Status, results and plans of WLCG

D. Bonacorsi

[INFN-CNAF, Bologna --- CMS Facilities / Infrastructure Ops coordinator]

CAVEAT: this talk is my personal view of the status and results of the WLCG collaboration.

Thanks to Jamie Shiers and Harry Renshall for providing base material for this talk.

Outline

WLCG and MoU

- Tiers' roles and responsibility in a WLCG-enabled world
- Technical participation of a Tier in WLCG

Service Challenges (SC)

- Roles of Tiers in SCs
- Roles of LHC experiments in SCs
- Some highlights on SC achievements, open issues

Common-VO Computing Readiness Challenge 2008 (CCRC'08)

- Same as above

Summary

WLCG

Worldwide LHC Computing Grid (WLCG)

Purpose: to provide the computing resources needed to process and analyze the data gathered by the LHC experiments

The LCG project is launched to:

- assemble at multiple computing centres the main offline data-storage and computing resources needed by LHC exps
- operate these resources in a shared Grid-like manner

Main goal: provide common tools and implement uniform means of accessing resources

MoU [*] for collaboration in the Deployment and Exploitation of the Worldwide LHC Computing Grid between "the Parties" [#]

[*] CERN-C-RRB-2005-01/Rev. 21 March 2006

[#] see next slide

After LCG Phase 1 (technology development and tests leading to a production prototype), the WLCG MoU governed the execution of the LCG Phase-2 (deployment and exploitation of LCG as a Service)

Defines the program of work, distribution of duties and responsibilities to the Parties as well as the Computing Resources Levels they will offer to the LHC exps (also organizational, managerial and financial guidelines).

The WLCG MoU "parties"

CERN:

- Tier-0:
  - receives raw data from exps' online computing farms and records them on permanent MSS + 1st-pass reco + distribute them to Tier-1's
- CERN Analysis Facility
  - functionality of a combined T1-T2 centre, except that it does not offer permanent storage of back-up copies of raw data

All Institutions participating in the provision of the WLCG with a T1 and/or T2 computing centre (federations included):

- Tier-1's:
  - provide a *distributed permanent back-up* of the raw data, storage and mgmt of data needed during the analysis process + offer a grid-enabled data service
  - data intensive analysis and reprocessing
  - may undertake national or regional support tasks, and contribute to Grid Operations Services
- Tier-2's
  - provide well-managed, grid-enabled disk storage and concentrate on tasks such as simulation, end-user analysis and high-performance parallel analysis

