

Service Challenge 4 - Ramping up to Full Production Data Rates

Target: September 2006

LHCC Referees Meeting, May 2006
[Updated with ALICE & LHCb TO-T1 transfers]

Jamie Shiers, CERN

Agenda

- a. Status of the Tier-0 to Tier-1 throughput tests

- b. Progress towards the Level-1 milestone to achieve nominal LHC rates to tape by end September

[Detail in hidden slides]

SC4 - Executive Summary

We have shown that we can drive transfers at full nominal rates to:

- Most sites simultaneously;
- All sites in groups (modulo network constraints - PIC);
- At the target nominal rate of 1.6GB/s expected in pp running

In addition, several sites exceeded the disk - tape transfer targets

- There is no reason to believe that we cannot drive all sites at or above nominal rates for sustained periods.

But

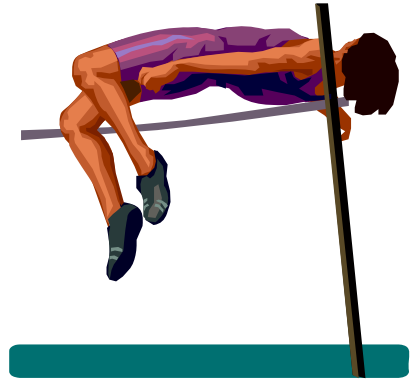
- There are still major operational issues to resolve - and most importantly - a full end-to-end demo under realistic conditions

Heat

Nominal Tier0 - Tier1 Data Rates (pp)

<i>Tier1 Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Target</i>
IN2P3, Lyon	9%	13%	10%	27%	200
GridKA, Germany	20%	10%	8%	10%	200
CNAF, Italy	7%	7%	13%	11%	200
FNAL, USA	-	-	28%	-	200
BNL, USA	-	22%	-	-	200
RAL, UK	-	7%	3%	15%	150
NIKHEF, NL	(3%)	13%	-	23%	150
ASGC, Taipei	-	8%	10%	-	100
PIC, Spain	-	4% (5)	6% (5)	6.5%	100
Nordic Data Grid Facility	6%	6%	-	-	50
TRIUMF, Canada	-	4%	-	-	50
TOTAL					1600MB/s

A Brief History...



- SC1 - December 2004: did **not** meet its goals of:
 - Stable running for ~2 weeks with 3 named Tier1 sites...
 - But more sites took part than foreseen...
- SC2 - April 2005: met **throughput** goals, but still
 - No reliable file transfer service (or real services in general...)
 - Very limited functionality / complexity
- SC3 "classic" - July 2005: added several components and **raised** bar
 - SRM interface to storage at all sites;
 - Reliable file transfer service using gLite FTS;
 - Disk - disk targets of 100MB/s per site; 60MB/s to tape
 - Numerous issues seen - investigated and debugged over many months
- SC3 "Casablanca edition" - Jan / Feb re-run
 - Showed that we had resolved many of the issues seen in July 2005
 - Network bottleneck at CERN, but most sites at or above targets
 - Good step towards SC4(?)

SC4 Schedule

- Disk - disk Tier0-Tier1 tests at the full nominal rate are scheduled for April. (from weekly con-call minutes...)
- 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)

Excellent report produced by IN2P3, covering disk and tape transfers, together with analysis of issues.

Successful demonstration of both disk and tape targets.

- Dropped based on experience of first week of disk - disk tests

Tier1 - Tier1 & Tier1 - Tier2 Transfers

- Tier1 - Tier1 transfers: ATLAS ESD mirroring; distribution of AOD and TAG datasets
- Tier1 - Tier2 transfers: MC archiving, analysis data download
- WLCG Q2 2006 Milestone - May:
 - All T1 sites to define channels to all other T1s and supported T2s and demonstrate functionality of transfers between sites.

➤ Some sites have established - and tested - these 'FTS channels', (e.g. GridPP, Spain, ... example for others?) but the process is long...

¿ Q: who should organise these? Tier1s? Experiments? Both?

Functionality

Performance

Meeting the LCG challenge

Example: Tier-2 individual transfer tests

Initial focus was on getting SRMs understood and deployed....


	Receiving											
	RAL Tier-1	Lancaster	Manchester	Edinburgh	Glasgow	Birmingham	Oxford	Cam	Durham	QMUL	IC-HEP	RAL-PPD
RAL Tier-1		-800Mb/s	350Mb/s	156Mb/s	166 Mb/s	289 Mb/s	252 Mb/s			118 Mb/s	84Mb/s	397 Mb/s
Lancaster												
Manchester	150 Mb/s											
Edinburgh	440Mb/s											
Glasgow	331Mb/s											
Birmingham	461 Mb/s											
IC-HEP												
Oxford	456 Mb/s											
Cambridge	74 Mb/s											
Durham	193 Mb/s											
QMUL	172 Mb/s											
IC-HEP												
RAL-PPD	388 Mb/s											

- Big variation in what sites could achieve
 - Internal networking configuration issues
 - Site connectivity (and contention)
 - SRM setup and level of optimisation
- Rates to RAL were generally better than from RAL
 - Availability and setup of gridFTP servers at Tier-2s
 - SRM setup and level of optimisation
- Scheduling tests was not straightforward
 - Availability of local site staff
 - Status of hardware deployment
 - Availability of Tier-1
 - Need to avoid first tests during certain periods (local impacts)

Example rates from throughput tests

Meeting the LCG challenge

Example: Tier-1 & Tier-2 combined transfer tests

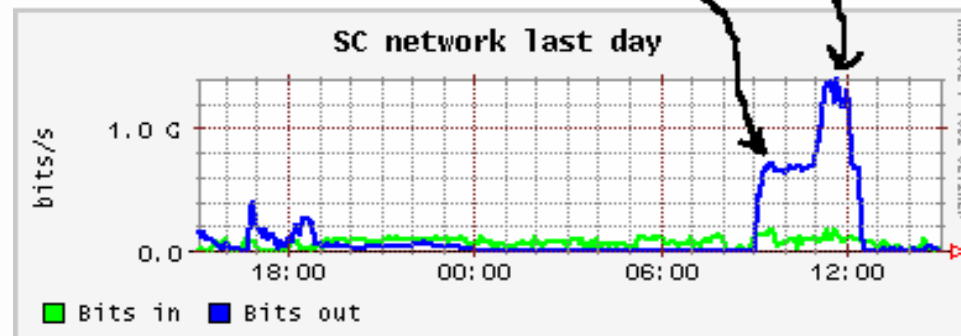


- Early attempts revealed unexplained dropouts
- Dropouts later traced to firewall
- A rate cap at RAL was introduced for later tests

- Tests repeated to check RAL capping
- Rate was stretched further by using an OPN link to Lancaster

Lancaster OPN traffic

Steady transfer at 700Mb/s
no dropouts.






