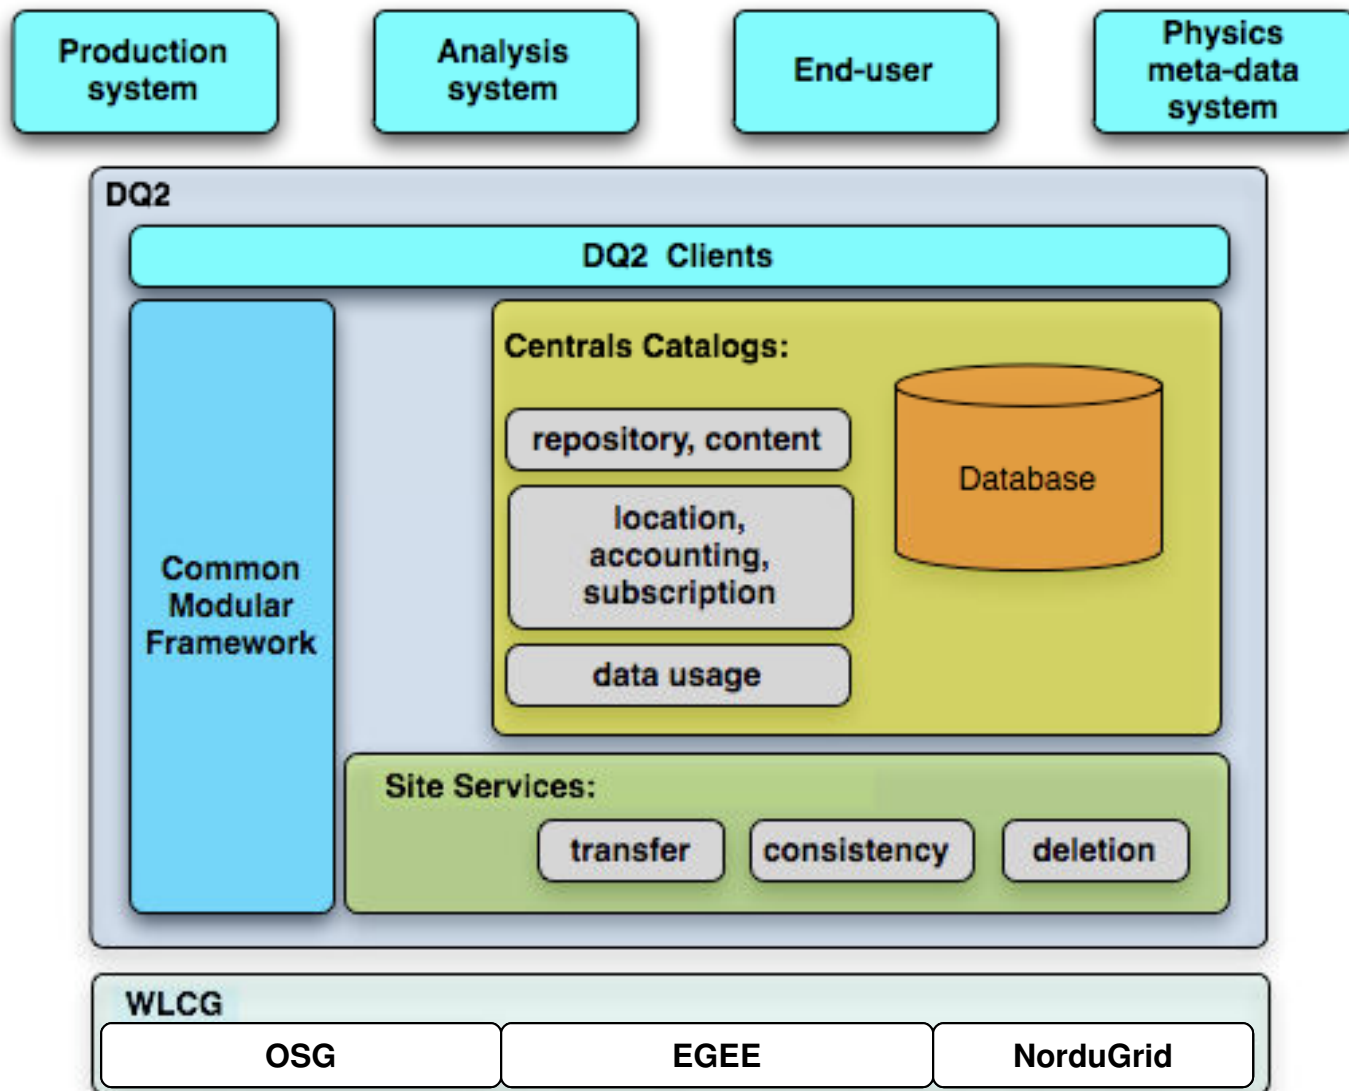


- Tier-0 facility (CERN):
 - Archival and distribution of primary RAW detector data
 - First pass processing of the primary event stream
 - Distribute the derived data to the Tier-1s
- 10 Tier-1 data centers:
 - Long-term access to all data
 - Reprocessing capacity
- ~100 Tier-2 institutes:
 - Analysis capacity for users & physics-groups
 - Monte-Carlo simulation