Tiers architecture

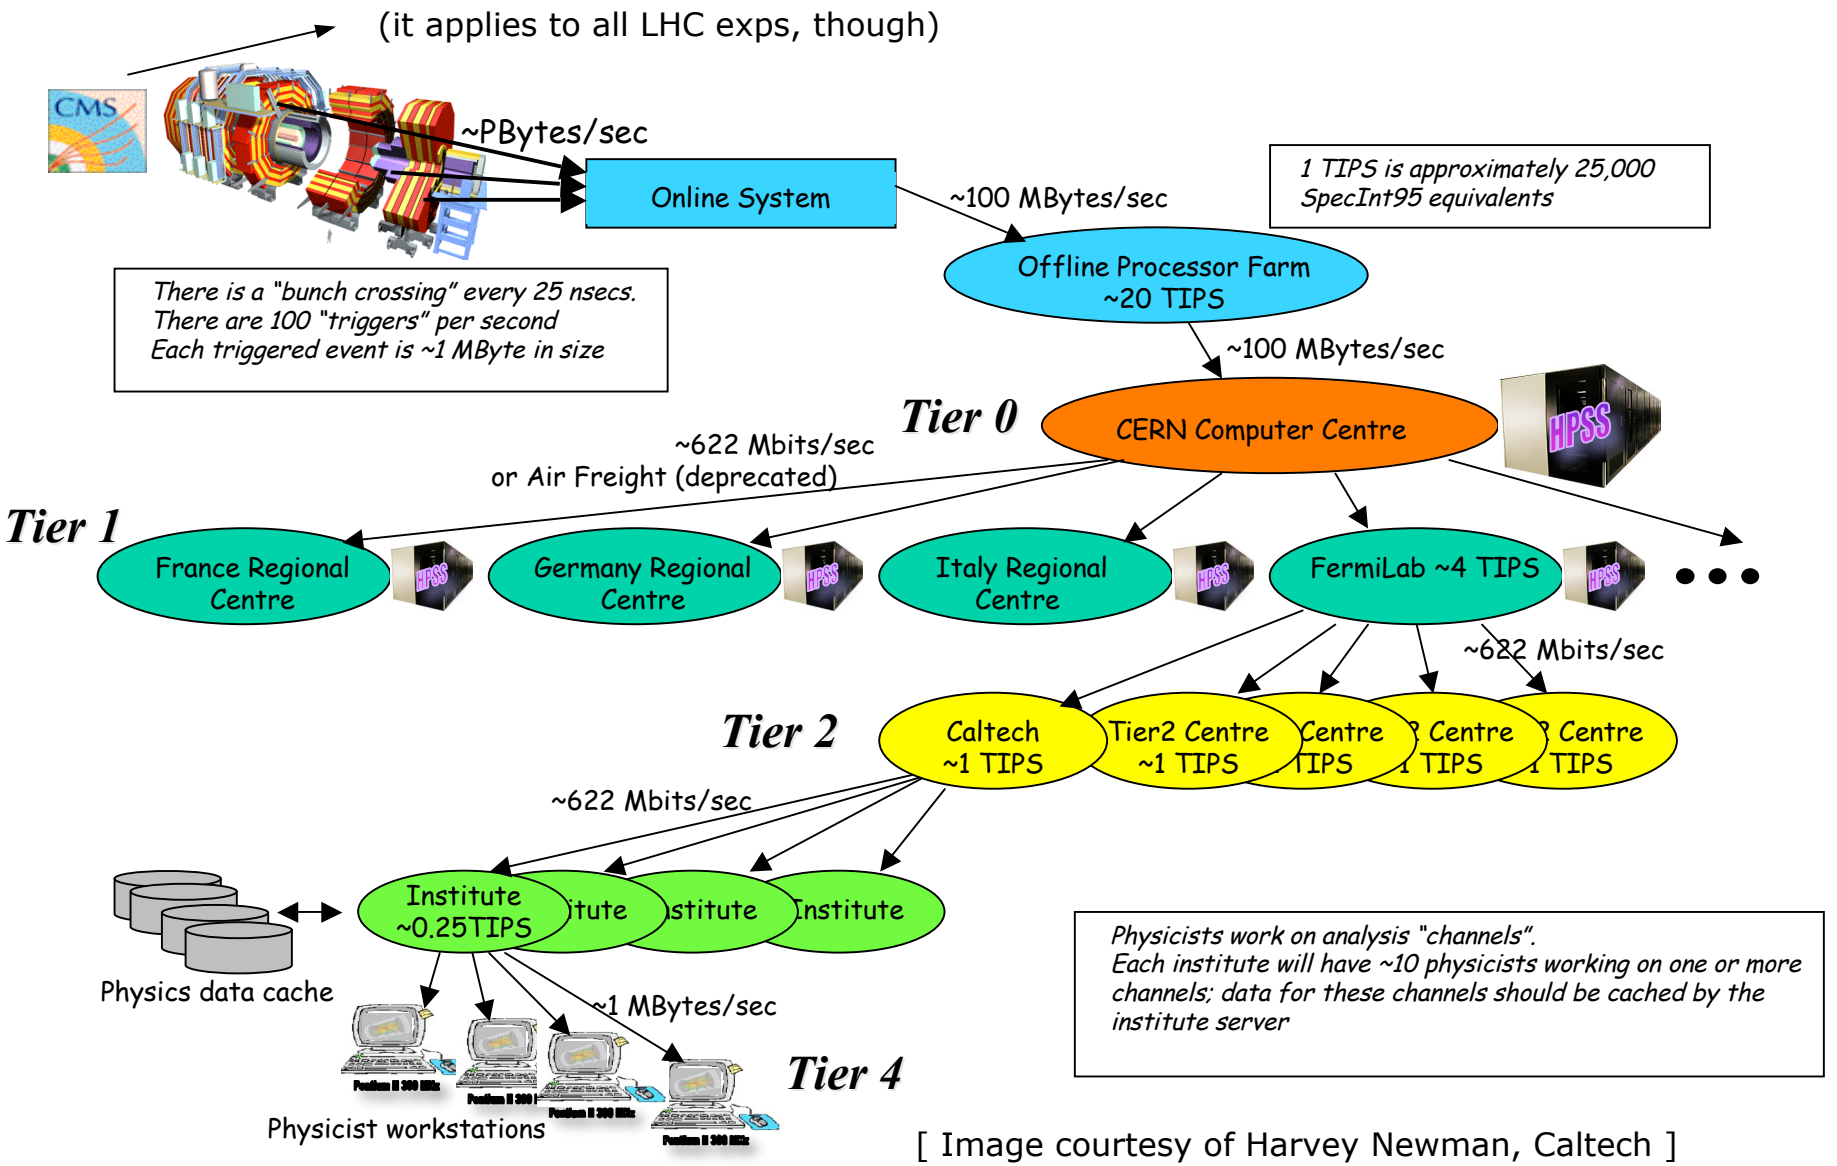

[Image courtesy of Harvey Newman, Caltech]

Minimal computing capacities for **Tier-0**

(i.e. minimal Computing Resources and Service Levels to qualify for WLCG membership)

In support of offline computing systems of LHC exps according to their Computing Models, CERN shall mainly supply the following services:

