



May 2006, Jamie Shiers

Establishing T0-T1 Transfers as Production Services

Recovery from Scheduled Down-Times

Executive Summary

This document addresses the following recommendations from the Review of Tier0 and Tier1 Site Monitoring and Operation in Service Challenge 4.

10. Recovery from scheduled downtimes of individual Tier1 sites for both short (~4 hour) and long (~48 hour) interventions at full nominal data rates needs to be demonstrated. Recovery from scheduled downtimes of the Tier0 – and thus affecting transfers to all Tier1s – up to a minimum of 8 hours must also be demonstrated. A plan for demonstrating this capability should be developed in the Service Coordination meeting before the end of May.
11. Continuous low-priority transfers between the Tier0 and Tier1s must take place to exercise the service permanently and to iron out the remaining service issues. These transfers need to be run as part of the service, with production-level monitoring, alarms and procedures, and not as a “special effort” by individuals.

Demonstrating that these transfers can be performed as stable, reliable production services is considered a pre-requisite and hence this issue is addressed first.

Tier0-Tier1 Transfers as a Production Service

As previously discussed, disk-disk transfers will restart as soon as individual Tier1 sites have upgraded to gLite 3.0. Contrary to previous “challenges”, these will only be monitored (by a human) during normal working hours from the CERN end. Problems seen with transfers to individual sites will be reported through GGUS and logged, together with the date and time of problem resolution. Average daily rates to each site will be published on the Service Challenge Wiki and compared with the nominal rates for that site. (Red – below 85% of nominal; yellow – up to 99%; green – 100% and more).

It is particularly important that sites with a slow ramp-up during the April SC4 tests – as shown in table 2 below – are able to demonstrate rapid ramp-up. By definition, this always concerns both the source site – CERN – and the destination. This concerns NDGF, ASGC, CNAF and FZK in particular, but in all cases the target should be ramp-up to full nominal rates in a matter of hours, and not days.

It is important that sites that have not yet implemented basic monitoring and alarms for the infrastructure (box-level monitoring etc.) and services (SRM ok?) – which from the SC4 experience in April appears to be the case – do so rapidly, so that any issues are spotted and resolved in a timely manner.

These issues will be followed through the weekly operations meeting, with daily follow-up via e-mail as required.



Recovery from Site Interruptions

Once Tier0-Tier1 transfers are running as stable, production services, demonstration of recovery from short and long interruptions (4 and 48 hours respectively) of individual Tier1 sites, as well as up to a minimum of 8 hours for the Tier0, can commence.

Given the schedule for experiment-driven transfers, these cannot realistically be scheduled prior to mid-July.

As stable services are a pre-requisite, this issue will be addressed further once these have been established.

Site	Ramp-up	Stability	Weekly Report	Weekly Meeting	Monitoring / Operations	Interventions	Average
CERN	2-3	2	3	1	2	1	2
ASGC	4	4	2	3	4	3	3
TRIUMF	1	1	4	2	1-2	1	2
FNAL	2	3	4	1	2	3	2.5
BNL	2	1-2	4	1	2	2	2
NDGF	4	4	4	4	4	2	3.5
PIC	2	3	3	1	4	3	3
RAL	2	2	1-2	1	2	2	2
SARA	2	2	3	2	3	3	2.5
CNAF	3	3	1	2	3	3	2.5
IN2P3	2	2	4	2	2	2	2.5
FZK	3	3	2	2	3	3	3

Table 1 - Summary of Site Production Readiness from SC4 Disk-Disk Throughput Phase