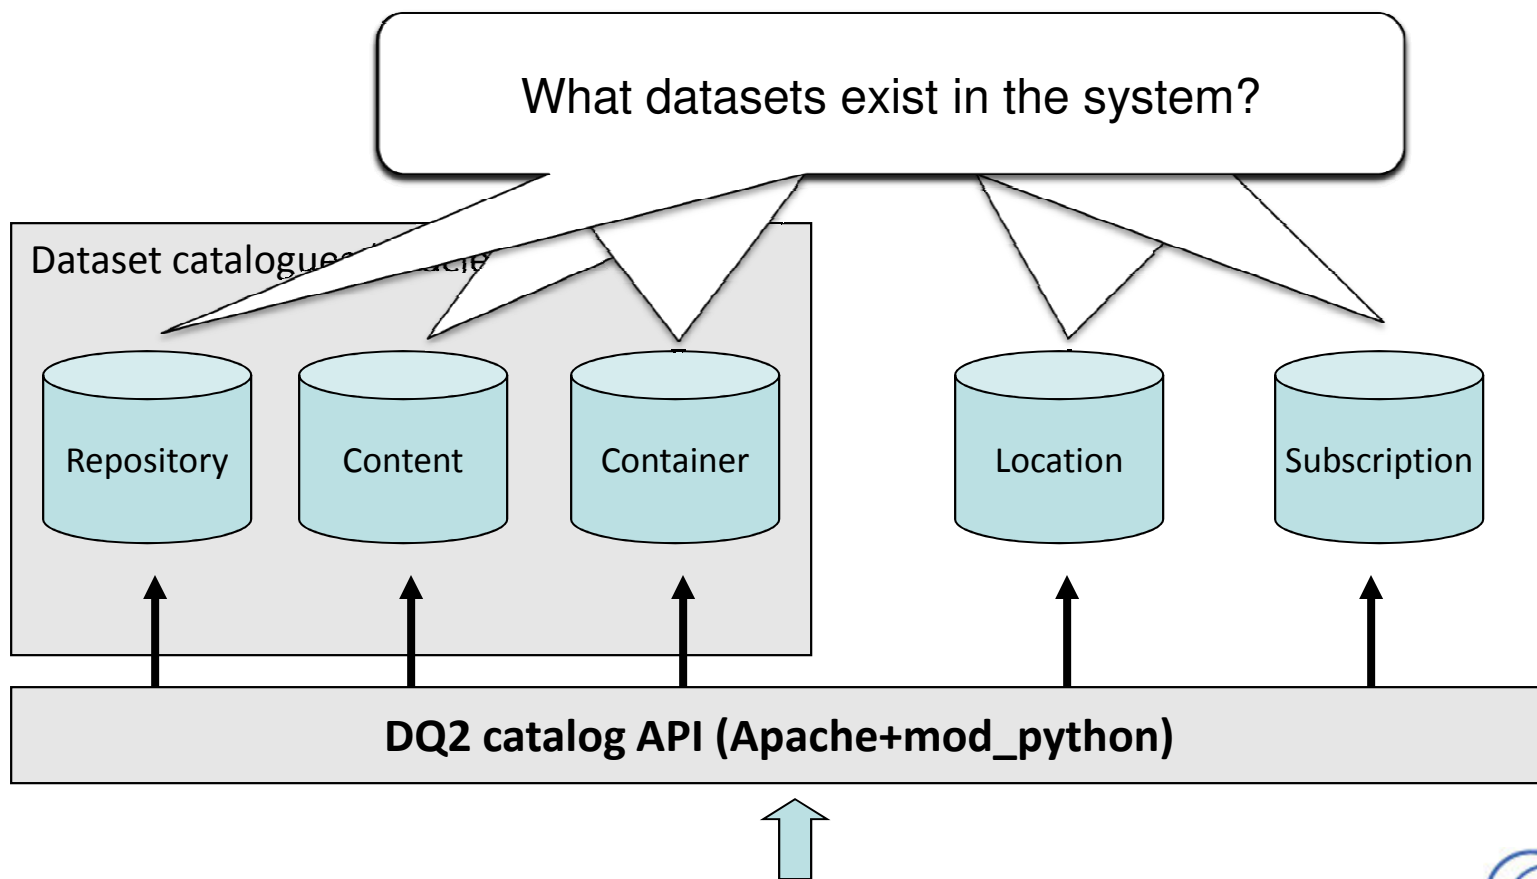


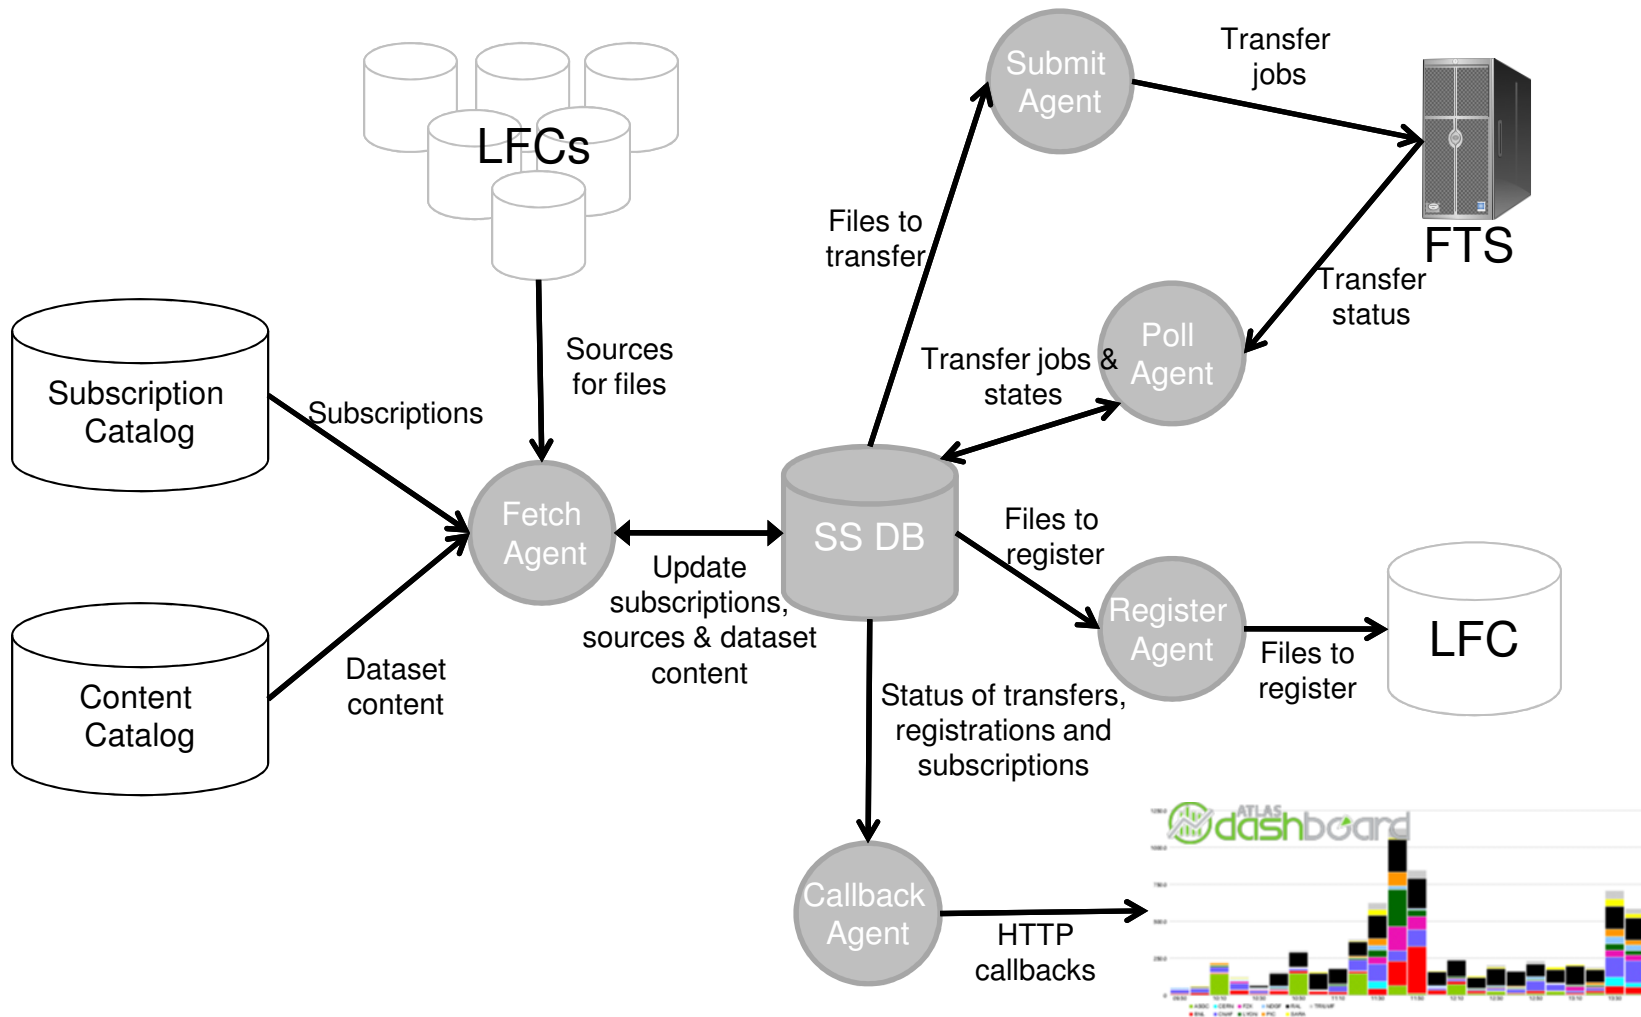
- Cloud model
- 3 different Grid flavors with different middleware:
 - Open Science Grid (OSG): US
 - NorduGrid (NDGF): Nordic European countries