- operation of the **Tier-0 facility**, providing:
  - high bandwidth network connectivity from exp. area to offline computing facility;
  - recording / permanent storage (MSS) of one copy of *raw data* throughout exp lifetime;
  - distribution of an agreed share of *raw data* to each Tier-1 centre;
  - 1st pass calibration and alignment processing, including sufficient buffer storage of associated calibration samples for up to 24 hrs + event reco according to policies agreed with exps;
  - storage of the *reco data* on disk and in MSS;
  - distribution of an agreed share of the *reco data* to each Tier-1 centre;
- operation of the **CERN Analysis Facility**, providing:
  - data-intensive analysis, high-performance access to current versions of the exps' real/simulated datasets;
  - Eventually end-user analysis.
    - i.e. all functionalities of a combined T1/T2, except from permanent storage of back-up copies of raw data
- provision of base services for Grid Coordination and Operation:
  - Overall management and coordination of the LHC Grid
    - integration, certification, distribution, support for software required for Grid operations
- support, at several levels:
  - network issues, databases, tools, libraries, infrastructures, VOs-management, etc…

Minimal computing capacities for **Tier-1's**

(i.e. minimal Computing Resources and Service Levels to qualify for WLCG membership)

Tier-1 centres form an integral part of data handling service of LHC exps

They undertake to provide services:

on a long-term basis (initially at least 5 yrs)

upgrading also, to keep pace with expected growth of LHC data volumes and analysis activities

with high level of reliability/availability + rapid responsiveness to problems

Wide pletora of services provided by each Tier-1 to LHC exps they serve:

acceptance of agreed share of *raw data* from Tier-0, keeping up with DAQ;
acceptance of agreed share of *1st-pass reco data* from Tier-0;
acceptance of processed and simulated data from other WLCG centres;
recording + archival storage of accepted share of *raw data* ("distributed back-up");
recording + maintenance of processed/simulated data on MSS
provision of managed storage space for permanent/tmp storage of files and dbs;
provision of access to the stored data by other WLCG centres;
operation of a data-intensive analysis facility;
provision of other services according to agreed exps' requirements;
ensure high-capacity network bandwidth and services for data exchange with Tier-0

(as part of an overall plan agreed amongst exps, Tier-0 and Tier-1's);

ensure network bandwidth and services for data exchange with Tier-1's and Tier-2's

(as part of an overall plan agreed amongst the exps, Tier-1's and Tier-2's);

administration of databases (and more…) required by exps at Tier-1's.

All storage/computational services shall be "grid enabled"

according to standards agreed between LHC exps and the regional centres

Minimal computing capacities for **Tier-2**'s

(i.e. minimal Computing Resources and Service Levels to qualify for WLCG membership)

Tier-2 services shall be provided by centres (or federations of centres):

as a guideline, Tier-2's are expected to be capable of fulfilling at least a few % of the resource requirements of the LHC exps that they serve.

Services provided by each Tier-2 are:

provision of managed storage space for permanent/tmp storage of files and dbs;

provision of access to the stored data by other WLCG centres;

operation of an end-user analysis facility;

provision of other services, e.g. simulation, according to agreed exp requirements;

ensure network bandwidth and services for data exchange with Tier-1's (as part of an overall plan agreed amongst exps, Tier-1's and Tier-2's);

As for T1's: "grid-enabling" of all storage/computational services

Technical participation to WLCG

The technical participation of the "Parties" is defined separately in terms of:

- **Computing Resources levels** that they pledge to provide to one/more LHC exps that they serve;
- **Computing Services levels** that they pledge to the WLCG collaboration.

In both cases having secured the necessary funding…

- for iron objects and warm bodies

Crucial to define (and apply) metrics:

- associate with each element a set of key qualitative measures:
  - e.g. Computing Resource Level associated to "Networking" shall include I/O throughput and average availability in terms of [time running]/[scheduled up-time]
  - e.g. Computing Service level imply reliability, availability, responsiveness to problems, …

Let's dig into both resources and services…

'Resources': the pledged Computing Capacities

CERN Tier0	Pledged	Planned to be pledged			
	2006	2007	2008	2009	2010
CPU (kSI2K)	2400	4800	12500	15900	26200
Disk (Tbytes)	230	450	1300	1300	1800
Tape (Tbytes)	1500	3400	13600	23600	33900
Nominal WAN (Mbits/sec)	80000	100000	120000	140000	160000

CERN Analysis Facility	Pledged	Planned to be pledged			
	2006	2007	2008	2009	2010
CPU (kSI2K)	1000	4320	10000	14600	15000
Disk (Tbytes)	540	1600	4200	5200	5300
Tape (Tbytes)		1480	3400	5700	5900

CNAF, Italy	Pledged	Planned to be pledged			
	2006	2007	2008	2009	2010
CPU (kSI2K)	1800	2400	5500	8000	11500
Disk (Tbytes)	850	1200	2500	4000	5800
Tape (Tbytes)	850	1000	2100	4100	6000
Nominal WAN (Mbits/sec)	5000	10000	20000	30000	40000

Italy, INFN Tier-2 Federation	Pledged	Planned to be pledged			
	2006	2007	2008	2009	2010
CPU (kSI2K)	1000	2500	5000	8000	11000
Disk (Tbytes)	200	600	1500	2500	3500
Nominal WAN (Mbits/sec)					

'Services': required components at Tiers

To Service	Only some VOs?	Class
SRM		C
LFC	LHCb	C
LFC	ALICE, ATLAS	H
FTS		C
CE		C
RB		C
Global BDII		C
Site BDII		H
Myproxy		C
VOMS		H→C
R-GMA		H
databases		C/H