GridPP

UK Computing for Particle Physics


Meeting the LCG challenge

Tier-1 & Tier-2 combined transfer tests-rerun

48 Hour Aggregate Test from RAL to Tier2s over SJ4 and UKLIGHT



48 Hour Aggregate Test to RAL from Tier2s over SJ4 and UKLIGHT



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

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

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

Easter Sunday: > 1.6GB/s including DESY



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

SC4 Disk - Disk 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

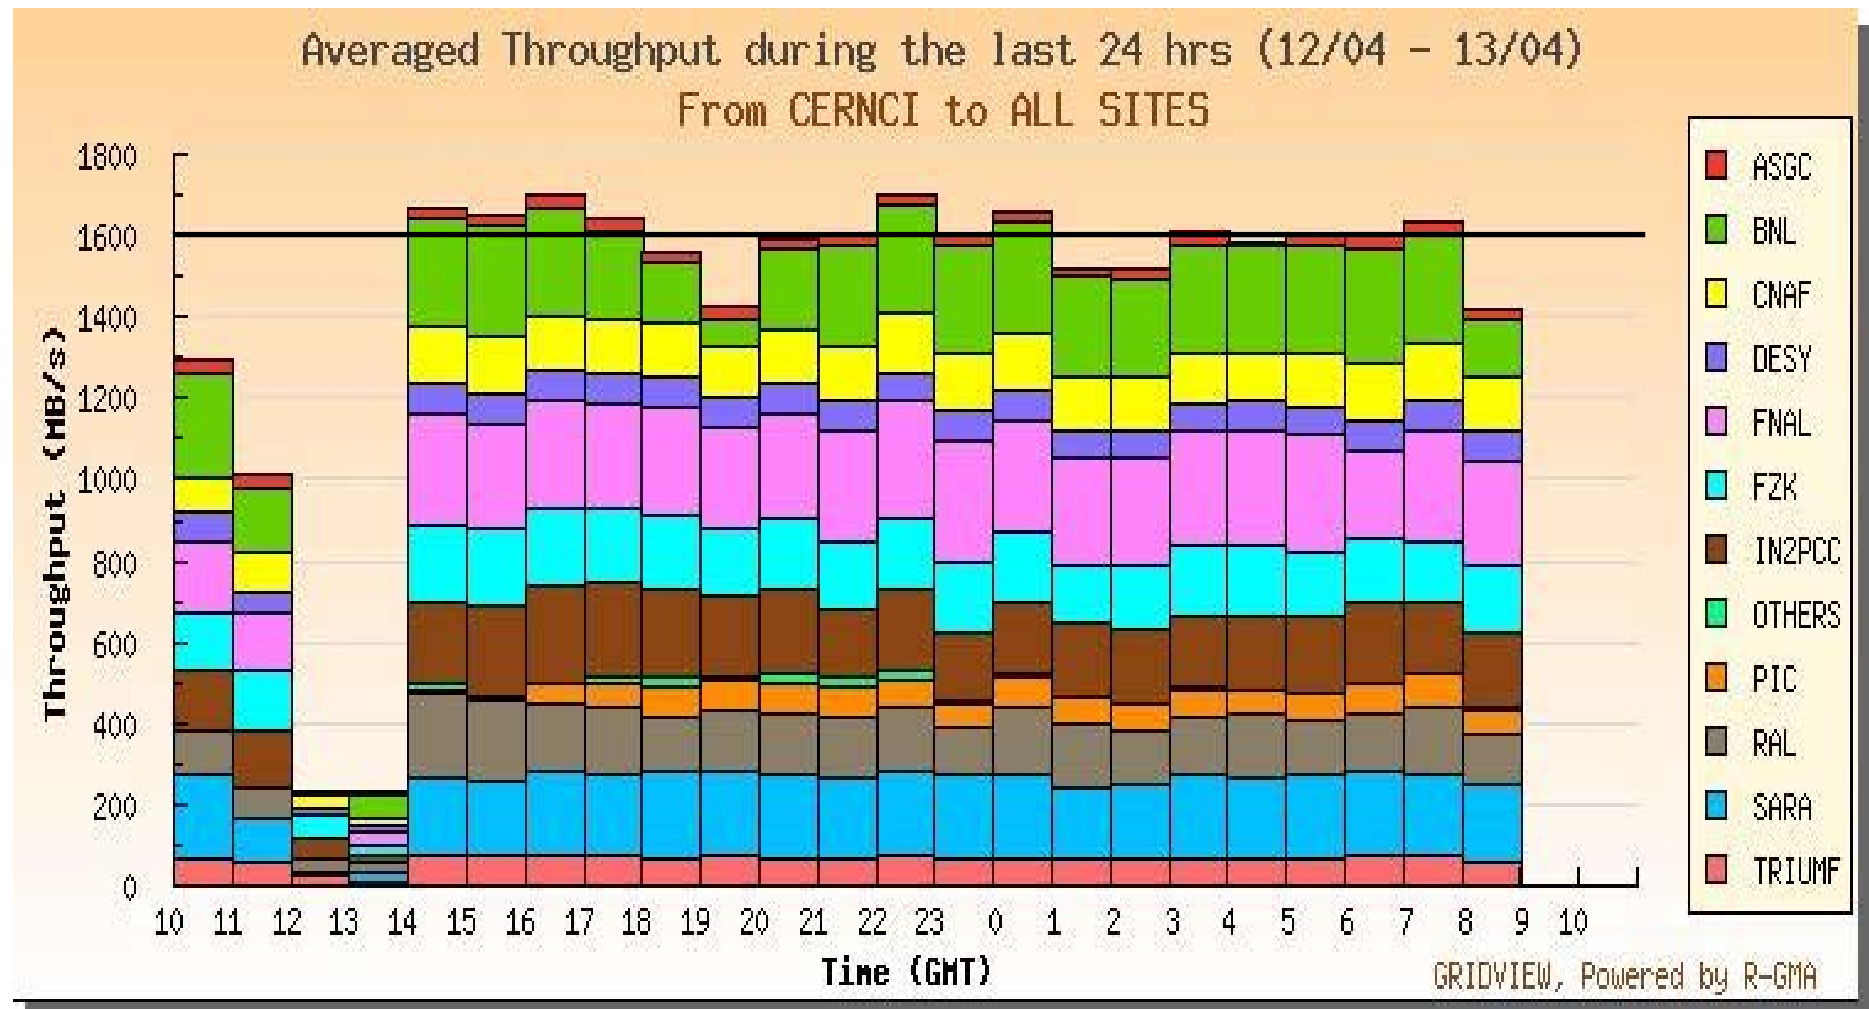
- 😊 **Things clearly improved --- both since SC3 and during SC4:**
 - Some sites meeting the targets!
 - Some sites 'within spitting distance' - optimisations? Bug-fixes?
 - **See SC4blog for examples of these issues and progress**
 - Some sites still with a way to go...