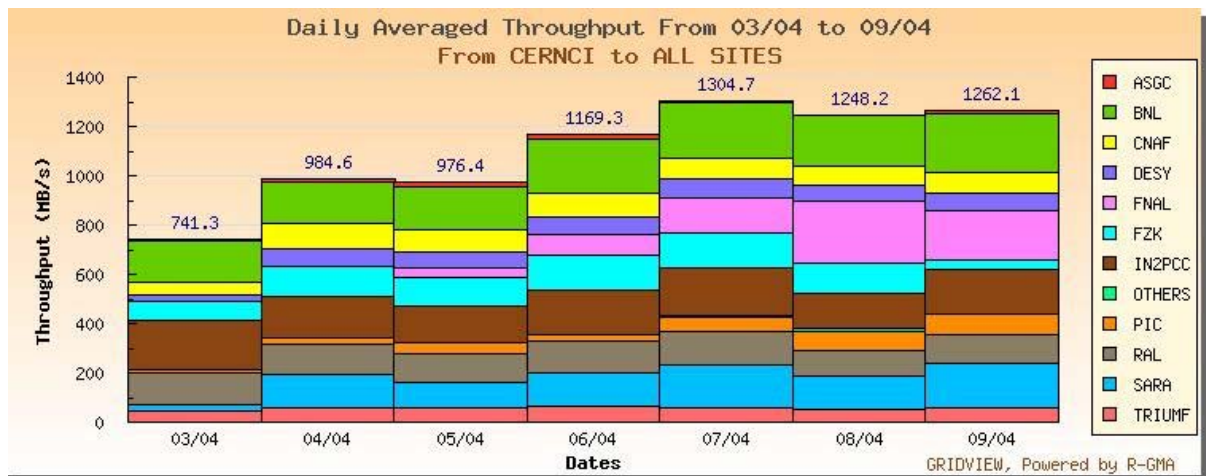


Figure 1 - Data Transfer Rates during Week 1 of S4 Disk - Disk Throughput Tests



May 2006, Jamie Shiers




Figure 2 - Data Transfer Rates during Week 1 of SC4 Disk - Disk Throughput Tests



Site/Date	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Av. (Nom.)
ASGC	0	7	23	23	0	0	12	22	33	25	26	21	19	22	17(100)
TRIUMF	44	42	55	62	56	55	61	62	69	63	63	60	60	62	58(50)
FNAL	0	0	38	80	145	247	198	168	289	224	159	218	269	258	164(200)
BNL	170	103	173	218	227	205	239	220	199	204	168	122	139	284	191(200)
NDGF	0	0	0	0	0	14	0	0	0	0	14	38	32	35	10(50)
PIC	0	18	41	22	58	75	80	49	0	24	72	76	75	84	48(100¹)
RAL	129	86	117	128	137	109	117	137	124	106	142	139	131	151	125(150)
SARA	30	78	106	140	176	130	179	173	158	135	190	170	175	206	146(150)
CNAF	55	71	92	95	83	80	81	82	121	96	123	77	44	132	88(200)
IN2P3	200	114	148	179	193	137	182	86	133	157	183	193	167	166	160(200)
FZK	81	80	118	142	140	127	38	97	174	141	159	152	144	139	124(200)

Table 2 - Summary of Achieved Transfer Rates (MB/s)

Appendix

[1] The Service Challenge 4 ‘blog’ -

<https://twiki.cern.ch/twiki/bin/view/LCG/ServiceChallengeFourBlog>.

[2] Service Challenge 4 Disk-Disk transfer results -

<https://twiki.cern.ch/twiki/bin/view/LCG/AprilDiskDiskTransferTargetsAndStatus>.

[3] The GridPP wiki - https://wiki.gridpp.ac.uk/wiki/Main_Page.

[4] TRIUMF SC3 disk – tape status page:

<http://grid.triumf.ca/status/sc4/sc4-disktape.html>.

[5] BNL tape plots -

http://www.atlasgrid.bnl.gov/dcache_tapewrite_monitoring/plots/.

[6] IN2P3 tape plots - <http://netstat.in2p3.fr/weathermap/graphiques/lcgmss.html>.

[7] IN2P3 preliminary analysis of SC4 throughput (Lionel Schwarz) -

<http://agenda.cern.ch/fullAgenda.php?ida=a057189>.

[8] FNAL tape plots -

<http://cmsdcam.fnal.gov:8090/dcache/outplot?lvl=1&filename=billing-2006.04.daily.bwrsc3.png>.

[9] Scheduling of Service Interruptions at WLCG Sites -

<http://agenda.cern.ch/askArchive.php?base=agenda&categ=a061501&id=a061501s0t11%2Fmoreinfo%2FSC4-scheduled-maintenance-May16.pdf>. Can also be found via <https://twiki.cern.ch/twiki/bin/view/LCG/TalksAndDocuments>.

¹ The agreed target for PIC is 60MB/s, pending the availability of their 10Gb/s link to CERN.