 - EGEE: Europe and rest of the world






- Datasets: Collection of files acting as transfer and organization units
- Subscriptions: Replication requests of ATLAS datasets






- Monitoring for:
 - Shifters: 24/7 follow-up of DDM activity
 - Site/Cloud operators: Overview of site/cloud activity
 - VO Managers: Overview of the whole activity
 - End users: State of subscriptions


Throughput (MB/s)



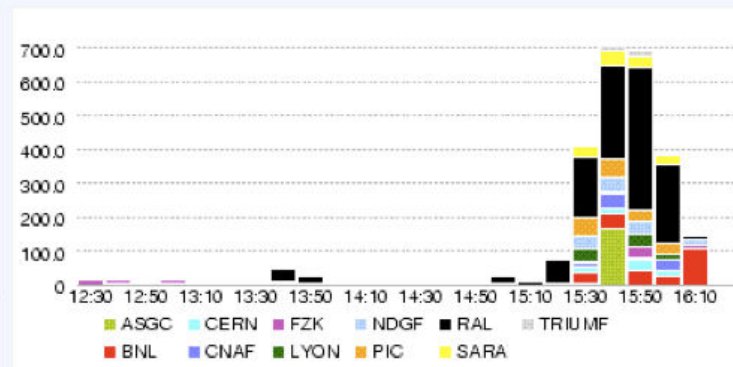
Data Transferred (GBytes)