- **Bottleneck due to size of FTS tables and consequent query time (hidden slides)**
 - Outstanding action for some time to implement 'partitioning'
 - Manual DB clean-up had clear effect - periodic cleanup now implemented

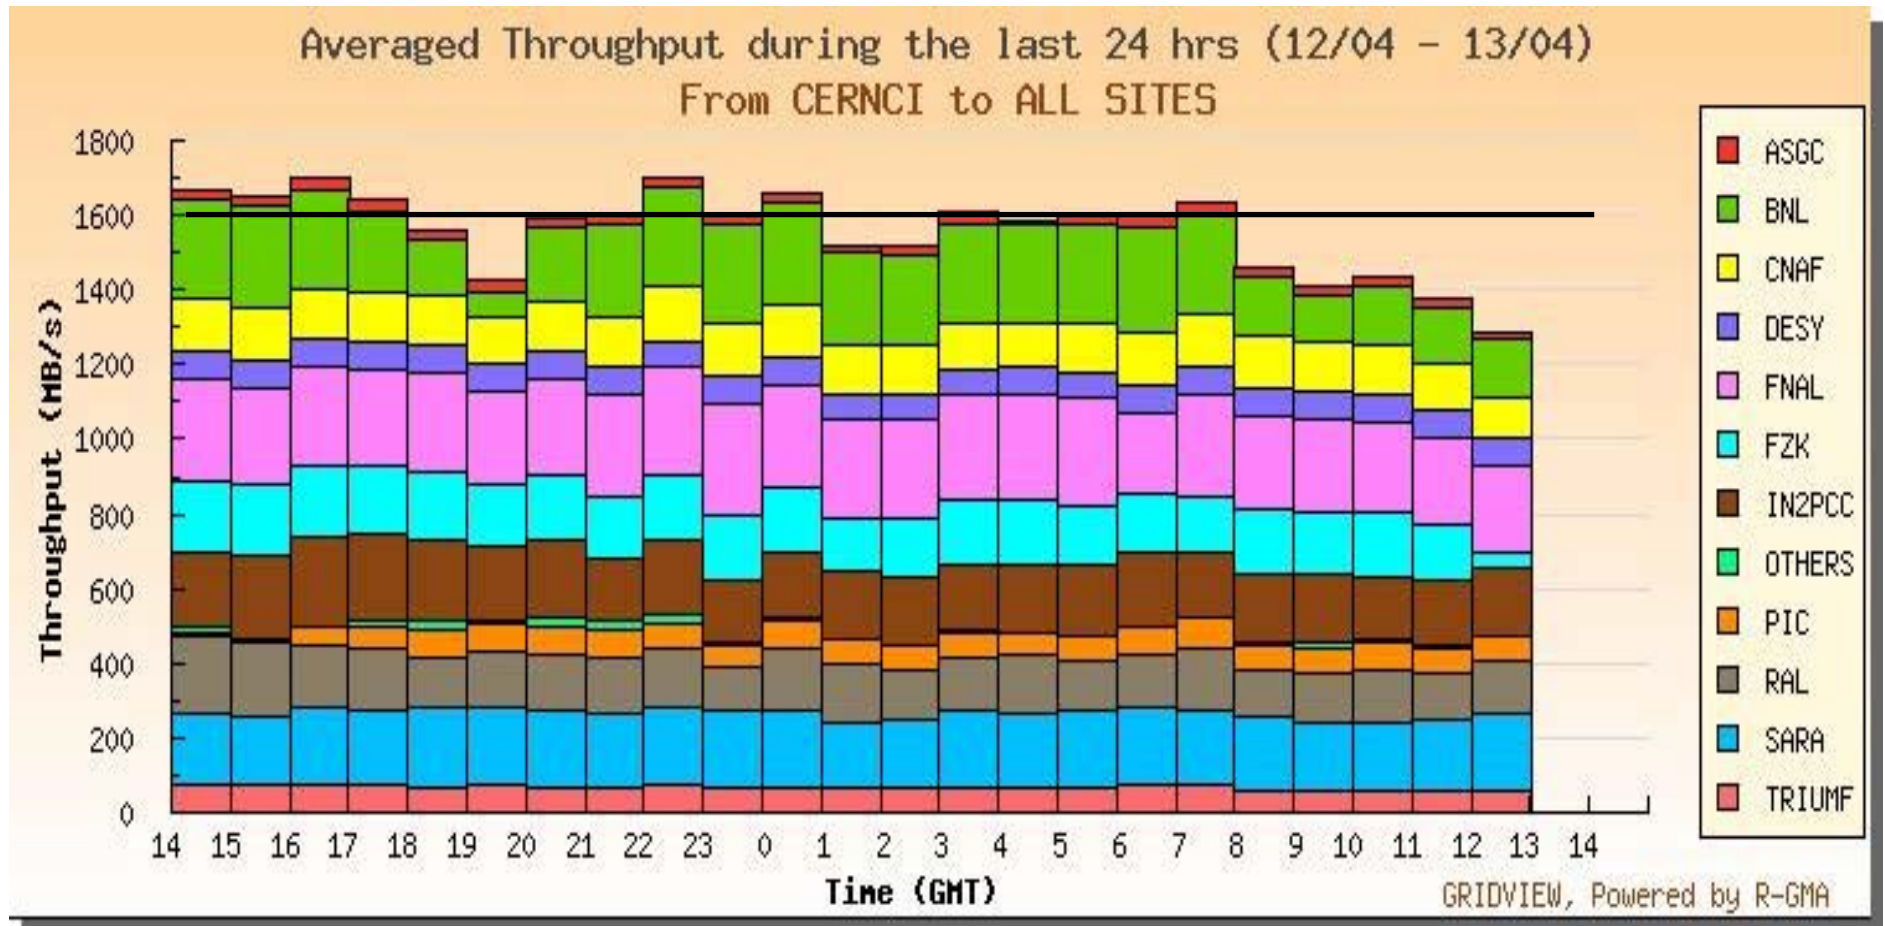
- **Other site by site tuning required - more hidden bottlenecks?**

- 💣 **"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.**

Effect of FTS DB Cleanup



24 hours since DB Cleanup



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](#)
[PHESS](#)
[PSSGroup](#)
[Plugins](#)
[SPI](#)
[SRMDev](#)
[Sandbox](#)
[SupComp05](#)
[TWiki](#)

[TWikiGuest](#)

Week two (April 10 on)

