

HEPiX Rome - April 2006

The High Energy Data Pump

A Survey of State-of-the-Art Hardware & Software Solutions

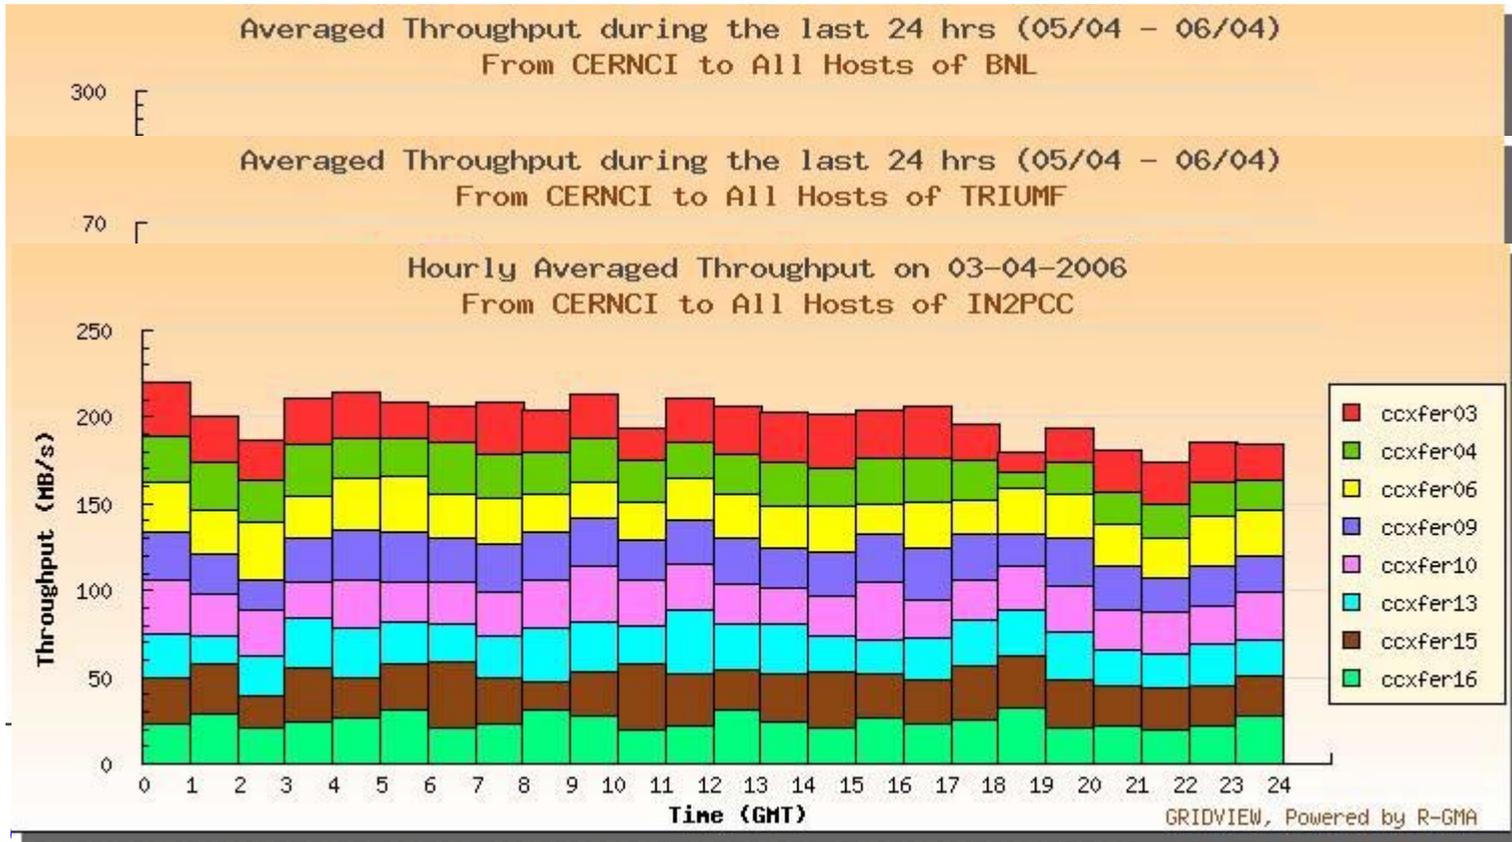
Martin Gasthuber / DESY

Graeme Stewart / Glasgow

Jamie Shiers / CERN

Summary

-
-
-
-
-
-
-
-



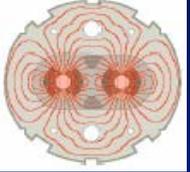
width

al

order!