T1 Service	Only some VOs?	Class
SRM 2.1		H/M
LFC	ALICE, ATLAS	H/M
FTS		H/M
CE		H/M
Site BDII		H/M
R-GMA		H/M

T2 Service	Only some VOs?	Class
SRM 2.1		M/L
LFC	ATLAS, ALICE	M/L
CE		M/L
Site BDII		M/L
R-GMA		M/L

Class	Description	Downtime	Reduced	Degraded	Availability
C	Critical	1 hour	1 hour	4 hours	99%
H	High	4 hours	6 hours	6 hours	99%
M	Medium	6 hours	6 hours	12 hours	99%
L	Low	12 hours	24 hours	48 hours	98%
U	Unmanaged	None	None	None	None

SERVICE LEVELS definition:

Downtime defines time between start of problem and restoration of service at min capacity (i.e. basic function but capacity < 50%)

Reduced defines time between the start of the problem and the restoration of a reduced capacity service (i.e. >50%)

Degraded defines the time between the start of the problem and the restoration of a degraded capacity service (i.e. >80%)

Availability defines the sum of the time that the service is down compared with the total time during the calendar period for the service. Site wide failures are not considered as part of the availability calculations. *99% means a service can be down up to 3.6 days a year in total.* 98% means up to a week in total.

None means the service is running unattended

'Services': the MoU availability targets

(to be reviewed by the operational boards of the WLCG Collaboration)

T0

Service	*Maximum delay in responding to operational problems*			*Average availability[2] measured on an annual basis*	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction or distribution of data to Tier-1 Centres during accelerator operation	6 hours	6 hours	12 hours	99%	n/a
Networking service to Tier-1 Centres during accelerator operation	6 hours	6 hours	12 hours	99%	n/a
All other Tier-0 services	12 hours	24 hours	48 hours	98%	98%
All other services[3] – prime service hours[4]	1 hour	1 hour	4 hours	98%	98%
All other services – outwith prime service hours	12 hours	24 hours	48 hours	97%	97%

T2's

Service	*Maximum delay in responding to operational problems*		*Average availability[2] measured on an annual basis*
	Prime time	*Other periods*	
End-user analysis facility	2 hours	72 hours	95%
Other services	12 hours	72 hours	95%

T1's

Service	*Maximum delay in responding to operational problems*			*Average availability measured on an annual basis*	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outwith accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services – prime service hours[6]	2 hour	2 hour	4 hours	98%	98%
All other services – outwith prime service hours	24 hours	48 hours	48 hours	97%	97%

The response times in the above table refer only to the maximum delay before action is taken to repair the problem. The mean time to repair is also a very important factor that is only covered in this table indirectly through the availability targets. All of these parameters will require an adequate level of staffing of the services, including on-call coverage outside of prime shift.

Experiments and sites

Centre	ALICE	ATLAS	CMS	LHCb
Canada, TRIUMF		X		
France, CC-IN2P3	X	X	X	X
Germany, FZK-GridKA	X	X	X	X
Italy, CNAF	X	X	X	X
Netherlands LHC/Tier1	X	X		X
Nordic Data Grid Facility	X	X	X	
Spain, PIC Barcelona		X	X	X
Taipei, ASGC		X	X	
UK, RAL	X	X	X	X
USA, BNL		X		
USA, FNAL			X	

VO-view ↔ WLCG-view

The 2 views are not diverging of course… but are different

1. One VO may rely more on a given T1 than all other VO's
   - ALICE mostly rely on GridKa; ATLAS on BNL,IN2P3,NL; CMS on FNAL,CNAF; LHCb on IN2P3,NL
   - WLCG focuses on providing tools and help for multi-VO load-balanced usage of Tiers
2. The WLCG planned tests must take VO schedules into account
3. The VO schedules are tight, and may be well different from each other

The challenge is also in finding a productive synergy

NEVER FORGET WE HAVE THE SAME MAIN GOAL!

WLCG Service Challenges

What is a 'SC' ?

A mechanism by which the readiness of the overall LHC computing infrastructure to meet the exps' requirements is measured and if(/where) necessary corrected

Aim is to understand what it takes to run a **real and wide Grid (set of) service(s)**

joined effort with Tiers community to

- to trigger resources deployment
- to drive activity planning
- to encourage distributed know-how based on realistic use patterns
- to ramp-up essential grid services to target levels of reliability, availability, scalability, end-to-end performance

A long and hard path…

… done in several steps… (see next slide)

WLCG SCn, n=1,4

Service Challenge 1 and 2 (SC1, SC2)

building up the necessary data mgmt infrastructure to perform reliable To-T1 transfers and permanent production services with the appropriate throughput: basically build up **To-T1 infrastructure and services** to handle production transfers and data flows

no T2 involved, no exp-specific sw and *no offline use-cases* included

SC2 met its throughput goal (100 MB/s/site, 500 MB/s sustained out of CERN) , but not its service goals