Completed File Transfers



Total Number Transfer Errors




- User access to data is chaotic by nature
- Behavior can not be foreseen
- Computing model: Users should send their jobs to the data
- Users want their data immediately:
 - dq2-get&direct storage access vs. SS subscriptions
 - dq2-get abuse can lead to
 - Storage Element overload
 - Network congestion
 - DDM service degradation



Single user brings down a site




- Not all data will be accessed with equal regularity:
 - Tackle user behavior and anticipate future workload by tracing the user's operations
 - Predict workload using ARMA models
 - Distribute important data before it is needed: HOTDISK spacetoken
 - Sites are responsible to replicate files in this spacetoken to different pools




CERN

Site	Group	Role	Used by role(TB)	Booked by role(TB)	ΣUsed(TB)	ΣBooked(TB)	DQ2 Used(TB)	SRM Used(TB)	SRM Total(TB)
CERN-PROD	DET-MUON	/atlas/det-muon/role=production	0.99	1	1.05	1	19.49	130.39	201.28
		/atlas/role=production	0.06	0					
	PERF-EGAMMA	/atlas/perf-egamma/role=production	0.23	1	0.25	1			
		/atlas/role=production	0.01	0					
	PERF-FLAVTAG	/atlas/role=null	0.07	0	0.07	1			
		/atlas/perf-flavtag/role=production	0.0	1					
	PERF-IDTRACKING	/atlas/role=production	0.09	0	0.09	1			
		/atlas/perf-idtracking/role=production	0.0	1					
	PERF-JETS	/atlas/perf-jets/role=production	0.57	1	0.64	1			
		/atlas/role=null	0.07	0					
	PHYS-EXOTICS	/atlas/phys-exotics/role=production	0.0	1	0.0	1			
	PHYS-GENER	/atlas/phys-gener/role=production	2.98	10	3.18	10			
		/atlas/role=production	0.2	0					
	PHYS-HI	/atlas/phys-hi/role=production	0.13	3	0.13	3			
	PHYS-HIGGS	/atlas/role=production	0.0	0	0.0	1			
		/atlas/phys-higgs/role=production	0.0	1					
	PHYS-SUSY	/atlas/role=production	0.0	0	6.26	20			
		/atlas/phys-susy/role=production	6.26	20					
	PHYS-TOP	/atlas/role=production	4.59	0	4.59	2			
		/atlas/phys-top/role=production	0.0	2					
PROJ-CTB	-	0	0	0.0	0				
PROJ-SIMCOS	-	0	0	0.0	0				
PROJ-SIT	/atlas/role=null	0.0	0	0.02	0				
	/atlas/role=production	0.01	0						
TRIG-DAQ	/atlas/trig-daq/role=production	3.22	1	3.22	1				


- ATLAS DDM is successfully handling all experiment data since 2005
- Major challenges passed successfully: ATLAS DDM is capable of handling far more subscriptions than defined in the Computing Model
- User workload will become the critical factor once we have real data




- “Distributed Data Management in the ATLAS Experiment”, Mario Lassnig
- “Distributed Data Management in ATLAS”, Ricardo Rocha
- “ATLAS, the Grid and the UK”, Roger Jones
- “The ATLAS Computing Model”, D. Adams et al.
- “ATLAS STEP09”, Graeme Stewart