Communication and Collaboration Fundamentals

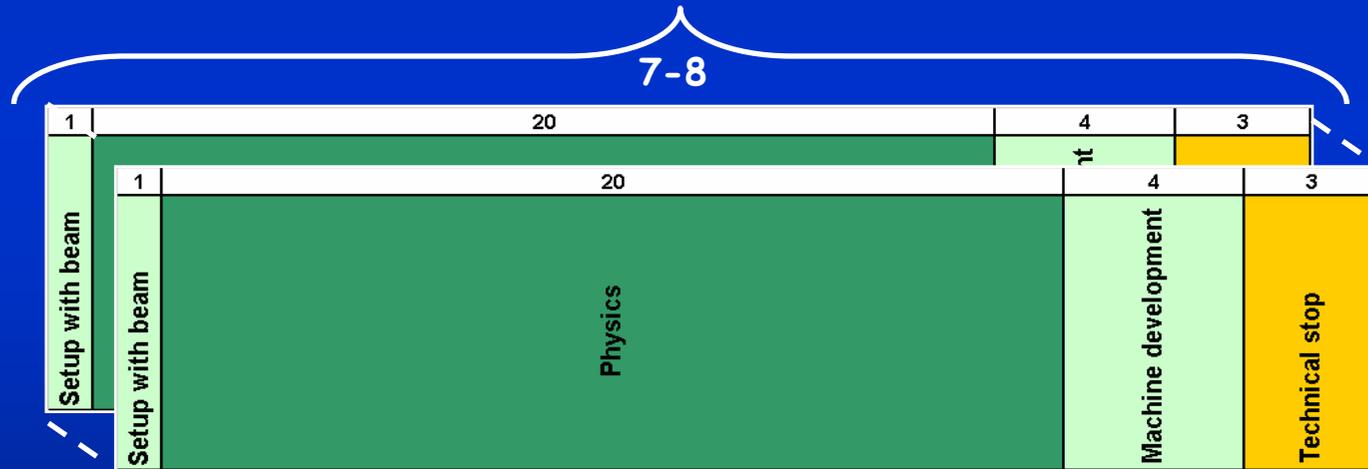
- The schedule of our future follows (7 - 8 periods of 1 month)...



Breakdown of a normal year

- From Chamonix XIV -

Service upgrade slots?

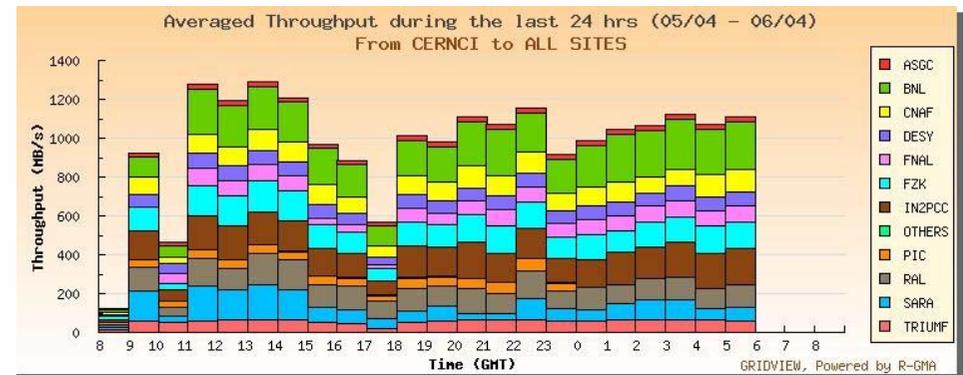


~ 140-160 days for physics per year
 Not forgetting ion and TOTEM operation
 Leaves ~ 100-120 days for proton luminosity running
 ? Efficiency for physics 50% ?
 ~ 50 days ~ 1200 h ~ 4×10^6 s of proton luminosity running / year

