

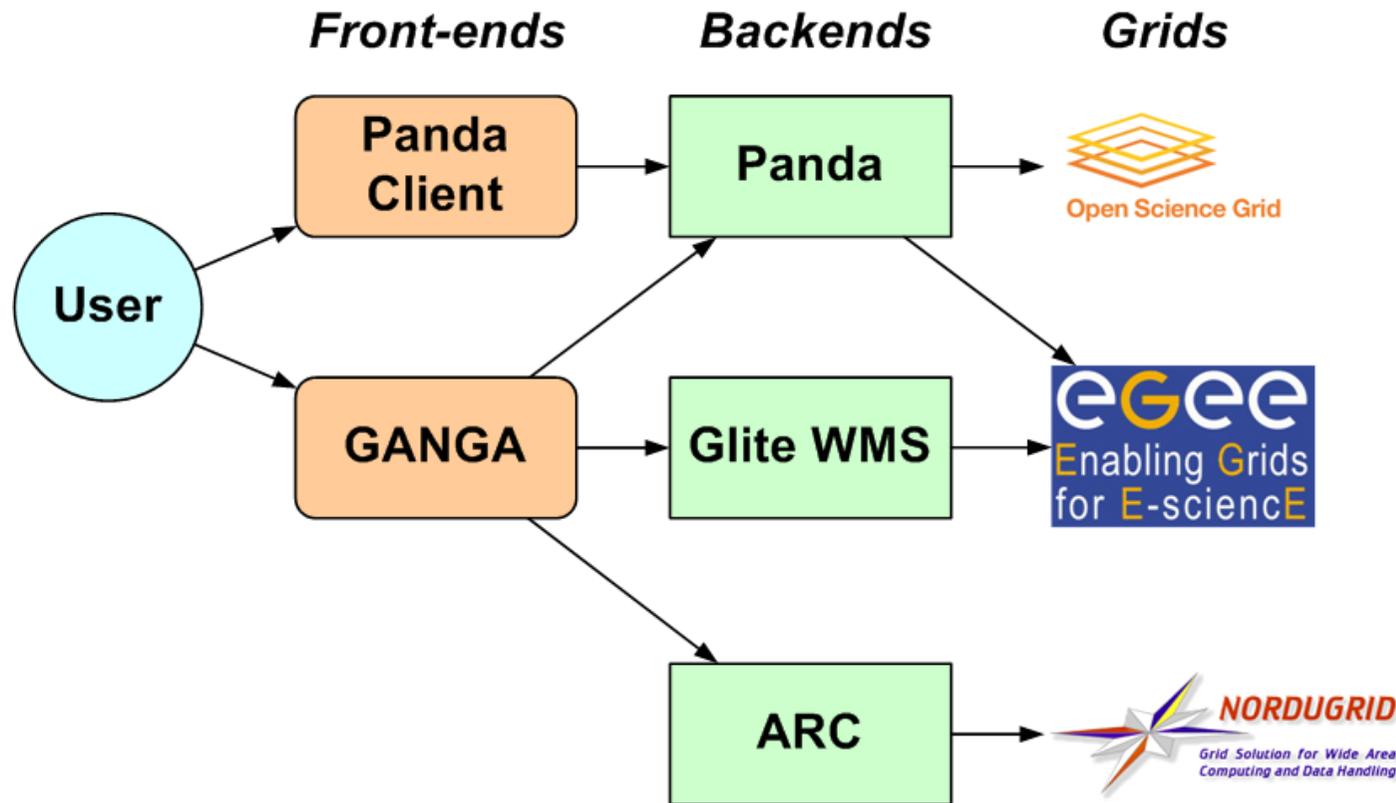


ATLAS Distributed Analysis Experiences in STEP'09

Dan van der Ster
for the DA stress testing team and ATLAS
Distributed Computing

*WLCG STEP'09 Post-Mortem
10 July 2009*

Distributed Analysis in ATLAS



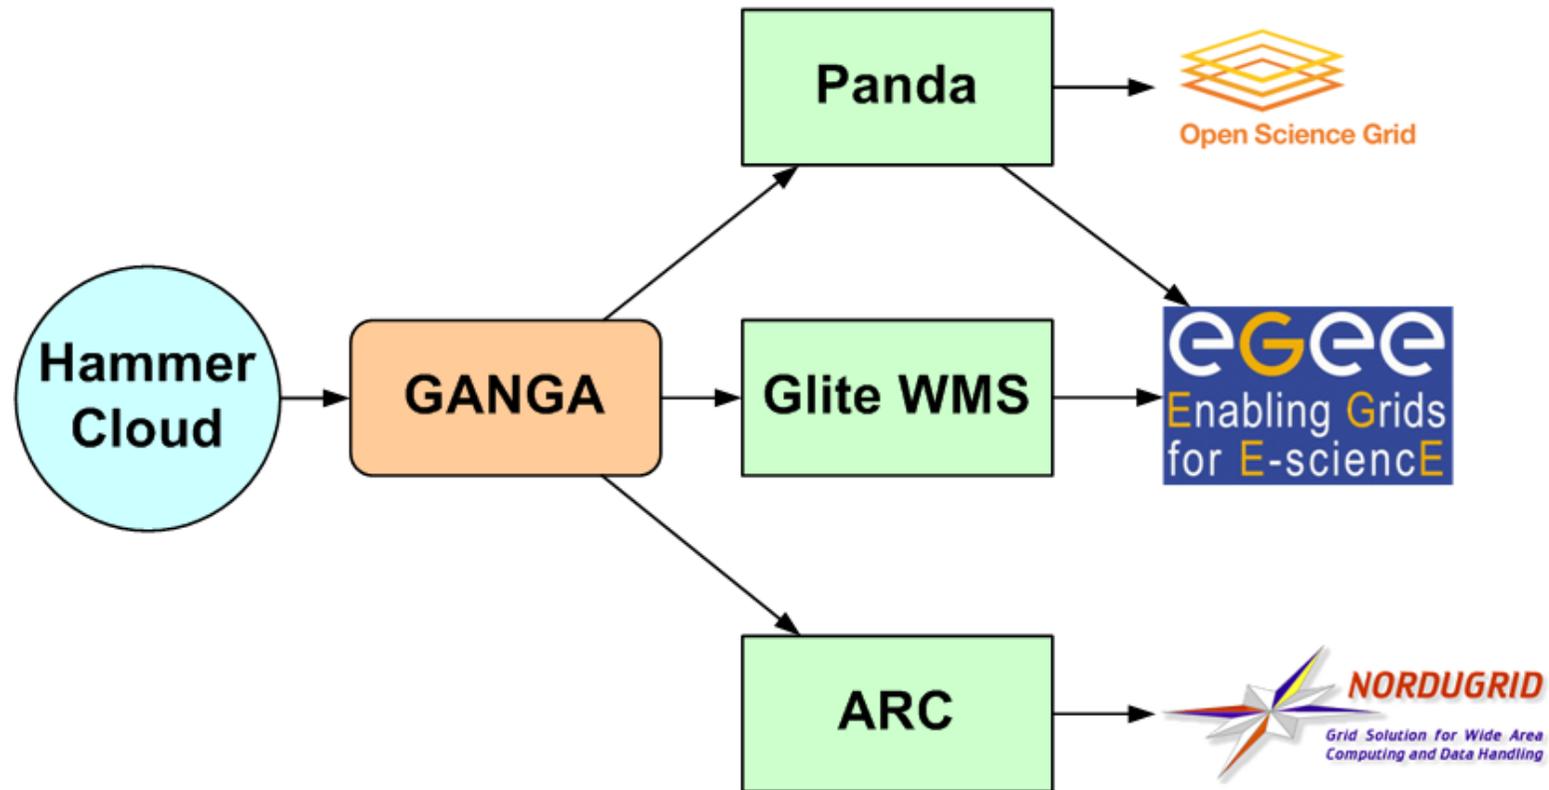
- Classic Athena analyses:
 - various input data sizes: ESD, AOD, DPD
 - TAG for direct event access
- Calibration & Alignment: RAW data and remote DB access
- Private MC production
- everything else (ROOT, arbitrary executables...)