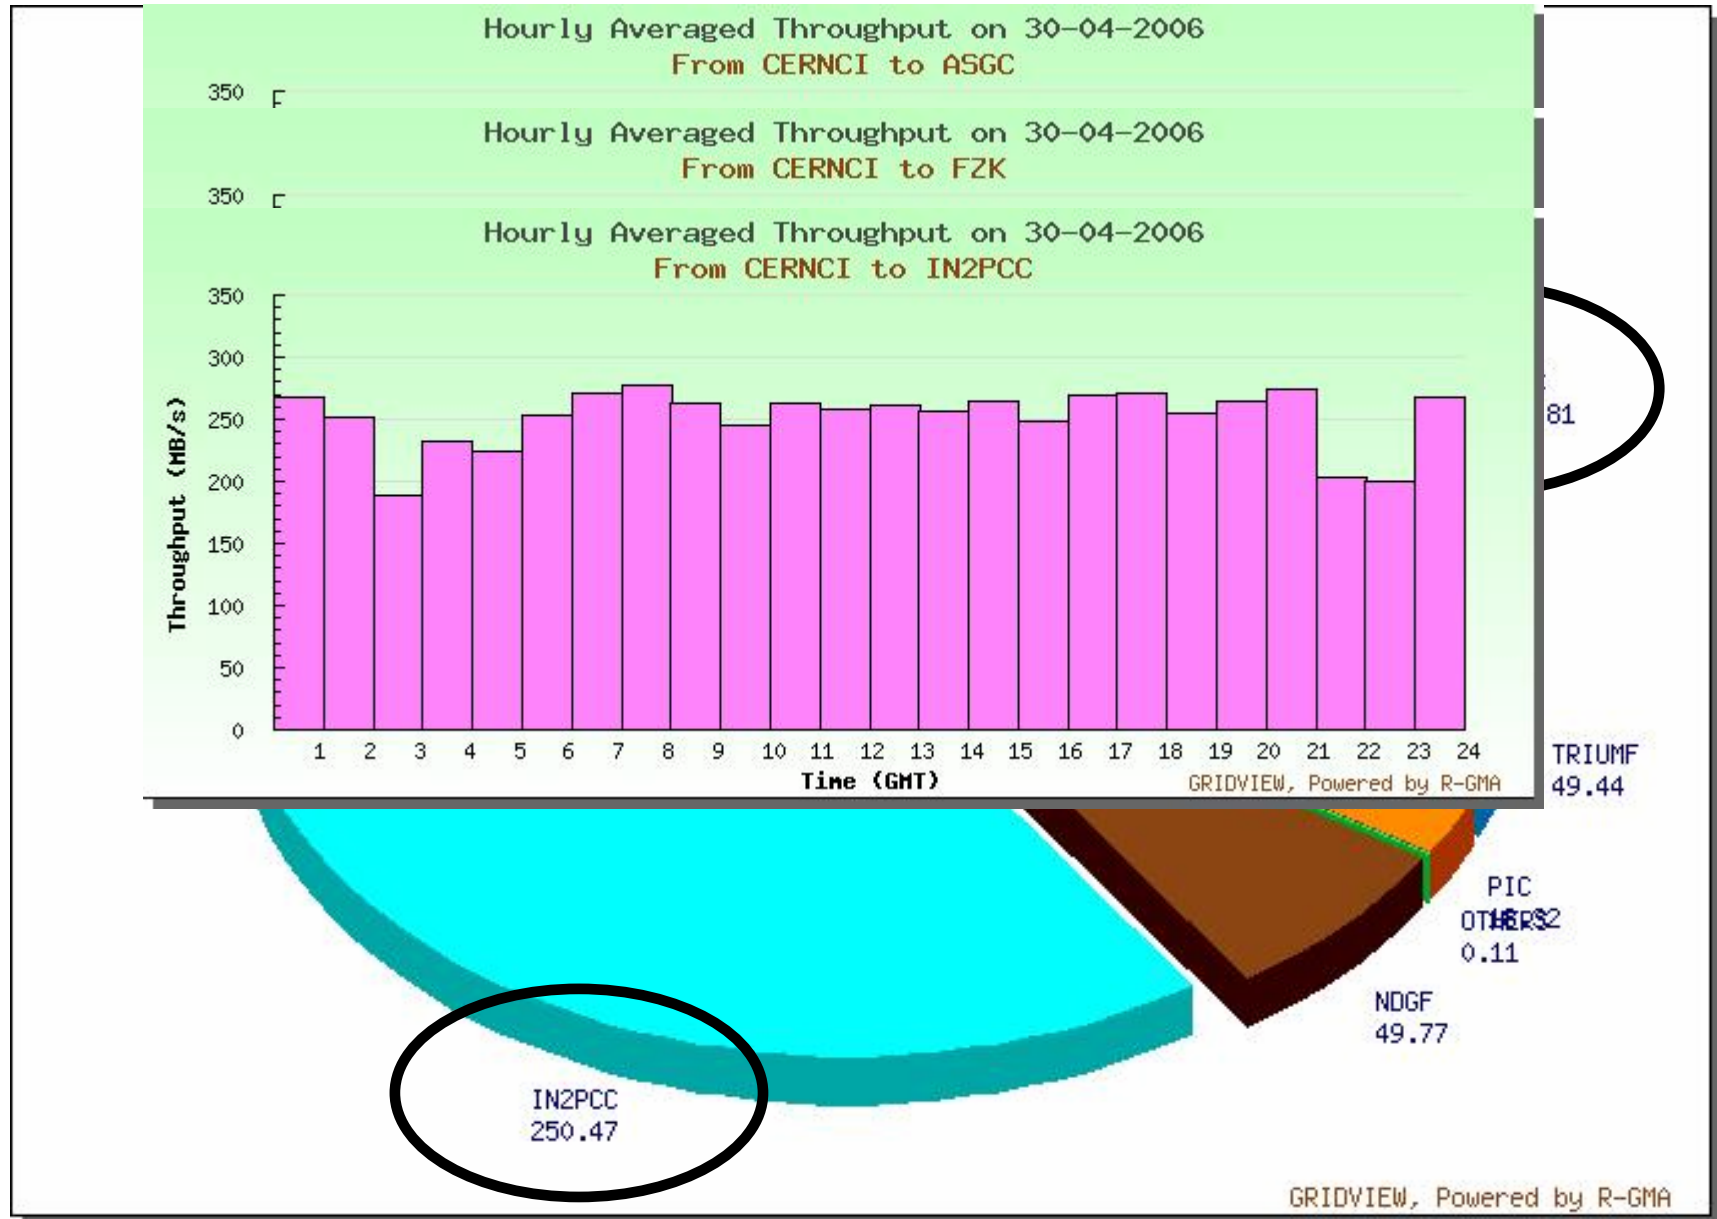
- Week 2 average to sum of Tier1 sites is 1262 MB/s - 79% of the target.