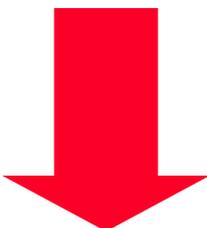
Achieved Rates - Daily Averages for Disk-Disk Transfers

- Rates achieved as per [Gridview](#) - see [SC4blog](#) for more information on events, interventions, changes etc.
- Note that the official goal is to **meet or exceed** the target rate *every* day (or make up for it shortly after - **more information below**). In addition, the startup (ramp-up to full nominal rates) needs to be rapid - essentially a step function at the beginning of each LHC running period.

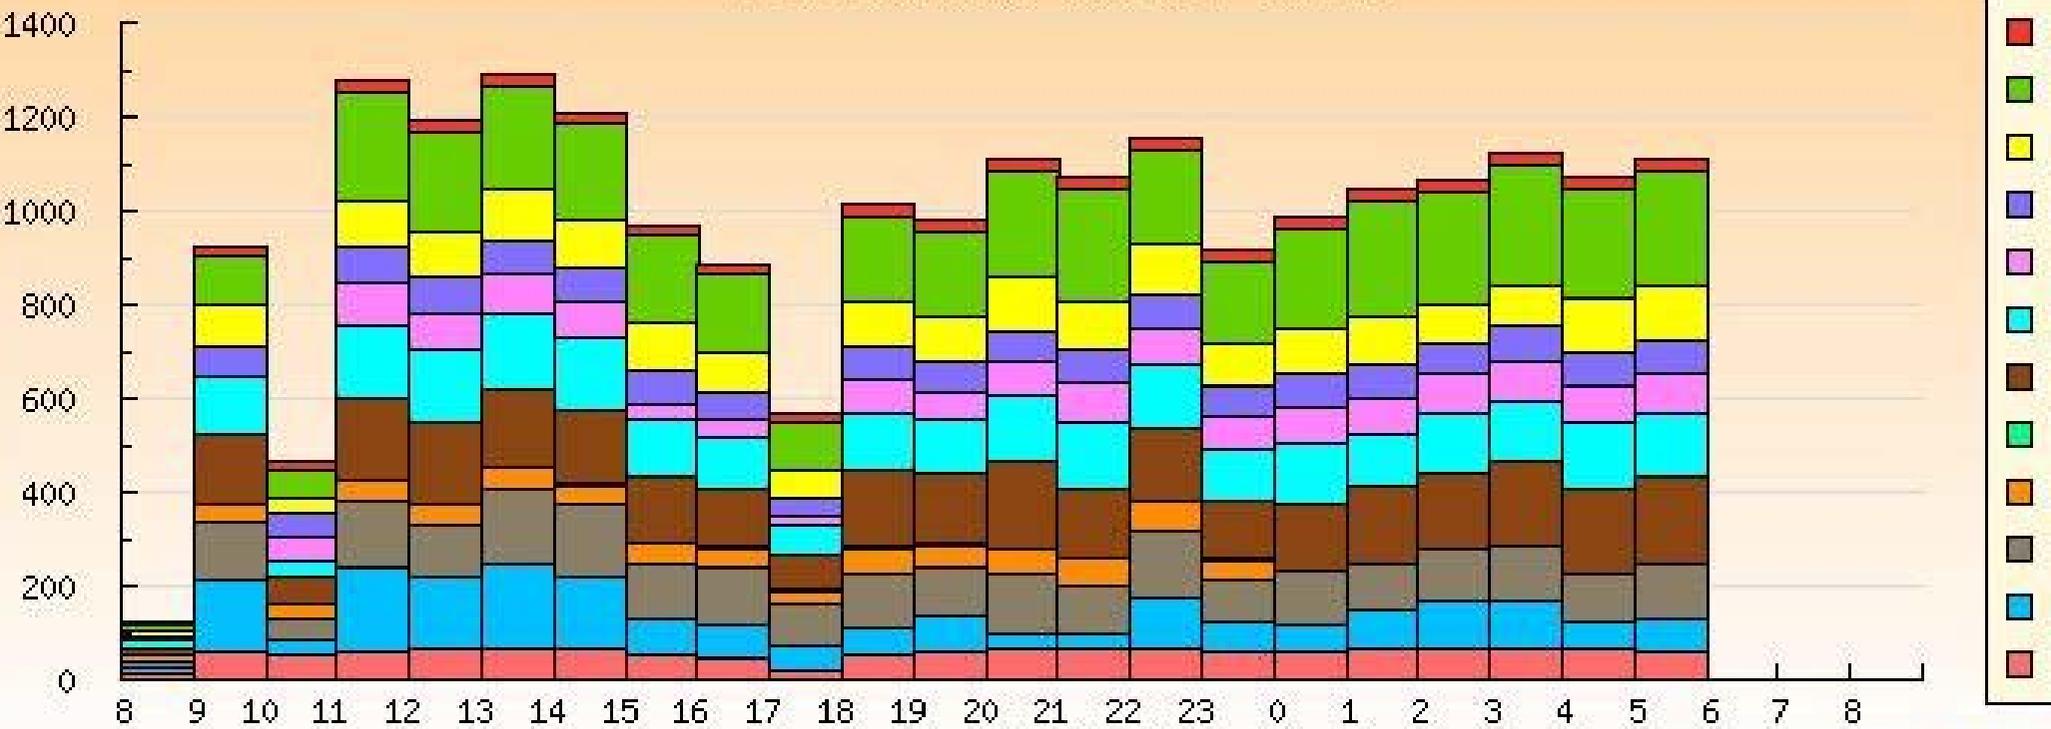
Site	Disk-Disk	Apr3	Apr4	Apr5	Apr6
ASGC	100	-	7	23	
TRIUMF	50	44	42	55	
BNL	200	170	103	173	
FNAL	200	-	-	38	
NDGF	50	-	-	-	
PIC	60	-	18	41	
RAL	150	129	86	117	
SARA	150	30	78	106	
IN2P3	200	200	114	148	
FZK	200	81	80	118	
CNAF	200	55	71	92	
TOTAL	1600	709	599	911	