DA Tests During STEP'09

- Stress testing with HammerCloud
 - "Infinite" real analysis jobs to participating sites (all T2's + a few T1's)
 - Most (~all) of our STEP'09 results drawn from these tests
- Coordinated user activity
 - US and DE challenges
- Background (normal) user activity
 - Some sites were quite busy without HC's help

HC: Workflows tested



- Classic Athena analyses:
 - various input data sizes: **ESD**, **AOD**, **DPD**
 - **TAG** for direct event access
- Calibration & Alignment: **RAW** data and remote **DB access**
- **Private MC** production
- everything else (**ROOT**, **arbitrary executables...**)



HC Testing Details

Four AOD analysis jobs:

WMS: 30GB per job, 2 Local protocol, 2 FileStager

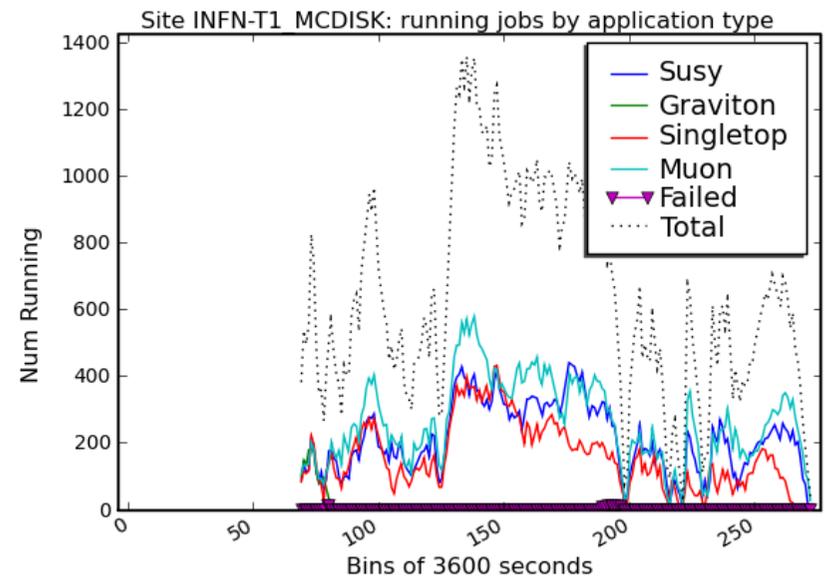
Panda: 5GB per job, switched to 10GB ~Sat.

ARC/NG: 5GB per job

Note: we expect ~4-6 times more Panda jobs than WMS jobs at the shared sites.

All analyses run at ~15-19Hz with local fast files.

CNAF running 3-4 analyses throughout STEP09. "Graviton" job (using FS) was turned off at WMS sites because of large log files and memory leak. (F. Galezzi)



HC: Global Summary

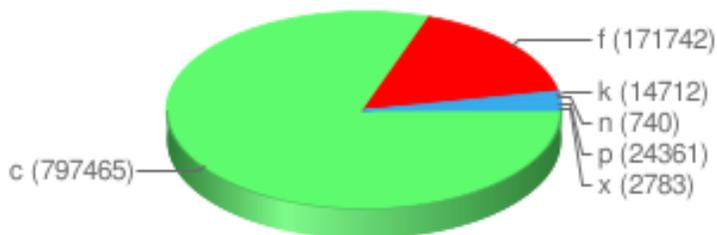
~1M jobs submitted -- 82.3% succeeded

Total 26.3B events processed

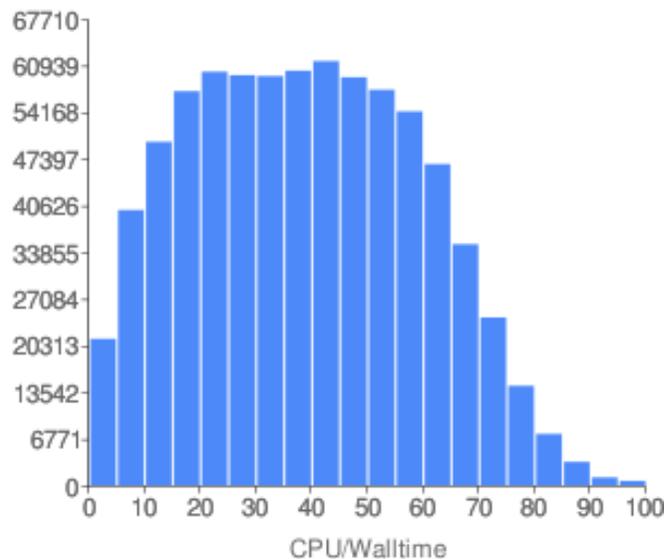
Mean Events/s = 7.7Hz

Mean CPU/Walltime = 0.39

Overall Efficiency

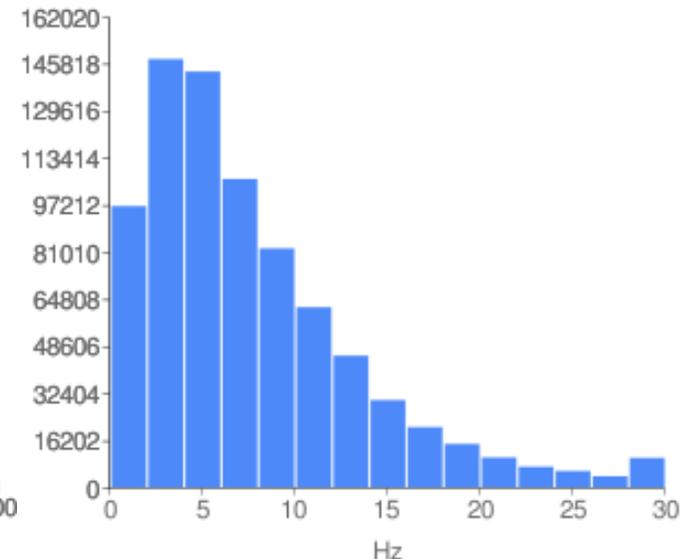


Overall CPU/Walltime



$\mu=39.3 \sigma=23.2$

Overall Events/s



$\mu=7.7 \sigma=6.5$

Global Efficiency: 82.3%

Global Rate: 28.6kHz



Tabulated Results

All the following results are available at:

<http://gangarobot.cern.ch/st/step09summary.html>

Definitions:

Efficiency = #completed/(#completed+#failed)

Hz = avg(#events/(stoptime-starttime))

Note: potential bias between Panda/WMS

CPU/Walltime:

For gLite: = "Percent of CPU this job got..." as reported by time athena...