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)	

Site by Site Debugging




Site by Site Debugging

Daily Averaged Throughput (in MB/s) on 01/05
From CERNCI to ALL SITES



CNAF Disk-Disk Re-run

- May 2: completed the upgrade of Castor2 to version 2.0.4-0;
- May 3: execution of local write tests and also remote transfers (but with few concurrent file transfers). Results showed a good local Castor 2 to LSF interaction. Power down problem in the Tier 1



Castor2-specific problem) tests with 50 concurrent files: 1800 Mb/s (225 MB/s)

- Stable run at 223 MB/s for approximately 1 and a half day (until May 6, 8 pm). Network traffic statistics graph is attached.

Concerns - April 25 MB

- Site maintenance and support coverage during throughput tests
 - After **5 attempts**, have to assume that this will not change in immediate future - better design and build the system to handle this
 - (This applies also to CERN)
- **Unplanned schedule changes, e.g. FZK missed disk - tape tests**
 - Some tests since last Friday
- Monitoring, showing the data rate to tape at remote sites and also of overall status of transfers
- Debugging of rates to specific sites [which has been done...]
- **Future throughput tests using more realistic scenarios**

SC4 Blog

May 2006

- 02/05 00:30 ASGC had a 1-hour dip with many SRM timeouts, otherwise doing 100 MB/s or better. BNL were doing 90 MB/s, then ran out of tape and decided to switch the channel off for the time being, given that the first disk and tape phases of SC4 have ended. FZK had a 1-hour dip to 120 MB/s during the night, a few dips to about 200 MB/s, running at about 240 MB/s most of the time. IN2P3 doing 250 MB/s or better most of the time. NDGF dropped to zero during the night due to no write pool being available, then came back to a steady 60 MB/s. PIC still at 20 MB/s with many SRM timeouts. TRIUMF stable at 50 MB/s. *Maarten*

April 2006

- 01/05 02:20 ASGC OK at 120 MB/s. BNL stable at 90 MB/s. DESY at 70 MB/s, then set inactive at 19 GMT in preparation of high-speed transfer tests with FZK. FZK/GridKa averaging 230 MB/s, doing 240 MB/s or better most of the time, falling to 200 MB/s a few times per day. IN2P3 averaging about 250 MB/s, with a drop to 200 MB/s between 21 and 23 GMT, just like yesterday, possibly due to a daily backup or so. NDGF OK at 60 MB/s. PIC still at 20 MB/s due to many SRM timeouts. TRIUMF OK at 50 MB/s. *Maarten*
- 30/04 02:20 ASGC stable at 120 MB/s. BNL doing 90+ MB/s. DESY 70 MB/s. **GridKa** doing 250 MB/s or better most of the time, but occasionally falls slightly below 200 MB/s. IN2P3 slightly above 250 MB/s most of the time, but occasionally dropping to about 200 MB/s. NDGF stable at 50 MB/s. PIC at one third of their usual rate due to many SRM timeouts. **RAL** dropped to zero around 9 GMT due to a problem with the OPN. TRIUMF stable at 50 MB/s. *Maarten*

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

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

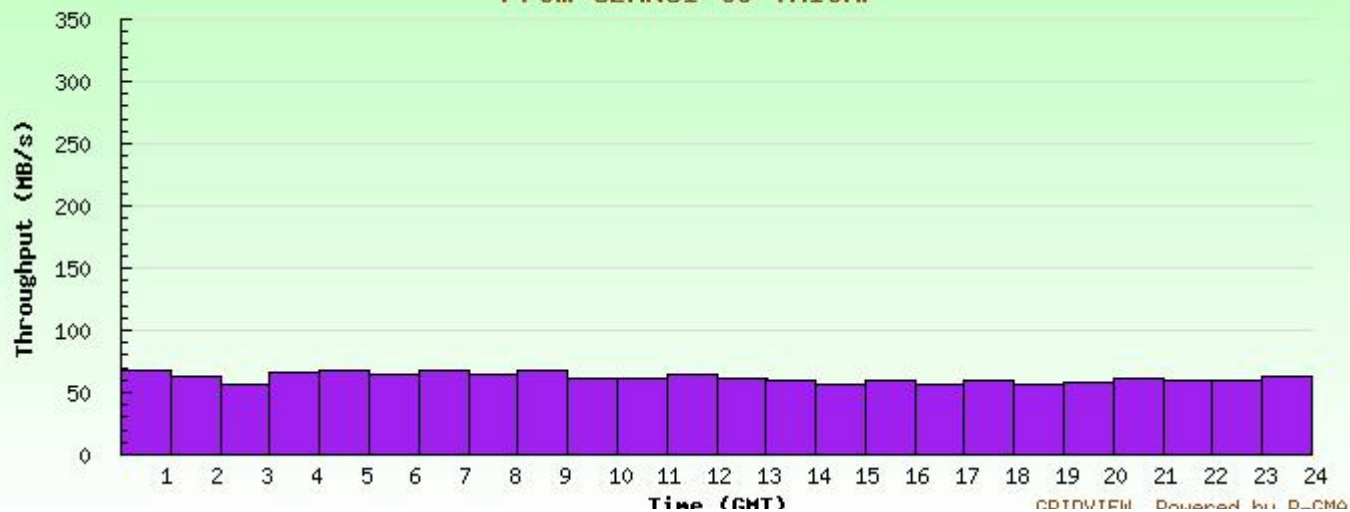
Hourly Averaged Throughput on 16-04-2006
From CERNCI to NDGF