Service Challenge 3 and 4 (SC3, SC4)

bringing **T2s** into loop and fully address **exps use-cases**

towards full production services:

add site/exps goals, e.g. focus on DM, batch prod and real data access

addressing problems beyond "set-up (throughput)" goals: add a "service phase"

WLCG deployment schedule 2006-2008

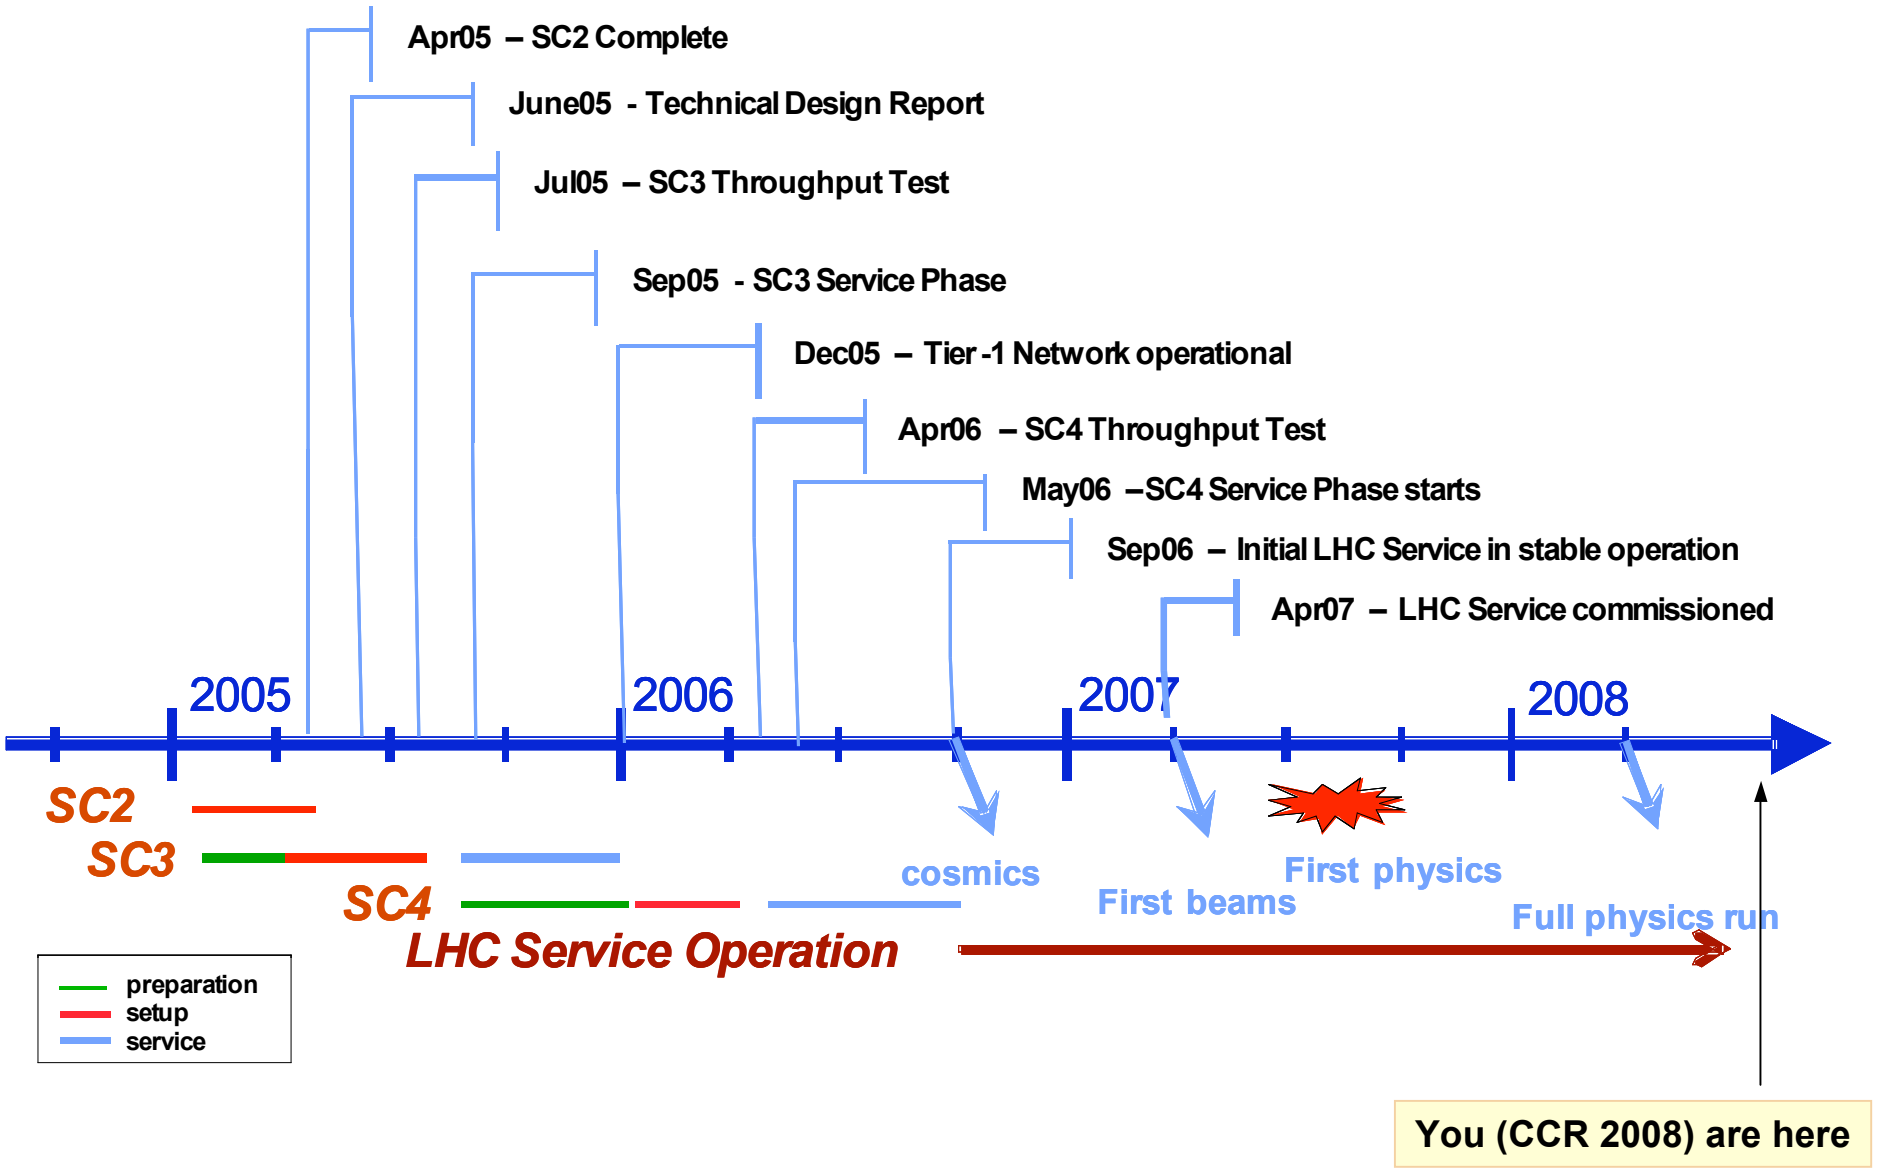

Service Challenge 3 (SC3)

The first SC with exps-oriented objectives

when: Jul 05 - Dec 05 (+ Jan 06)

who: To, all T1's, small nb of T2's

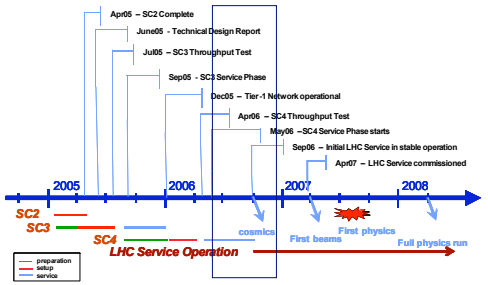

Set-up ("throughput") phase (1 month: Jul 05)

Throughput targets for each T1 are 150 MB/s network-disk and 60 MB/s network/tape, with CERN capable of supporting 1 GB/s for the transfers to