For Panda: = 100 * cpuConsumptionTime / cpuConversion / (endTime - startTime)

Cloud Summary: Throughput

Cloud	# Jobs	# Successful	# Failed	Efficiency	#files	#events ▲	Hz	CPU/Wall
FR	166427	144658	21769	0.869	557571	5050911395	9.3	44.5
DE	176076	135218	40858	0.768	557395	4991372555	7.9	40.8
US	163004	153240	9764	0.940	465393	4169999722	6.0	33.9
UK	181394	145222	36172	0.801	439084	3984651767	7.1	39.8
IT	59163	52990	6173	0.896	311061	2798011153	6.5	32.6
ES	72562	62565	9997	0.862	236690	2150478621	10.0	44.3
NL	66632	37171	29461	0.558	154452	1352903529	7.9	35.0
TW	24178	19544	4634	0.808	86293	833817261	15.2	48.4
CA	41890	32306	9584	0.771	87117	757520054	4.2	34.4
NG	16730	14551	2179	0.870	20179	172708698	7.3	

Cloud Summary: Efficiency

Cloud	# Jobs	# Successful	# Failed	Efficiency ▲	#files	#events	Hz	CPU/Wall
US	163004	153240	9764	0.940	465393	4169999722	6.0	33.9
IT	59163	52990	6173	0.896	311061	2798011153	6.5	32.6
NG	16730	14551	2179	0.870	20179	172708698	7.3	
FR	166427	144658	21769	0.869	557571	5050911395	9.3	44.5
ES	72562	62565	9997	0.862	236690	2150478621	10.0	44.3
TW	24178	19544	4634	0.808	86293	833817261	15.2	48.4
UK	181394	145222	36172	0.801	439084	3984651767	7.1	39.8
CA	41890	32306	9584	0.771	87117	757520054	4.2	34.4
DE	176076	135218	40858	0.768	557395	4991372555	7.9	40.8
NL	66632	37171	29461	0.558	154452	1352903529	7.9	35.0

Site Summaries

We now look at the site performance.

The metric used to rank the sites is debatable.

I will present sites ordered by normalized throughput:
Millions of events processed / AOD data share

Sites with share < 0.1 have been filtered out for now
Complete tables are in the "backup" slides.

My interpretation of the metric:

If you perform well it means:

- (a) you performed well
- (b) your data share is too small
- (c) both

If you do not perform well it means:

- (a) you have an excuse e.g. many background user jobs
- (b) you did not perform well



Site Summary (1)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▼	Hz	CPU/W
ES	PIC	41117	0.948	1177	0.10	64.0	11776	10.8	46.3
FR	IN2P3-LAPP	33661	0.970	1233	0.15	64.0	8222	11.7	53.4
DE	CSCS-LCG2	29541	0.918	970	0.17	70.0	5707	6.7	43.1
UK	UKI-SCOTGRID-GLASGOW	78327	0.974	1846	0.40	148.0	4615	8.4	42.6
DE	GOEGRID	13798	0.776	673	0.15	76.0	4488	5.7	26.0
NL	RU-PROTVINO-IHEP	14690	0.896	550	0.15	60.0	3667	11.2	38.6
ES	LIP-LISBON	10172	0.545	320	0.10	77.0	3205	11.4	39.7
FR	BEIJING-LCG2	13211	0.846	318	0.10	54.0	3189	8.0	51.2
FR	IN2P3-LPC	12857	0.913	468	0.15	84.0	3123	7.2	47.1
US	NET2	29906	0.954	960	0.33	117.0	2909	6.6	32.9
US	MWT2	35195	0.968	1051	0.40	481.0	2628	6.7	37.9
FR	GRIF-LAL	22070	0.826	613	0.24	105.0	2555	9.8	42.6
US	AGLT2	64325	0.954	1528	0.60	358.0	2547	6.0	27.1
UK	UKI-NORTHGRID-LIV-HEP	10717	0.831	349	0.15	81.0	2328	5.9	50.2
DE	CYFRONET-LCG2	24271	0.696	386	0.17	99.0	2274	7.4	32.7
DE	DESY-ZN	26759	0.977	1115	0.50	157.0	2231	11.3	56.7
UK	UKI-LT2-RHUL	30258	0.702	505	0.23	185.0	2196	6.2	36.8
IT	INFN-T1	32878	0.908	1886	1.00	-1.0	1886	6.5	33.2

Great normalized throughput. But check your efficiency and event rate.

Site Summary (2)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▲	Hz	CPU/W
IT	INFN-NAPOLI-ATLAS	12899	0.949	527	0.30	156.0	1757	7.5	37.2
DE	LRZ-LMU	15814	0.711	278	0.16	146.0	1740	8.7	40.1
DE	MPPMU	10641	0.553	270	0.16	89.0	1689	6.8	34.6
UK	UKI-SOUTHGRID-RALPP	9565	0.862	418	0.25	76.0	1675	4.5	24.3
ES	UAM-LCG2	10179	0.888	393	0.25	126.0	1572	8.7	46.1
TW	TW-FTT	19085	0.803	756	0.50	115.0	1512	17.6	49.6
FR	TOKYO-LCG2	36470	0.986	1501	1.00	293.0	1501	11.0	46.6
DE	DESY-HH	18723	0.942	681	0.50	180.0	1362	9.3	43.9
NL	TR-10-ULAKBIM	8616	0.866	230	0.20	80.0	1153	4.1	32.5
NL	JINR-LCG2	11110	0.468	221	0.20	63.0	1107	8.6	29.1
UK	UKI-NORTHGRID-SHEF-HEP	6227	0.821	102	0.10	36.0	1020	6.6	50.7
US	SLACXRD	14716	0.889	337	0.34	215.0	993	4.0	64.8
IT	INFN-ROMA1	7848	0.864	270	0.30	131.0	901	6.3	25.6
FR	GRIF-SACLAY	8672	0.935	213	0.26	63.0	822	8.9	45.2
DE	UNI-FREIBURG	8805	0.604	126	0.17	106.0	743	7.7	45.5
CA	ALBERTA-LCG2	5370	0.884	185	0.25	24.0	740	4.8	28.3
IT	INFN-FRASCATI	2748	0.877	71	0.10	38.0	715	4.3	29.5
ES	IFAE	7588	0.928	161	0.25	68.0	647	5.7	29.7
UK	UKI-NORTHGRID-MAN-HEP1	7948	0.633	83	0.15	40.0	559	3.8	45.8

Good normalized throughput. Again, check efficiency and event rate.

Site Summary (3)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▲	Hz	CPU/W
UK	UKI-NORTHGRID-MAN-HEP2	9347	0.382	71	0.15	40.0	479	2.8	32.0
FR	GRIF-LPNHE	10312	0.857	217	0.50	140.0	435	7.9	34.2
DE	FZK-LCG2	16286	0.681	424	1.00	-1.0	424	4.1	23.0
US	SWT2	8053	0.788	134	0.33	40.0	406	2.0	16.7
CA	SFU-LCG2	10832	0.950	270	0.72	157.0	375	5.8	26.6
FR	RO-02-NIPNE	2810	0.432	37	0.10	51.0	373	2.9	19.2
ES	IFIC-LCG2	2372	0.753	87	0.25	68.0	350	12.4	63.8
NL	ITEP	864	0.991	34	0.10	10.0	347	7.5	38.5
UK	UKI-NORTHGRID-LANCS-HEP	6819	0.565	78	0.30	140.0	260	4.8	24.5
FR	IN2P3-CC	10486	0.761	212	1.00	-1.0	212	3.1	26.7
DE	HEPHY-UIBK	2056	0.467	19	0.10	36.0	193	4.5	19.5
NG	NDGF-T1	16730	0.870	172	1.00	-1.0	172	7.3	
US	BNL-OSG2	10809	0.911	157	1.00	-1.0	157	7.1	34.5
NL	NIKHEF-ELPROD	5174	0.875	143	1.00	-1.0	143	7.5	39.6
IT	INFN-MILANO	2790	0.612	42	0.30	78.0	140	1.7	19.2
TW	AUSTRALIA-ATLAS	4548	0.891	67	0.50	119.0	135	6.0	42.4
NL	RRC-KI	5339	0.402	67	0.50	171.0	135	2.9	18.3
DE	WUPPERTALPROD	5783	0.210	19	0.17	63.0	117	4.7	29.7
FR	RO-07-NIPNE	1516	0.157	5	0.10	33.0	59	5.7	27.0
NL	SARA-MATRIX	1461	0.763	41	1.00	-1.0	41	3.3	37.2
UK	UKI-LT2-UCL-CENTRAL	473	0.222	3	0.10	16.0	35	3.8	28.3
NL	RU-PNPI	204	0.632	4	0.15	65.0	28	0.8	11.1
ES	LIP-COIMBRA	1134	0.160	9	0.50	115.0	19	6.4	22.1
TW	TAIWAN-LCG2	545	0.321	9	1.00	125.0	9	13.2	94.2