Hourly Averaged Throughput on 16-04-2006
From CERNCI to PIC

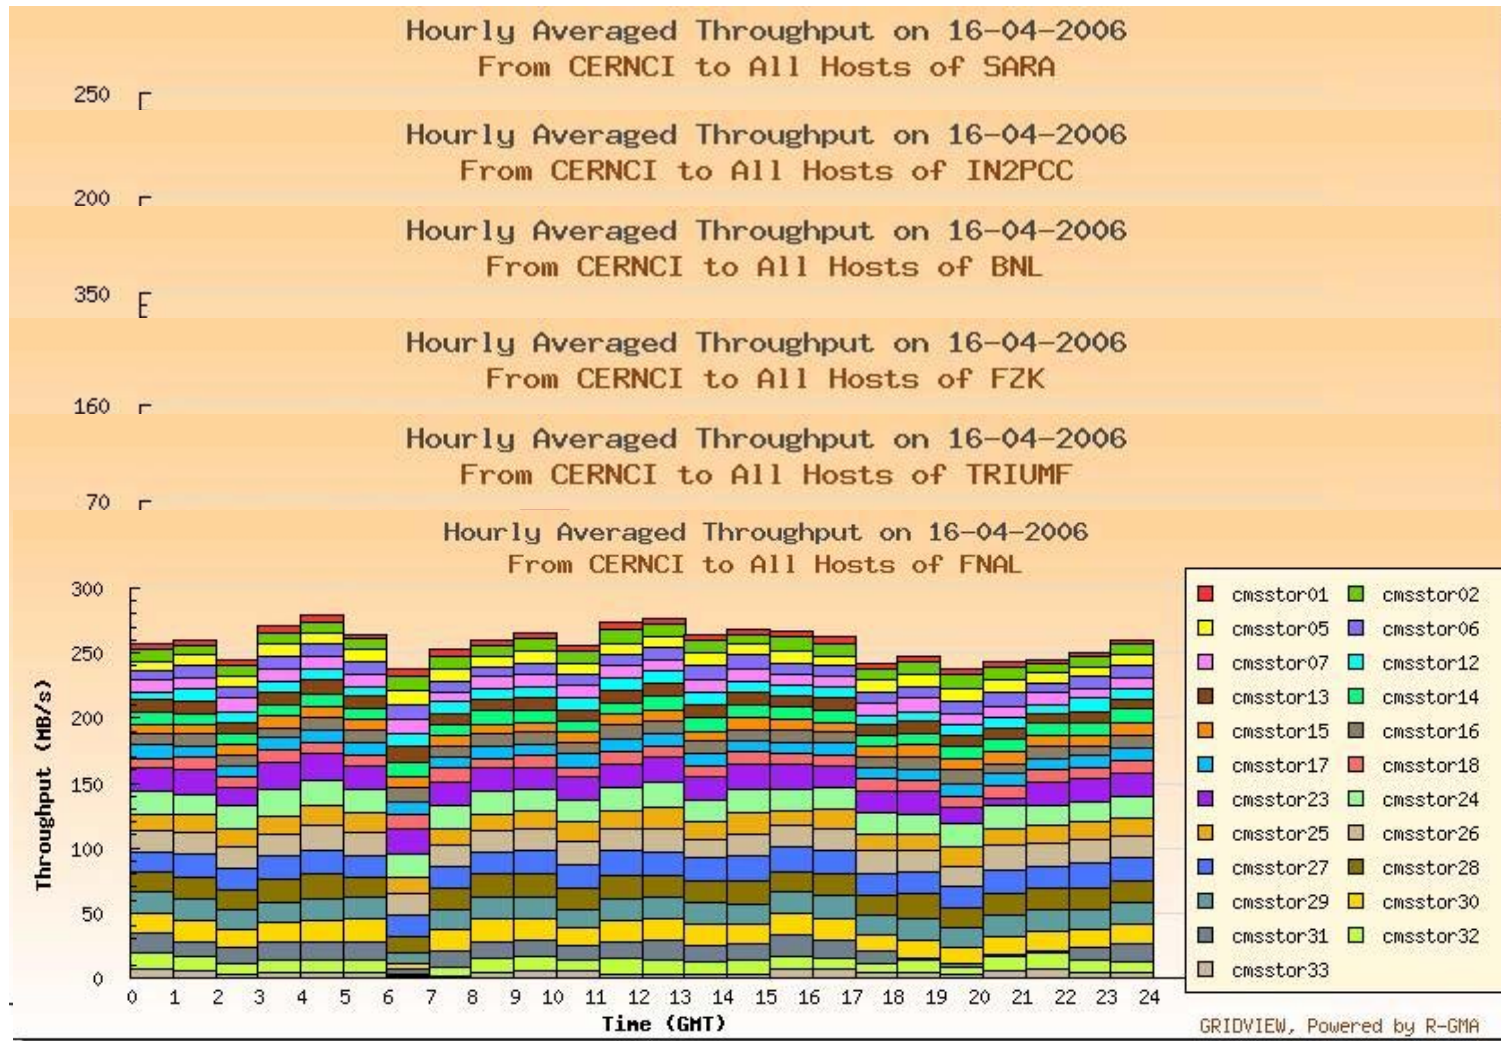
Hourly Averaged Throughput on 16-04-2006
From CERNCI to RAL

Hourly Averaged Throughput on 16-04-2006
From CERNCI to SARA

Hourly Averaged Throughput on 16-04-2006
From CERNCI to TRIUMF

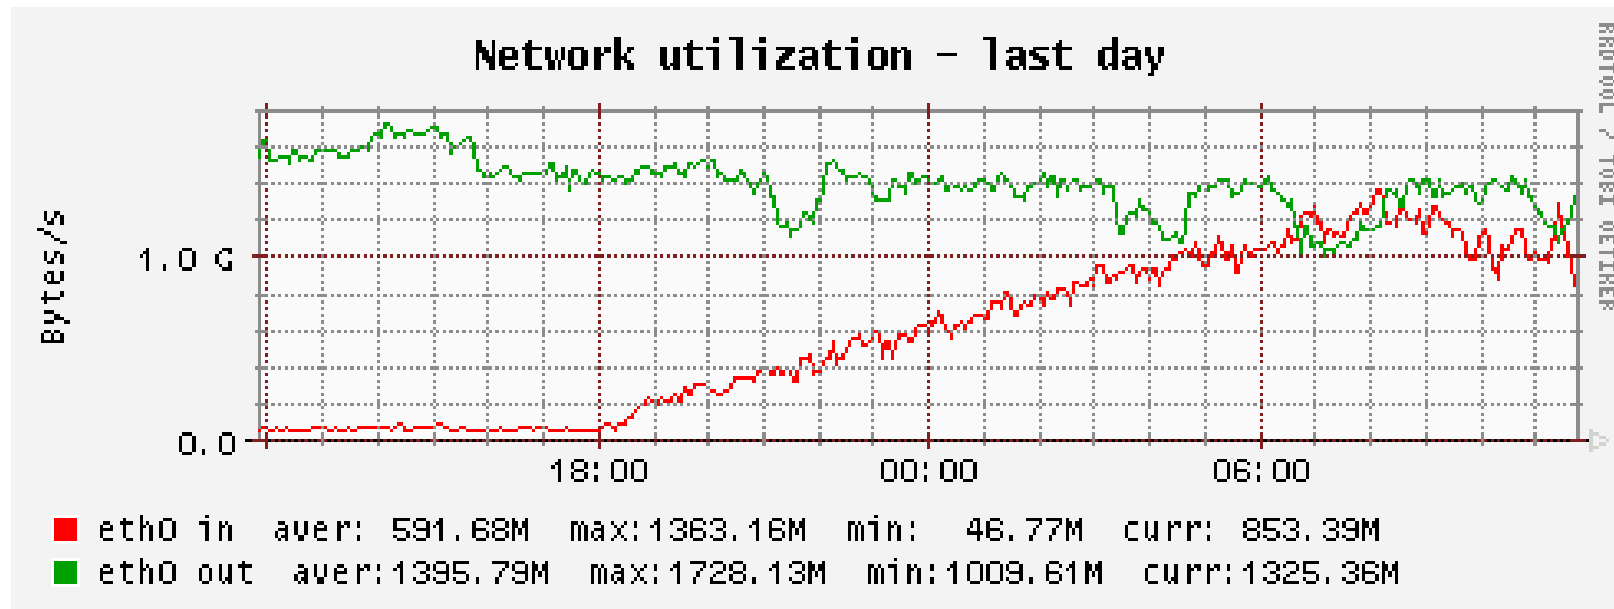


Site by Site Detail



SC4 Disk - Disk Transfers

- Q: Well, 80% of target is pretty good isn't it?