disk and 400 MB/s for those to tape at the T1s. Some T1s also support T2s, which may upload simulated data and download analysis data.

"Service" phase (4 months: Sep-Dec 05)

Stable operation during which exps are committed to carry out tests of their sw chains and computing models

Includes additional sw components, including a grid WMS, Grid catalog, mass storage mgmt services and a file transfer service

Re-run of the throughput phase (Jan 06)

SC3 throughput phase (Jul 05)

Jul 05 throughput tests did not meet targets (~50% higher than SC2)

on average: ~1/2 target, scarce stability

Network-disk target:
150MB/s/T1
1GB/s (CERN)

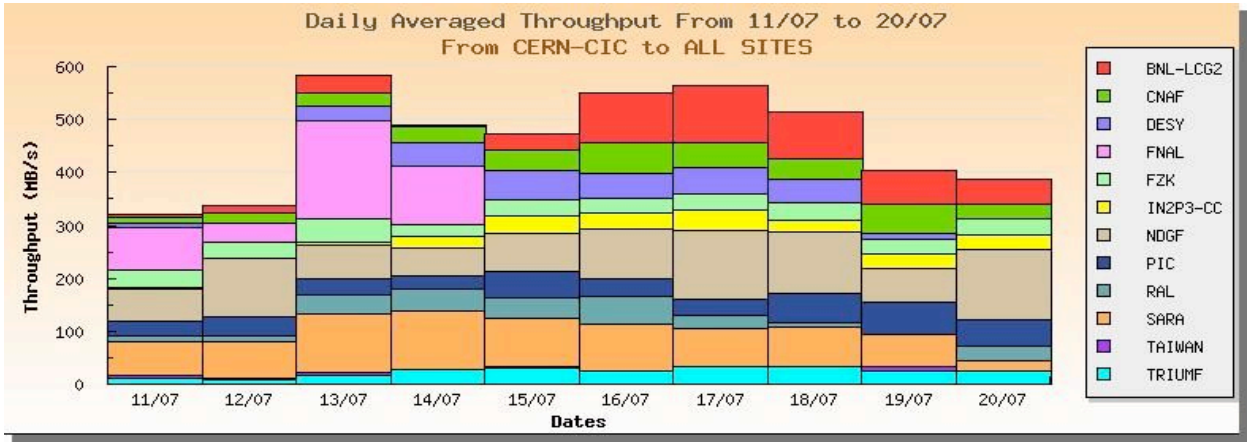

Network-tape target:
60MB/s/T1
(not all sites, though)