Low throughput

Reasons:

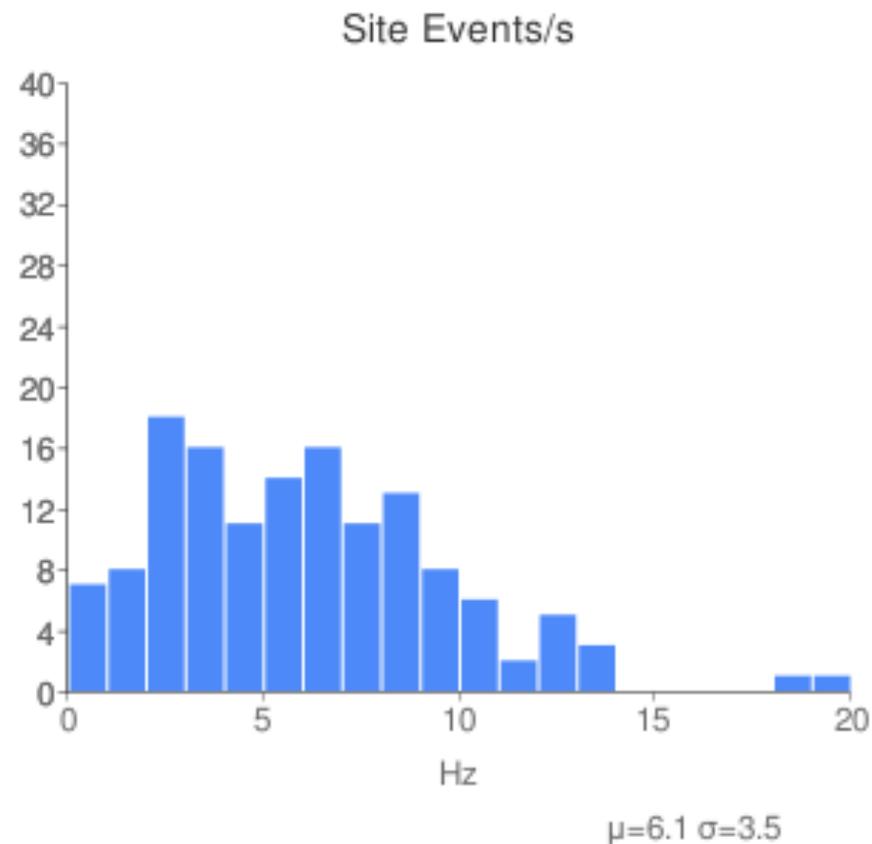
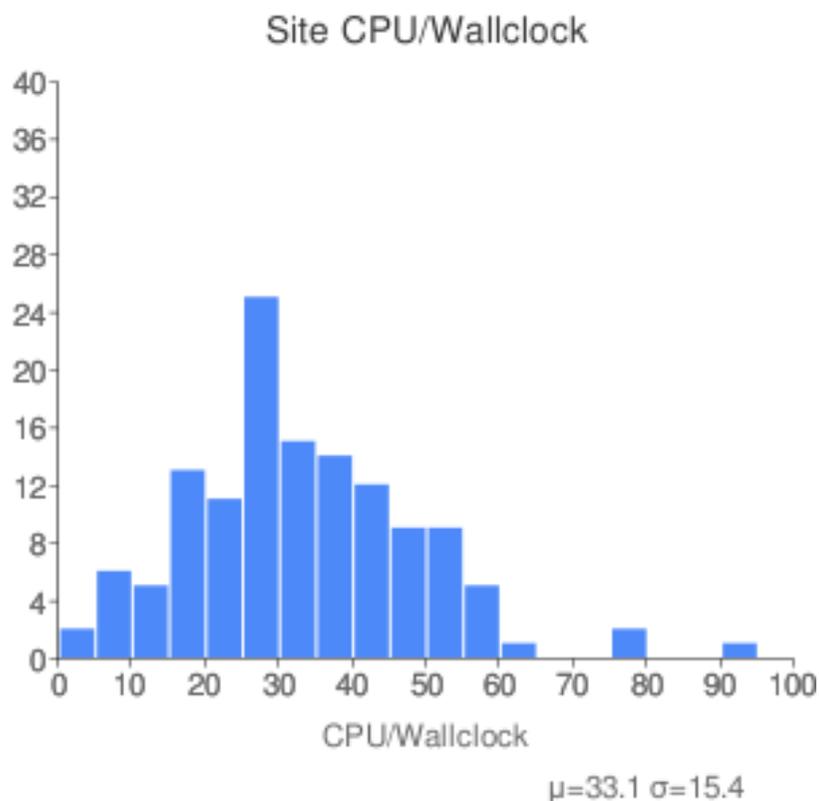
Many other jobs

Low efficiency

Low event rates

... ?

Performance



~All sites have I/O
bottleneck

*Mean rate is ~1/3rd
your laptop*



A Note on Efficiencies

We have always recorded the *success efficiency*, but probably more interesting would be the *time-weighted efficiency*:

time-weighted eff = time spent in successful jobs / total time

note that **success eff == time eff** iff $\text{mean}(\text{succ. time}) == \text{mean}(\text{fail time})$

During STEP09:

mean successful Panda job took 4584 s.

mean failed Panda job took 5721 s.

overall completion efficiency for Panda was 0.856

overall time efficiency was 0.827

for WMS we didn't record the times for failed jobs, but with hanging posix I/O connections failing jobs could waste *a lot* of time.

In future we should focus on the time-weighted efficiency metric to better understand how the sites are being utilized.



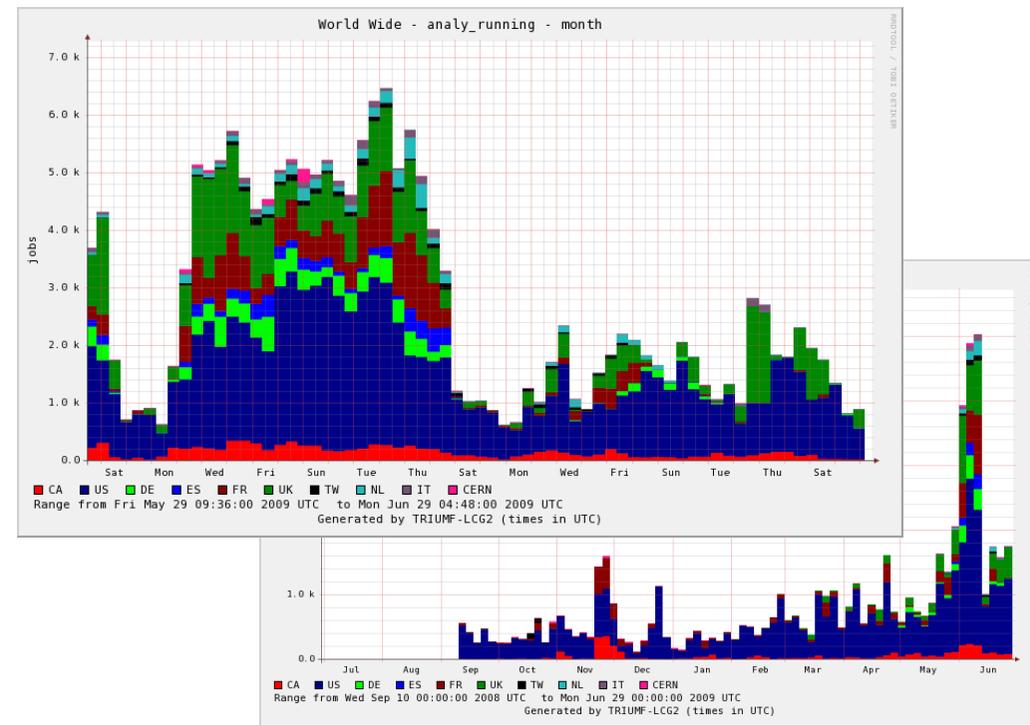
PanDA Results

Panda uses "copy mode" at most sites.

