

Summary of SC4 Disk-Disk Transfers

LCG MB, April 18 2006

Jamie Shiers, CERN

Goals of SC4 Disk-Disk Transfers

- Disk - disk Tier0-Tier1 tests at the full nominal rate are scheduled for April.
- The proposed schedule is as follows:
 - April 3rd (Monday) - April 13th (Thursday before Easter) - sustain an average daily rate to each Tier1 at or above the full nominal rate. (This is the week of the GDB + HEPiX + LHC OPN meeting in Rome...)
- Any loss of average rate $\geq 10\%$ needs to be:
 - accounted for (*e.g. explanation / resolution in the operations log*)
 - compensated for by a corresponding increase in rate in the following days
- We should continue to run at the same rates unattended over Easter weekend (14 - 16 April).

Achieved (Nominal) pp data rates (SC3++)

Centre	ALICE	ATLAS	CMS	LHCb	Rate into T1 (pp) Disk-Disk (SRM) rates in MB/s
ASGC, Taipei	-	✓	✓	-	80 (100) (have hit 140)
CNAF, Italy	✓	✓	✓	✓	200
PIC, Spain	-	✓	✓	✓	>30 (100) (network constraints)
IN2P3, Lyon	✓	✓	✓	✓	200
GridKA, Germany	✓	✓	✓	✓	200
RAL, UK	-	✓	✓	✓	200 (150)
BNL, USA	-	✓	-	-	150 (200)
FNAL, USA	-	-	✓	-	>200 (200)
TRIUMF, Canada	-	✓	-	-	140 (50)
SARA, NL	✓	✓	-	✓	250 (150)
Nordic Data Grid Facility	✓	✓	-	-	150 (50)

☺ Meeting or **exceeding** nominal rate (disk - disk)

☺ Met target rate for SC3 (disk & tape) re-run

Missing: rock solid stability - nominal tape rates

SC4 T0-T1 throughput goals: nominal rates to disk (April) and tape (July)

(Still) To come:
 Srm copy support in FTS;
 CASTOR2 at remote sites;
 SLC4 at CERN;
 Network upgrades etc.

All Wikis

- ACPP
- ADCgroup
- AISgroup
- ALICE
- ALPHA
- AliceSPD
- AthenaFCalTBAna
- Atlas
- CERNSearch
- CMS
- CS
- Controls
- DESgroup
- Dbaservices
- DefaultWeb
- EGEE
- ELFms
- ETICS
- EgeePtf
- FIOgroup
- HCC
- HRonDemand
- Know
- Lar
- LCG
- LCGAAWorkbook
- LHCatHome
- LHCOPN
- LHCb
- LHCgas
- LcgProcurementInfo
- LinuxSupport
- Main
- PHSS
- PSSGroup
- Plugins
- SPI
- SRMDev
- Sandbox
- SupComp05
- TWiki

TWikiGuest

Week two (April 10 on)

- Week 2 average to sum of Tier1 sites is 1262 MB/s - 79% of the target.

Site	Disk-Disk	Week1 Average	Week2 Average	Apr10	Apr11	April12	April13	April14	April15	April16	April17
TRIUMF	50	54	63	62	69	63	63	60	60	62	63
BNL	200	191	199	220	199	204	168	122	139	284	257
FNAL	200	101	231	168	289	224	159	218	269	258	261
PIC	60	49	78 (5 days)	49	-	24	72	76	75	84	82
RAL	150	118	136	137	124	106	142	139	131	151	160
SARA	150	120	178	173	158	135	190	170	175	206	213
IN2P3	200	165	157	86	133	157	183	193	167	166	167
FZK	200	104	142	97	174	141	159	152	144	139	130
CNAF	200	80	88	82	121	96	123	77	44	132	32
ASGC	100		24	22	33	25	26	21	19	22	24
NDGF	50		28 (5 days)	-	-	-	14	38	32	35	20
DESY	60	70	74	71	77	69	72	76	73	76	76
TOTAL (T1s)	1600			1096	1300	1175	1046	1266	1255	1539	1409


Week one (April 3 on)