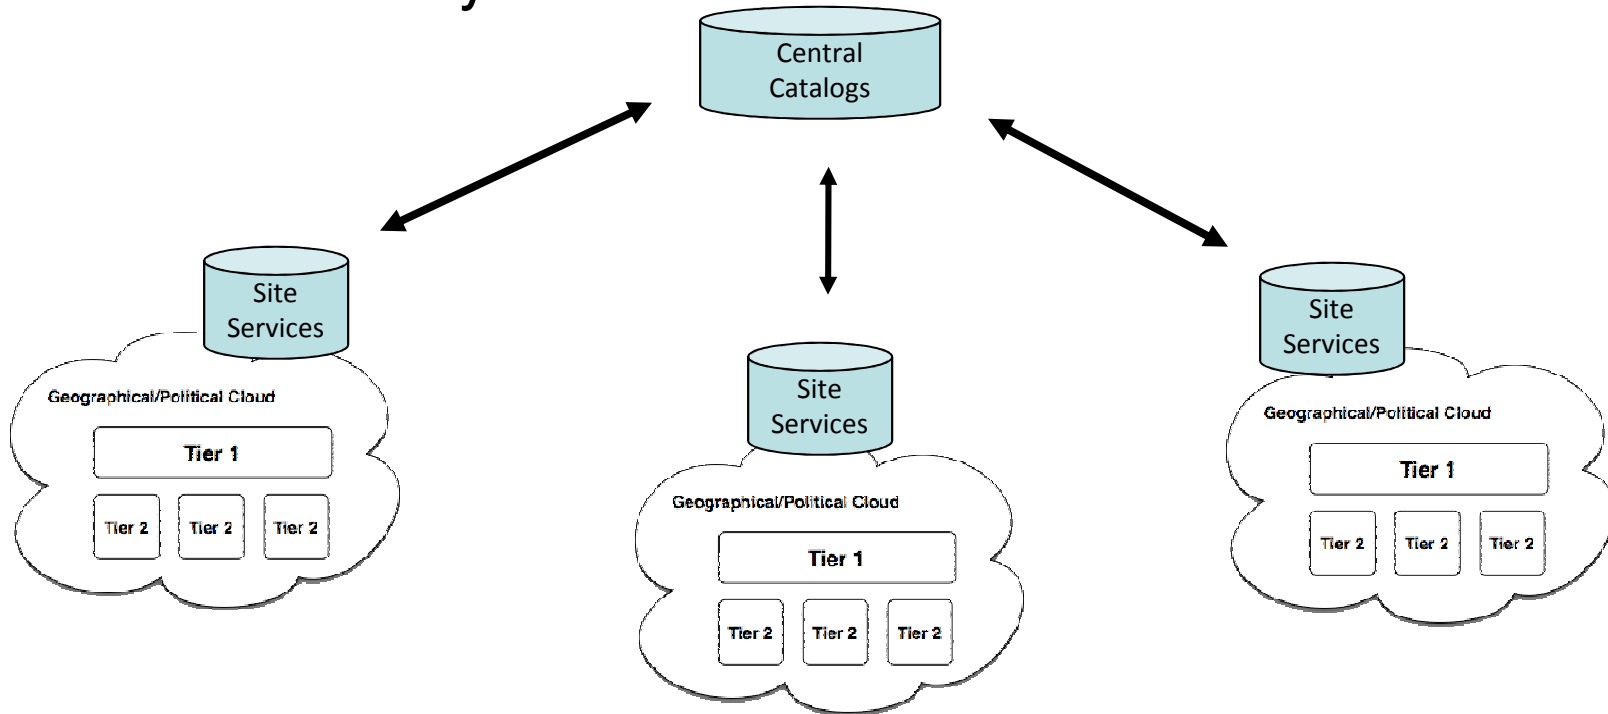
Backup slides




Name	Description	Size
Monte Carlo	Event generator output	
RAW	Detector output – Byte stream format	1.6 MB/ev
Event Summary Data (ESD)	Full output of reconstruction – Object format	1.0 MB/ev
Analysis Object Data (AOD)	Summary of reconstruction. Primary analysis data	0.2 MB/ev
Tag	Thumbnail of each event used for identifying interesting events at the analysis stage	0.01 MB/ev
Derived Physics Data (DPD)	Skimmed, slimmed, thinned events derived from AODs	0.01 MB/ev