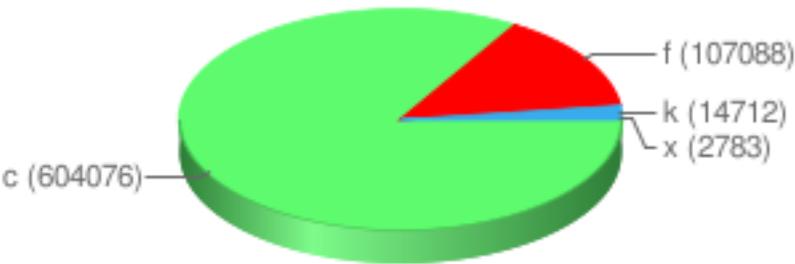
Record job numbers for Panda DA: peaked at ~6500 concurrent

84.9% success efficiency

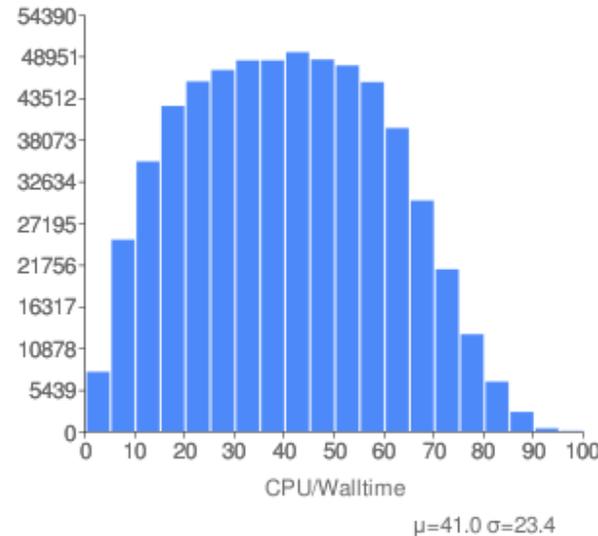
Slightly above average CPU/wall and Hz



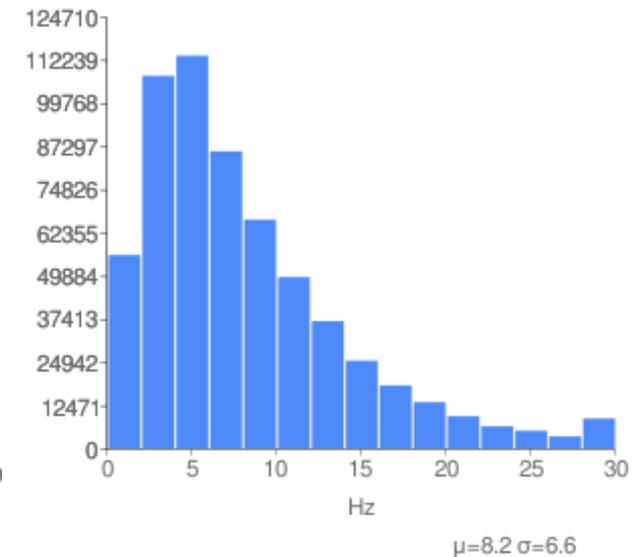
Panda Efficiency



Panda CPU/Walltime



Panda Events/s



Panda Analysis Conclusions

1. Pilots:

DE, IT didn't run enough Panda jobs, probably queue depth was not large enough

Need to better organize the pilot factories

2. Data access timeouts:

40% of the failures were "Get error: lcg-cp get was timed out after 1800 seconds"

Need a way to discover terminally ill vs. slow transfers

3. "Hidden failures" at xrootd sites. (athena didn't read all events)

10% errors at LYON, 2% at SLAC, and 1.4% at SWT2_CPB

Need reliable run report from athena, and post job validation



gLite WMS Results

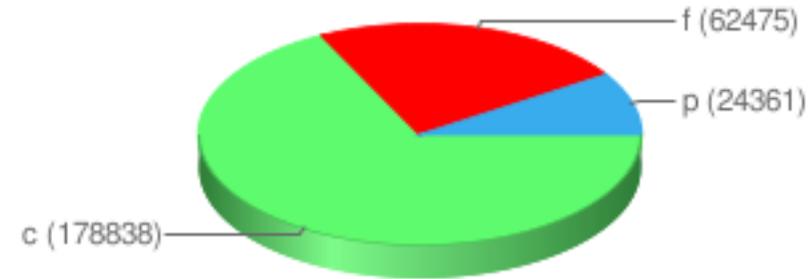
WMS jobs used mostly posix I/O:

For WMS analysis we have been more *adventurous* ...
FileStager "copy" mode was not used as much as we had intended.

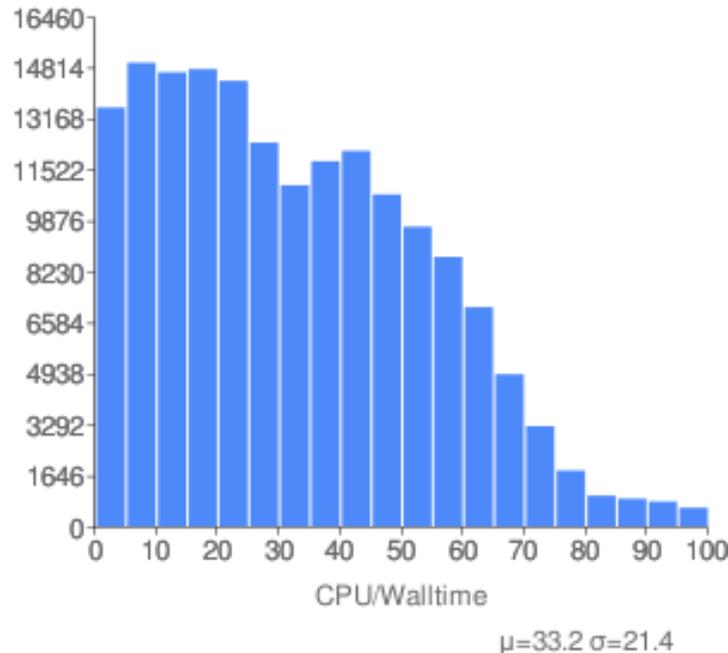
74% success efficiency
below avg. Hz and CPU/Wall

Lower WMS efficiency results from long jobs higher susceptibility to I/O flakiness and poor posix I/O performance