Site	Disk-Disk	Apr3	Apr4	Apr5	Apr6	Apr7	Apr8	Apr9	Weekly average	Average from startup	Target
TRIUMF	50	44	42	55	62	56	55	61	54	54 (>100%)	50
BNL	200	170	103	173	218	227	205	239	191	191 (>95%)	200
FNAL	200	-	-	38	80	145	247	198	101	141 (>70%)	200
PIC	60	-	18	41	22	58	75	80	49	42 (70%)	60
RAL	150	129	86	117	128	137	109	117	118	118 (~80%)	150
SARA	150	30	78	106	140	176	130	179	120	120 (80%)	150
IN2P3	200	200	114	148	179	193	137	182	165	165 (>80%)	200
FZK	200	81	80	118	142	140	127	38	104	104	200
CNAF	200	55	71	92	95	83	80	81	80	80	200
ASGC	100	-	7	23	23	-	-	12			100
NDGF	50	-	-	-	-	-	14	-			50
DESY	60	-	68	63	75	74	68	74		70	60
TOTAL (T1s)	1600	709	599	911	1089	1215	1179	1187		984 (61.5% of target)	


SC4 Blog

- 18/04 01:30 Easter Monday has been above 1.4 GB/s for most of the day, averaging about 1.5 GB/s, peaking at 1.6 GB/s right after the CNAF channel was switched on again. The problems of the day were with the CNAF channel, which was off except for a 6-hour spell, and with the BNL channel, which experienced many errors that were not very much reflected in the rate until 22:00 GMT. NDGF experimenting with dCache parameters. *Maarten*
- 17/04 02:30 Easter Sunday was the first full day averaging 1.6 GB/s! All channels were stable except for CNAF, whose LSF queue got full with stuck jobs a few times, requiring admin interventions. *Maarten*
- 16/04 03:50 A rocky night with...


Effect of FTS DB Cleanup



24 hours since DB Cleanup



Easter Sunday: > 1.6GB/s including DESY



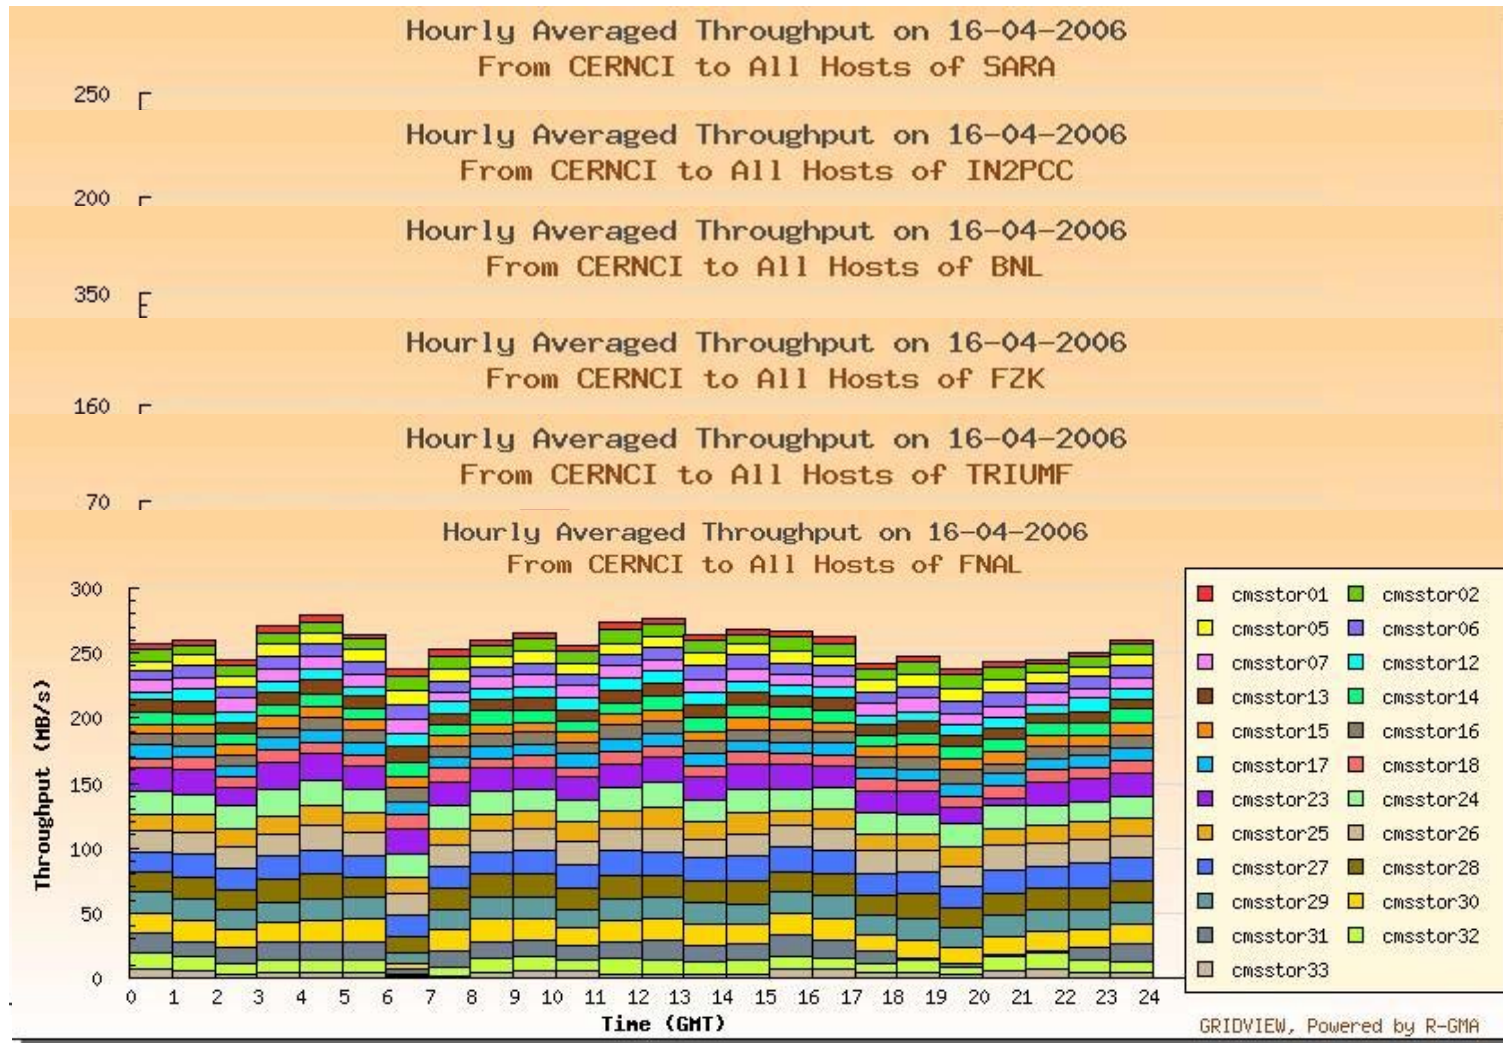
GridView reports 1614.5MB/s as daily average for 16-04/2006