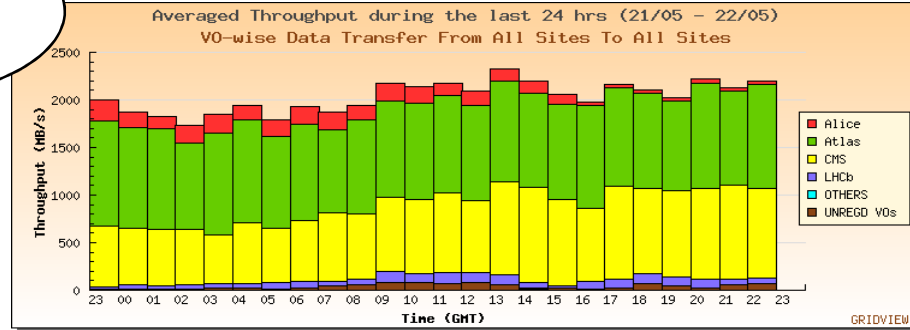



- Central Catalogues and Site Services hosted at CERN
- LFC and FTS at Tier-1s
- SRM at every site

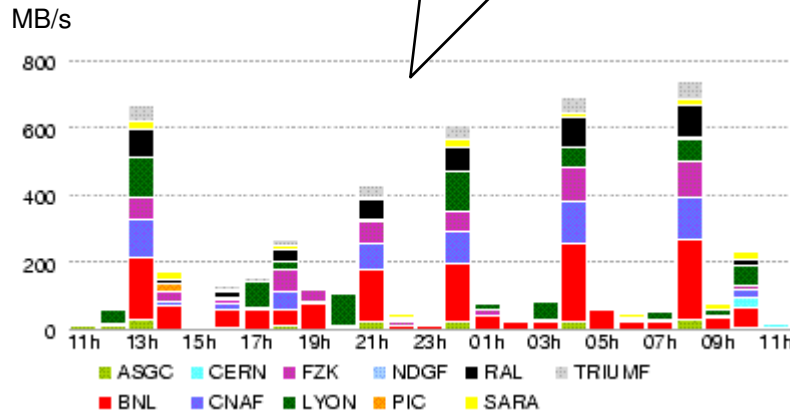





All Experiments in the game




Throughput tests: Burst subscriptions injected every 4 hours and immediately honored

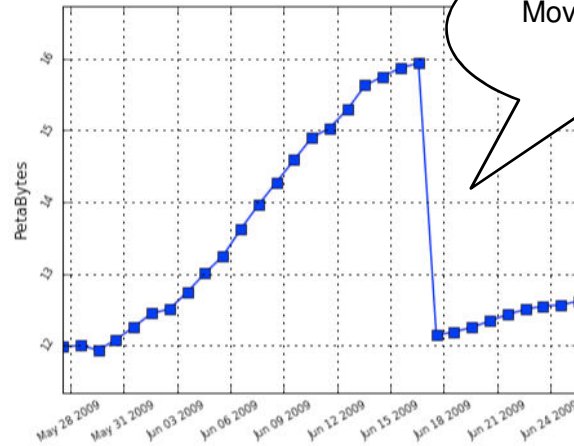


T0->T1s throughput





Total GRID disk usage according to dq2



Moved 4PB in 2 weeks

Traffic rates up to 5.5 GB/s

Activity Summary ('2009-06-12 17:40' to '2009-06-12 21:40')

Click on the cloud name to view list of sites

Cloud	Transfers			Registrations		
	Efficiency	Throughput	Successes	Datasets	Files	Transfer
ASGC	99%	397 MB/s	6286	833	6299	59
BNL	84%	1128 MB/s	23903	855	24000	4698
CERN	100%	334 MB/s	5712	133	5692	10
CNAF	98%	561 MB/s	7638	382	7652	150
FZK	85%	556 MB/s	6852	381	6858	1200
LYON	96%	620 MB/s	8510	749	8506	325
NDGF	84%	137 MB/s	913	102	907	170
PIC	93%	429 MB/s	3365	908	3361	256
RAL	99%	838 MB/s	13324	1656	13726	125
SARA	53%	262 MB/s	3186	134	3171	2882
TRIUMF	100%	297 MB/s	4704	169	4706	13