- 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...

SC4 Disk - Tape Transfers

- To reflect current tape hardware and infrastructure, nominal rates scaled to 50 - 75MB/s
 - What can be achieved with 'a few' current drives (~5?)
- Important to build experience with additional complexity of tape backend
 - Before adding Tier1 activities, such as **re-processing**
- Disk - tape had been exercised to a small extent in SC2 and SC3 parts 1 & 2
- Still see more spiky behaviour & poorer stability than disk - disk
- Now need to schedule POW to ramp-up to full nominal rates to tape by September

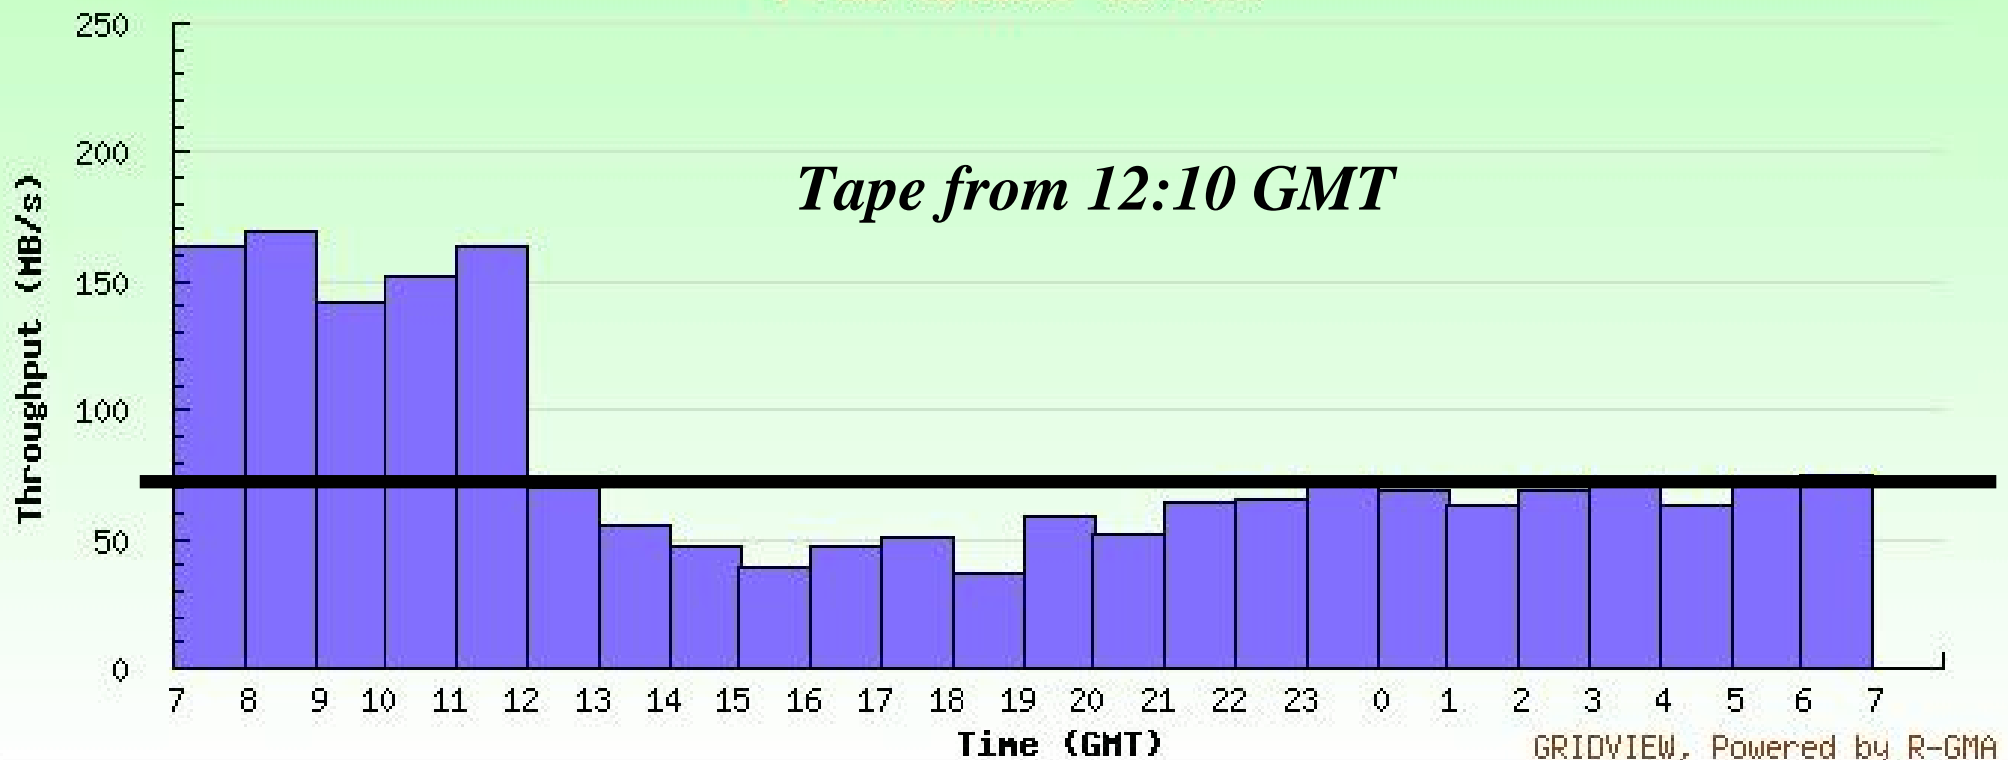
ATLAS Computing Model

- Data reprocessed 2-3 months after it's taken
 - All data reprocessed once per year
 - Done on Tier1 that stores RAW, on tape
 - Potential help from EF farm
 - In parallel with other Tier1 responsibilities
 - RAW, ESD, AOD from T0
 - SIM from Tier2s and other Tier1s, ...
 - Tape system load is critical
- 
- All done in conjunction with acceptance of data from Tier0