The LHC Computing Grid – (The Worldwide LCG)

- Hourly Averaged Throughput on 16-04-2006
From CERNCI to ASGC
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to BNL
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to CNAF
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to FNAL
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to FZK
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to IN2PCC
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to NDGF
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to PIC
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to RAL
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to SARA
- Hourly Averaged Throughput on 16-04-2006
From CERNCI to TRIUMF



Site by Site Detail



Next targets (backlog)

Centre	ALICE	ATLAS	CMS	LHCb	Rate into T1 (pp) Disk-Disk (SRM) rates in MB/s
ASGC, Taipei	-	✓	✓	-	150
CNAF, Italy	✓	✓	✓	✓	300
PIC, Spain	-	✓	✓	✓	150
IN2P3, Lyon	✓	✓	✓	✓	300
GridKA, Germany	✓	✓	✓	✓	300
RAL, UK	-	✓	✓	✓	225
BNL, USA	-	✓	-	-	300
FNAL, USA	-	-	✓	-	300
TRIUMF, Canada	-	✓	-	-	75
SARA, NL	✓	✓	-	✓	225
Nordic Data Grid Facility	✓	✓	-	-	75

Need to vary some key parameters (filesize, number of streams etc) to find sweet spot / plateau.
Needs to be consistent with actual transfers during data taking

Summary

- We did not sustain a daily average of 1.6MB/s out of CERN nor the full nominal rates to all Tier1s for the period
 - Just under 80% of target in week 2
- Things clearly improved --- both since SC3 and during SC4:
 - Some sites meeting the targets!
 - Some sites 'within spitting distance' - optimisations? Bug-fixes?
 - See blog for CNAF castor2 issues for example...
 - Some sites still with a way to go...
- "Operations" of Service Challenges still very heavy
 - Special thanks to Maarten Litmaath for working > double shifts...
 - Need more rigour in announcing / handling problems, site reports, convergence with standard operations etc.
- Remember the "Carminati Maxim":

"If its not there for SC4, it won't be for the production service in October..."

(and vice-versa)

Postscript

- Q: Well, 80% of target is pretty good isn't it?
- A: Yes and no.
- We are still testing a simple (in theory) case.
- How will this work under realistic conditions, including other concurrent (and directly related) activities at the T0 and T1s?
 - See Bernd's tests...
- We need to be running comfortably at 2-2.5GB/s day-in, day-out and add complexity step by step as things become understood and stable.
- And anything requiring >16 hours of attention a day is not going to work in the long term...