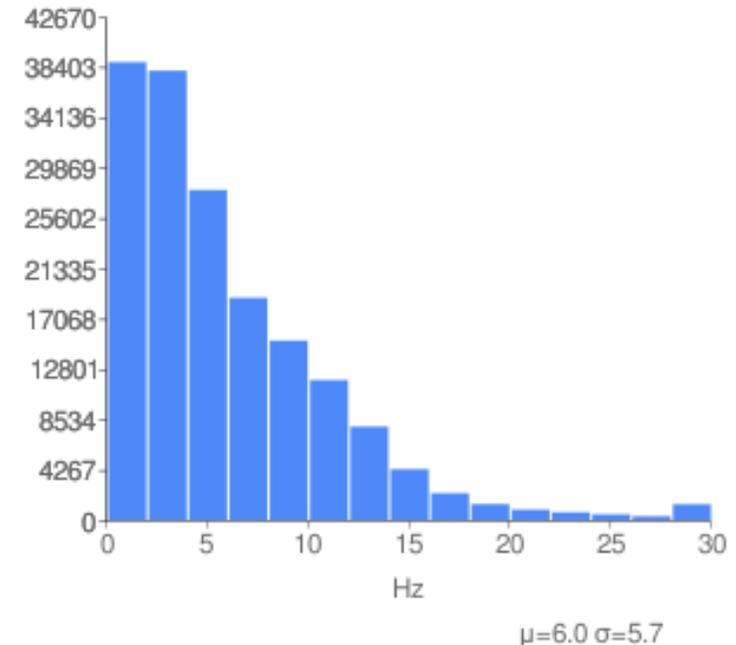
gLite WMS Efficiency



gLite WMS CPU/Walltime



gLite WMS Events/s



gLite WMS issues

1. Longer jobs lead to more failures with memory leaks, huge log files, etc...

Clients should default to fine & intelligent splitting

2. Ganga-MonAlisa link was down so we couldn't easily track the WMS jobs

Need more reliable central monitoring of the WMS jobs

3. Post analysis is difficult. Need to download log files and parse:

Post process on WN. Adopt error codes from Panda.

4. On average Posix I/O did not perform well. This is the default for WMS jobs, so we should change it.

Data access method should be same for Panda and WMS jobs, not user definable, and stored in central IS (it is already in Panda DB)

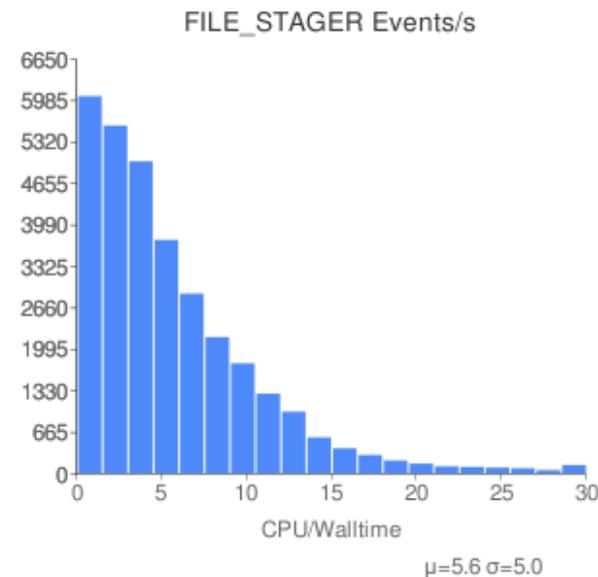
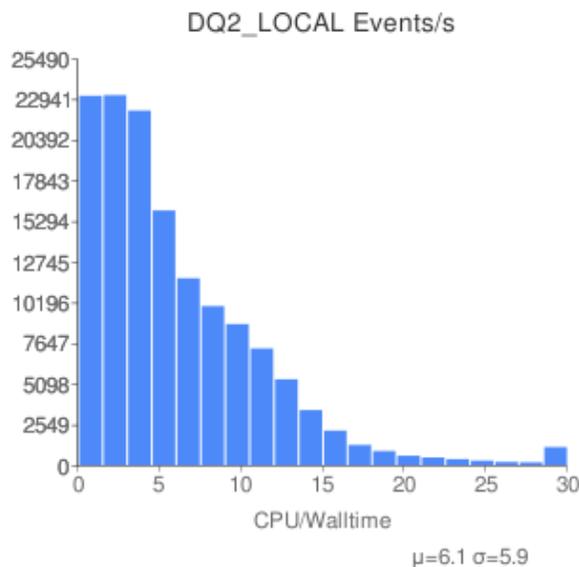
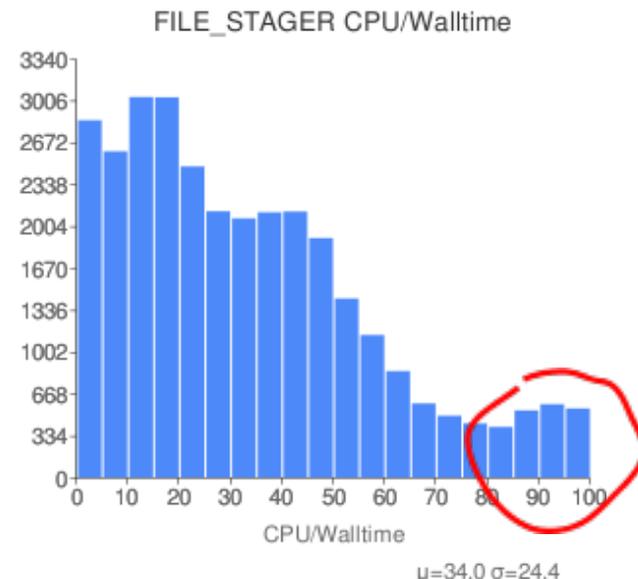
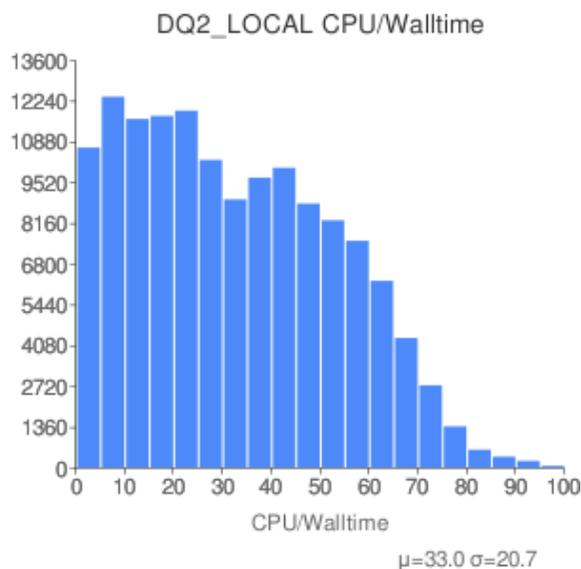


Data Access Methods

The "WMS" tests were mostly direct I/O.

By chance the 2 FS analyses were memory leaky and produced huge log files so we stopped these early.