BNL tape transfer rate -- Hourly (MB/sec) update: 2006-05-01 13:00

Averaged Throughput during the last 24 hrs (04/05 - 05/05)

From CERNCI to FZK



19 20 21 22 23 24 25 26 27 28 29 30 1 2
Apr06 May06

Next targets: nominal + 50% (backlog)

<i>Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Rate into T1 (pp) Disk-Disk (SRM) rates in MB/s</i>
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 to find sweet spot / plateau.**
 - filesize, number of files / streams etc.
- **Needs to be consistent with actual transfers during data taking**
 - Benefit from experiment T0-T1 tests in SC4 production phase

SC4 - Remaining Challenges

- Full nominal rates *to tape* at all Tier1 sites - sustained!
- Proven ability to ramp-up to nominal rates at LHC start-of-run
- Proven ability to recover from backlogs
 - T1 unscheduled interruptions of 4 - 8 hours
 - T1 scheduled interruptions of 24 - 48 hours(!)
 - 💣 **T0 unscheduled interruptions of 4 - 8 hours**
- Production scale & quality operations and monitoring
- **Monitoring and reporting is still a grey area**
 - I particularly like TRIUMF's and RAL's pages with lots of useful info!

Outline Plan for Tape Transfers

- Some sites - e.g. ASGC - still need to migrate to CASTOR2 (and more...)
- **Need deployment plans for new tape hardware and infrastructure**
- Do not expect all above to have completed by July 2006 - the original target
- However, history has told us that we never get it right first time...
- **Must continue transfers on a regular basis, ramping up progressively in rate towards full nominal, under realistic (data-taking) conditions**
- In parallel, continue to resolve other issues, related to operations, monitoring, rapid ramp-up, handling of backlogs etc.
- **WLCG Level-1 milestone is all Tier1 sites at full nominal rates to tape by end September**

Proposal on how to address this next

Disk - Tape Targets

- Realisation during SC4 that we were simply “turning up all the knobs” in an attempt to meet site & global targets
 - Not necessarily under conditions representative of LHC data taking
- Could continue in this way for future disk - tape tests *but*
- Recommend moving to realistic conditions as soon as possible
 - At least some components of distributed storage system not necessarily optimised for this use case (focus was on local use cases...)
 - ☛ If we **do** need another round of upgrades, know that this can take 6+ months!
- Proposal: benefit from ATLAS (and other?) Tier0+Tier1 export tests in June + Service Challenge Technical meeting (also June)
 - Work on operational issues can (must) continue in parallel
 - As must deployment / commissioning of new tape sub-systems at the sites
 - e.g. milestone on sites to perform disk - tape transfers at > (>>) nominal rates?
- This will provide some feedback by late June / early July
 - Input to further tests performed over the summer

ATLAS Tier0 + Tier1 Export Rates + Sites

Centre	ATLAS	SC4 proposal	Nominal
ASGC	60.0	60	100
CNAF	59.0	60	200
PIC	48.6	50	100
IN2P3	90.2	100	200
GridKA	74.6	80	200
RAL	59.0	60	150
BNL	196.8	200	200
TRIUMF	47.6	50	50
SARA	87.6	90	150
NDGF	48.6	50	50
FNAL	-	-	200

A more realistic test for July – rather than another dTeam test – would be to add other VO's in parallel to ATLAS (once understood).

See next slides for CMS, LHCb and ALICE goals...

(Background dTeam transfers on-going...)

- The only site not involved in these tests is FNAL.
- The rates to most sites (except ATLAS-only sites) is (much) lower than the full nominal rate for that site.
- All above sites *should* be able to (can) sustain these rates - both to disk and tape - *now*.

Combined Tier0 + Tier1 Export Rates

<i>Centre</i>	<i>ATLAS</i>	<i>CMS*</i>	<i>LHCb⁺</i>	<i>ALICE</i>	<i>Combined</i>	<i>Nominal</i>
ASGC	60.0	10	-	-	70	100
CNAF	59.0	25	23	?	108	200
PIC	48.6	30	23	?	103	100
IN2P3	90.2	15	23	?	138	200
GridKA	74.6	15	23	?	95	200
RAL	59.0	10	23	?	118	150
BNL	196.8	-	-	-	200	200
TRIUMF	47.6	-	-	-	50	50
SARA	87.6	-	23	?	113	150
NDGF	48.6	-	-	?	50	50
FNAL	-	50	-	-	50	200
Totals					~1150	1600

* CMS target rates double by end of year

+ Mumbai rates - scheduled delayed by ~1 month (start July)

Transfers: Who Drives them?

- Responsibility for generating data files / submitting transfers lies with the experiments
 - Team to provide > 8 x 5 coverage required? (Data-taking)
- **Monitoring of transfers also requires support at site level**
- **Clear need & motivation to automate / alarm as much of this as possible**
- ATLAS Tier0-Tier1 transfer tests resembles a data-taking period
- Do we want / need to establish a service coordinator / shift-crew?
- Should this be a medium-term requirement on sites?
 - Provide manpower for shifts / eventually rotate service coordinator responsibility
- **Man-power also from other groups / teams at CERN: ARDA / EIS?**

Components Involved (Simplified!)

Tier0 → Network (varies) →

Tier1 SRM → mass storage adapter →

Mass storage system → Tape subsystem

- **No combination of SRM implementation + MSS adaptor + MSS + Tape subsystem is the same!**

SC4 - Successes & Remaining Work

We have shown that we can drive transfers at full nominal rates to:

- Most sites simultaneously;
- All sites in groups (modulo network constraints - PIC);
- At the target nominal rate of 1.6GB/s expected in pp running


In addition, several sites exceeded the disk - tape transfer targets

- There is no reason to believe that we cannot drive all sites at or above nominal rates for sustained periods.

But

- There are still major operational issues to resolve - and most importantly - a full end-to-end demo under realistic conditions

Conclusions




- We have demonstrated - through the SC3 re-run and more convincingly through SC4 - that we can send data to the Tier1 sites at the required rates for extended periods
 - Disk - tape rates are reasonably encouraging but still require full deployment of production tape solutions across all sites to meet targets

➤ **Demonstrations of the needed data rates corresponding to experiment transfer patterns must now be proven**

- As well as an acceptable - and affordable - service level

- Moving from dTeam to experiment transfers will hopefully also help drive the migration to full production service
 - Rather than the current 'best' (where 'best' is clearly +ve!) effort

Conclusions



- There is already a need for continuous production reliable file transfer services
- In parallel, there is much work remaining to ramp-up in rate and reliability and to include the additional complexity of realistic LHC data taking and re-processing / analysis conditions
- We have made much progress over the past 18 months...
- ... but we still have a lot *more* to do in *less* than 1/3 of the time...
- Not to mention the parallel service deployment / debugging...