The red line is the target average daily rate!



Averaged Throughput during the last 24 hrs (05/04 - 06/04)
From CERNCI to ALL SITES



ServiceChallengeFourBlog (LCG→SC)

- 06/04 09:00 TRIUMF exceeded their nominal data rate of 50MB/s yesterday, despite the comments below. Congratulations! *Jamie*
- 05/04 23:59 A rough day with problems that are not yet understood (see the tech list), but we also reached the highest rate ever (almost 1.3 GB/s) and we got FNAL running with srmcopy. Most sites are below their nominal rates, and at that they need too many concurrent transfers to achieve those rates, so we still have some debugging ahead of us. CASTOR has been giving us timeouts on SRM get requests and Olof had to clean up the request database. To be continued... *Maarten*
- 05/04 16:30 The Lemon monitoring plots show that almost exactly at noon the output of the SC4 WAN cluster dropped to zero. It looks like the problem was due to an error in the load generator, which might also explain the bumpy transfers BNL saw. *Maarten*
- 05/04 11:02 Maintenance on USLHCNET routers completed. (During the upgrade of the Chicago router, the traffic was rerouted through GEANT). *Dan*
- 05/04 11:06 Database upgrade completed by 10am. DLF database was recreated from scratch. Backup scripts activated. DB Compatibility moved to release 10.2.0.2, automatic startup/shutdown of the database tested. *Nilo*
- 05/04 10:50 DB upgrade is finished and CASTOR services have restarted. SC4 activity can resume. *Miguel*
- 05/04 09:32 SC4 CASTOR services stopped. *Miguel*
- 05/04 09:30 Stopped all channels to allow for upgrade of Oracle DB backend to more powerful node in CASTOR. *James*
- 04/04 IN2P3 meet their target nominal data rate for the past 24 hours (200MB/s). Congratulations! *Jamie*

Conclusions

- We have to **practise** repeatedly to get sustained daily average data rates
- Proposal is to repeat an LHC running period (reduced - just ten days) every month
- Transfers driven as **dteam** with low priority
- We still have to add tape backend - not to mention the full DAQ-T0-T1 chain - and drive with experiment software

First data will arrive next year

NOT an option to get things going later