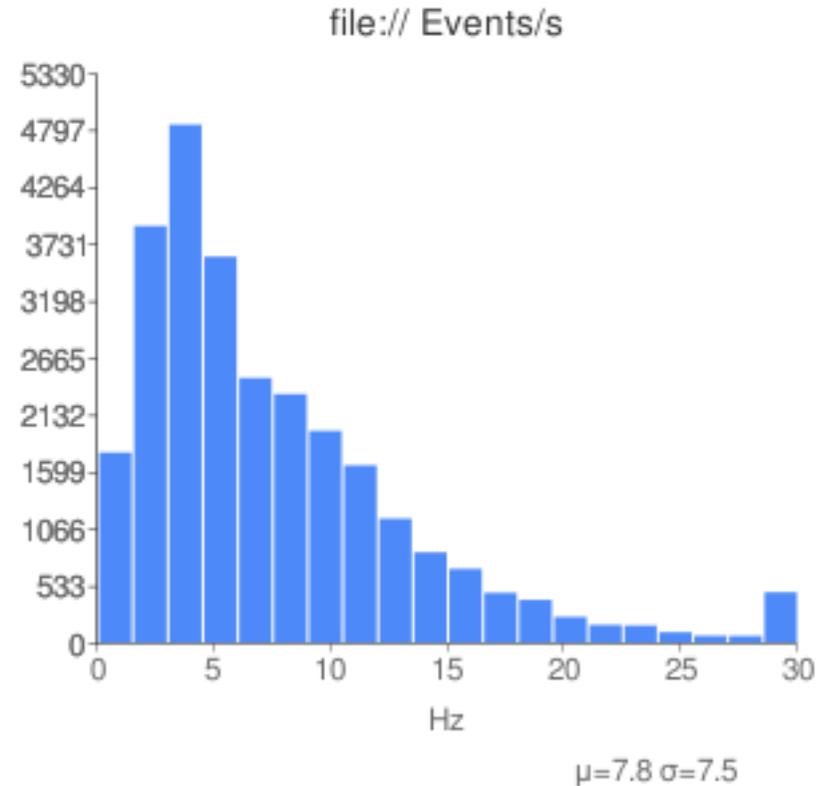
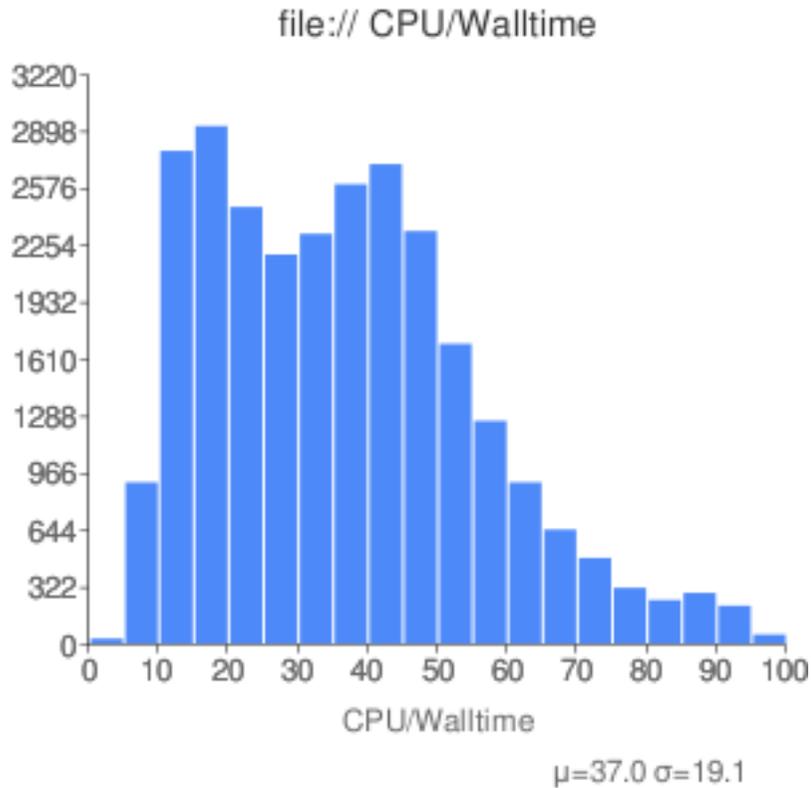
In previous tests FS vastly outperformed direct access at most sites; here it did not. Apparently, the usage of direct access at a site penalized the FS jobs running at the same site.



Should move to FileStager where beneficial and possible.



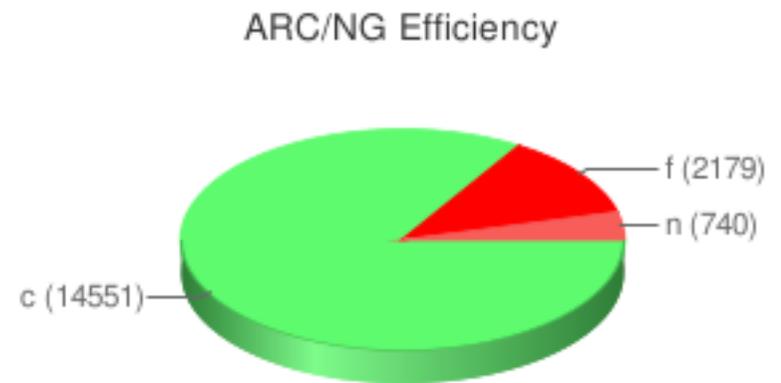
file:// protocol at CNAF, IFIC, LIP



Better than dcap/rfio and FileStager:
only a few very slow jobs
still the storage is the bottleneck



ARC/NG Results



Efficiency was high, but middleware issues led to low throughput at the sites (due to data stage-in delay to the site caches). Solutions are in development.

More stress testing is needed at NG sites

Real User Analysis During STEP'09

There was *a lot* of real analysis during STEP'09.

Coordinated efforts in US and DE (that I have results from).

US: 14 users processed a 1M evt sample (for 75M total events processed with ~75% efficiency)

Job success efficiency (%)

AGLT2	MWT2	NET2	SLAC	SWT2
59*	80	74	84	75

* one user made many attempts here before overcoming user config issue - from slides of J. Cochran



Coordinated Analysis in DE

User	Site/Cloud	N-Job fail	N-Job succe	Efficiency	N-Files	N-Evts	Wall-time	evt rate Hz	Cpu/Wall
1									
	DE	223	1335	85.69	6675	12894675	21815080	0.59	0.26
	FR	115	0	0					
2									
	DE	8162	8046	49.64	10515	73986440	50649545	1.46	0.43
	UK	0	85	100	848	211493	207363	1.02	0.43
	FR	329	219	39.96	2137	533588	404915	1.32	0.4
	IT	61	222	78.45	2210	551024	256054	2.15	0.71
	US	770							
3									
	DE	0	158	100	158	964496	84730	11.38	0.76
4									
	DE	22	1630	98.67	1630	1630000	239599	6.8	0.49

19.06.09

1

Günter Duckeck, LMU

Results from 4 users:

~90M events

varied efficiencies



Background User Analysis

We know that real analysis continued behind HC, but exact numbers aren't known worldwide (because of WMS monitoring lapse).

In US we know:

- 426309 user analysis jobs during STEP'09
 - 173648 (~40%) from HC
 - Remaining 60% from 198 other users
 - 172k succeeded, 80k failed (of which 40-50k killed by user)
- numbers from K. De.*

In general, looking at the overall efficiency of real user jobs is almost meaningless

- need to perform detailed error analysis to identify site/service vs. user errors, which is possible but onerous with existing tools.



Distributed Analysis User Support

- DA user support is provided by the DAST shifters.
 - 1 EU + 1 US shift per week
- Made no special preparations for STEP'09
 - i.e. normal shift rotation without extra effort
- The numbers:
 - 2-12 June (STEP'09) saw 76 threads in the DA help forum
 - No extraordinary problems or complaints
 - Pre STEP (22 May - 1 June): 73 threads
 - Post STEP (13-23 June): 76 threads
- The user support model succeeded during STEP'09, though we are still planning to increase the shifter effort to handle the influx of users coming with real data.