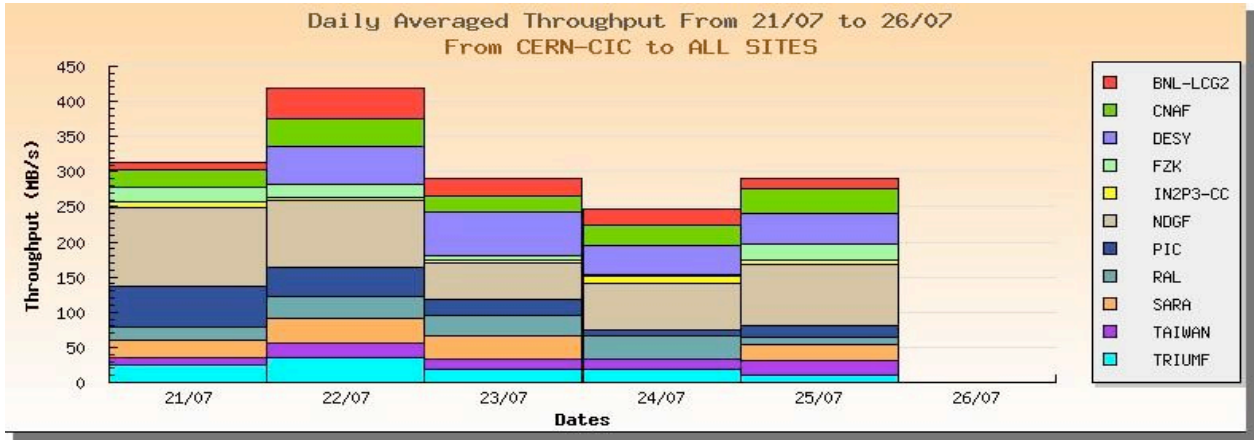

important step to gain experience with the services before SC4

since then: SRM on each site, dCache 1.6.6+, gLite FTS 1.4, CASTOR fixes, a lot of debugging in transfers, network upgrades: all sw upgrades released and deployed

SC3 To-T1 disk-disk re-run (Jan 06)

substantial improvement in recent SC3 throughput re-run…

- some sites exceeded the targets
- all sites clearly have a much better handle on the systems

… but still a lot of work to do…

- daily average rate still far away, demonstrate sustainability of met/exceeded targets, test recovery from problems, … and get rid of heroic efforts!