Conclusions

- The distributed analysis tests during STEP'09 were a success for ATLAS:
 - The services were exercised at record levels.
 - Global Efficiency: 82.3% Global Rate: 28.6kHz
 - Real users continued to get real work done
- The problems are known...
 - dcap/rfio access is troublesome at heavy load
 - our pilot factories need to be better organized
 - need better communication with athena
 - monitoring can be improved (as always)
 - and others.
- ...but most have known solutions.
- Testing continues now... no waiting for SEPT



Backup: all sites (1)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▼	Hz	CPU/W
CA	VICTORIA-LCG2	7941	0.775	193	0.01	35.5	19333	4.1	39.7
UK	UKI-SOUTHGRID-OX-HEP	4420	0.871	269	0.02	66.0	13455	4.4	21.7
ES	PIC	41117	0.948	1177	0.10	64.0	11776	10.8	46.3
FR	IN2P3-LAPP	33661	0.970	1233	0.15	64.0	8222	11.7	53.4
DE	CSCS-LCG2	29541	0.918	970	0.17	70.0	5707	6.7	43.1
CA	TORONTO-LCG2	17747	0.626	108	0.02	20.0	5439	2.5	41.5
UK	UKI-SCOTGRID-GLASGOW	78327	0.974	1846	0.40	148.0	4615	8.4	42.6
DE	GOEGRID	13798	0.776	673	0.15	76.0	4488	5.7	26.0
NL	RU-PROTVINO-IHEP	14690	0.896	550	0.15	60.0	3667	11.2	38.6
ES	LIP-LISBON	10172	0.545	320	0.10	77.0	3205	11.4	39.7
FR	BEIJING-LCG2	13211	0.846	318	0.10	54.0	3189	8.0	51.2
FR	IN2P3-LPC	12857	0.913	468	0.15	84.0	3123	7.2	47.1
FR	IN2P3-CPPM	12143	0.573	154	0.05	26.0	3091	4.7	24.3
US	NET2	29906	0.954	960	0.33	117.0	2909	6.6	32.9
UK	UKI-SOUTHGRID-CAM-HEP	5050	0.708	143	0.05	32.0	2864	8.4	33.7
US	MWT2	35195	0.968	1051	0.40	481.0	2628	6.7	37.9
FR	GRIF-LAL	22070	0.826	613	0.24	105.0	2555	9.8	42.6
US	AGLT2	64325	0.954	1528	0.60	358.0	2547	6.0	27.1
UK	UKI-NORTHGRID-LIV-HEP	10717	0.831	349	0.15	81.0	2328	5.9	50.2
DE	CYFRONET-LCG2	24271	0.696	386	0.17	99.0	2274	7.4	32.7
DE	DESY-ZN	26759	0.977	1115	0.50	157.0	2231	11.3	56.7
UK	UKI-LT2-RHUL	30258	0.702	505	0.23	185.0	2196	6.2	36.8
IT	INFN-T1	32878	0.908	1886	1.00	-1.0	1886	6.5	33.2

Backup: all sites (2)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▲	Hz	CPU/W
IT	INFN-NAPOLI-ATLAS	12899	0.949	527	0.30	156.0	1757	7.5	37.2
DE	LRZ-LMU	15814	0.711	278	0.16	146.0	1740	8.7	40.1
DE	MPPMU	10641	0.553	270	0.16	89.0	1689	6.8	34.6
UK	UKI-SOUTHGRID-RALPP	9565	0.862	418	0.25	76.0	1675	4.5	24.3
ES	UAM-LCG2	10179	0.888	393	0.25	126.0	1572	8.7	46.1
TW	TW-FTT	19085	0.803	756	0.50	115.0	1512	17.6	49.6
FR	TOKYO-LCG2	36470	0.986	1501	1.00	293.0	1501	11.0	46.6
FR	IN2P3-LPSC	2219	0.723	73	0.05	29.0	1463	5.8	21.6
UK	UKI-SOUTHGRID-BHAM-HEP	3249	0.948	72	0.05	35.0	1440	9.7	32.9
DE	DESY-HH	18723	0.942	681	0.50	180.0	1362	9.3	43.9
NL	TR-10-ULAKBIM	8616	0.866	230	0.20	80.0	1153	4.1	32.5
NL	JINR-LCG2	11110	0.468	221	0.20	63.0	1107	8.6	29.1
UK	UKI-NORTHGRID-SHEF-HEP	6227	0.821	102	0.10	36.0	1020	6.6	50.7
US	SLACXRD	14716	0.889	337	0.34	215.0	993	4.0	64.8
IT	INFN-ROMA1	7848	0.864	270	0.30	131.0	901	6.3	25.6
NL	WEIZMANN-LCG2	17806	0.115	41	0.05	115.0	830	7.7	38.4
FR	GRIF-SACLAY	8672	0.935	213	0.26	63.0	822	8.9	45.2
UK	UKI-LT2-QMUL	8440	0.276	39	0.05	123.0	797	3.7	29.4
DE	UNI-FREIBURG	8805	0.604	126	0.17	106.0	743	7.7	45.5
CA	ALBERTA-LCG2	5370	0.884	185	0.25	24.0	740	4.8	28.3
IT	INFN-FRASCATI	2748	0.877	71	0.10	38.0	715	4.3	29.5
ES	IFAE	7588	0.928	161	0.25	68.0	647	5.7	29.7
UK	UKI-NORTHGRID-MAN-HEP1	7948	0.633	83	0.15	40.0	559	3.8	45.8

Backup: all sites (3)

Cloud	Site	# Jobs	Eff.	evts	share	Size(TB)	evts/share ▲	Hz	CPU/W
UK	UKI-NORTHGRID-MAN-HEP2	9347	0.382	71	0.15	40.0	479	2.8	32.0
FR	GRIF-LPNHE	10312	0.857	217	0.50	140.0	435	7.9	34.2
DE	FZK-LCG2	16286	0.681	424	1.00	-1.0	424	4.1	23.0
US	SWT2	8053	0.788	134	0.33	40.0	406	2.0	16.7
CA	SFU-LCG2	10832	0.950	270	0.72	157.0	375	5.8	26.6
FR	RO-02-NIPNE	2810	0.432	37	0.10	51.0	373	2.9	19.2
ES	IFIC-LCG2	2372	0.753	87	0.25	68.0	350	12.4	63.8
NL	ITEP	864	0.991	34	0.10	10.0	347	7.5	38.5
NL	CSTCDIE	1368	0.379	16	0.05	39.0	331	7.3	35.4
UK	UKI-NORTHGRID-LANCS-HEP	6819	0.565	78	0.30	140.0	260	4.8	24.5
FR	IN2P3-CC	10486	0.761	212	1.00	-1.0	212	3.1	26.7
DE	HEPHY-UIBK	2056	0.467	19	0.10	36.0	193	4.5	19.5
NG	NDGF-T1	16730	0.870	172	1.00	-1.0	172	7.3	
US	BNL-OSG2	10809	0.911	157	1.00	-1.0	157	7.1	34.5
NL	NIKHEF-ELPROD	5174	0.875	143	1.00	-1.0	143	7.5	39.6
IT	INFN-MILANO	2790	0.612	42	0.30	78.0	140	1.7	19.2
NL	RRC-KI	5339	0.402	67	0.50	171.0	135	2.9	18.3
TW	AUSTRALIA-ATLAS	4548	0.891	67	0.50	119.0	135	6.0	42.4
DE	WUPPERTALPROD	5783	0.210	19	0.17	63.0	117	4.7	29.7
DE	PRAGUELCG2	870	0.294	3	0.05	23.0	65	0.7	8.9
FR	RO-07-NIPNE	1516	0.157	5	0.10	33.0	59	5.7	27.0
NL	SARA-MATRIX	1461	0.763	41	1.00	-1.0	41	3.3	37.2
UK	UKI-LT2-UCL-CENTRAL	473	0.222	3	0.10	16.0	35	3.8	28.3
NL	RU-PNPI	204	0.632	4	0.15	65.0	28	0.8	11.1
ES	LIP-COIMBRA	1134	0.160	9	0.50	115.0	19	6.4	22.1