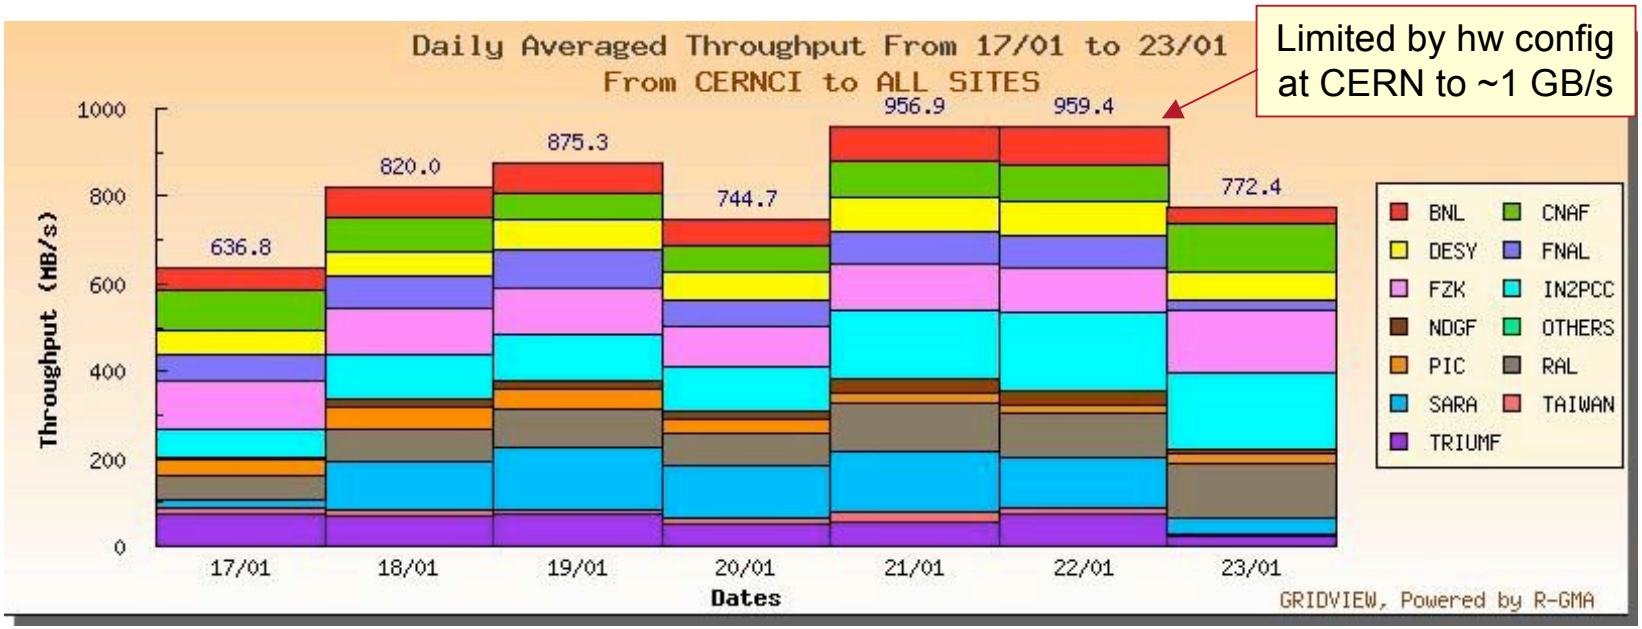

Example: CNAF inbound

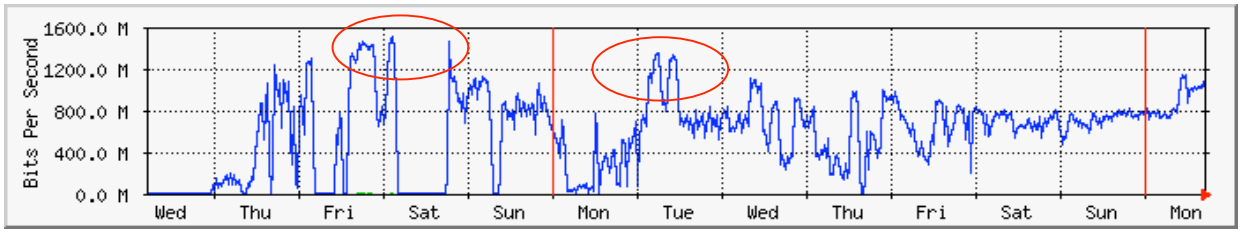

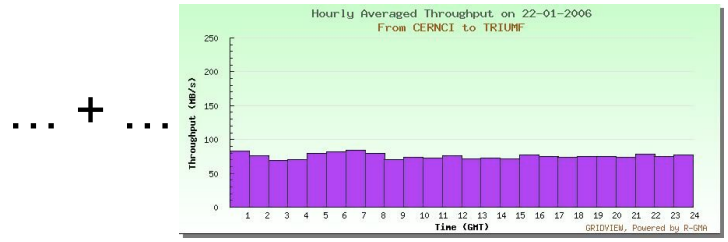

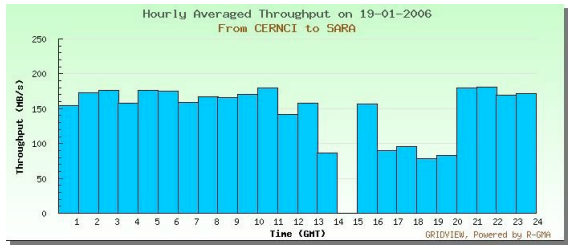

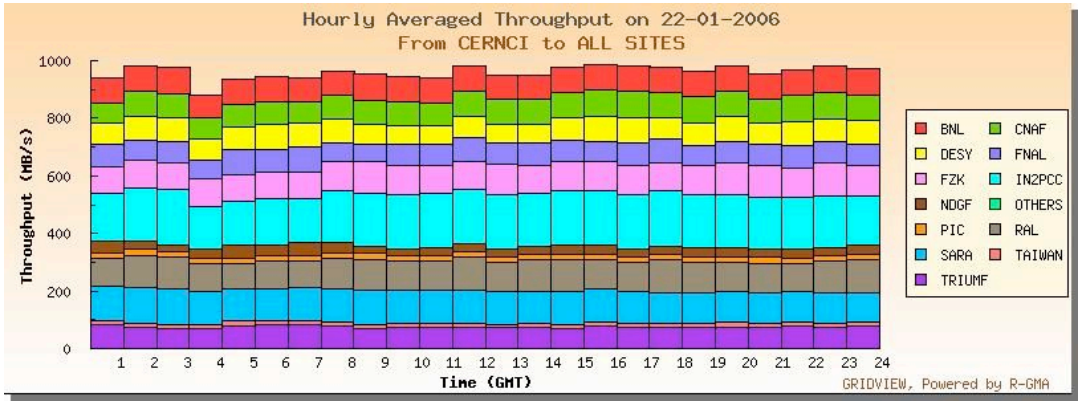

Day-view

Week-view

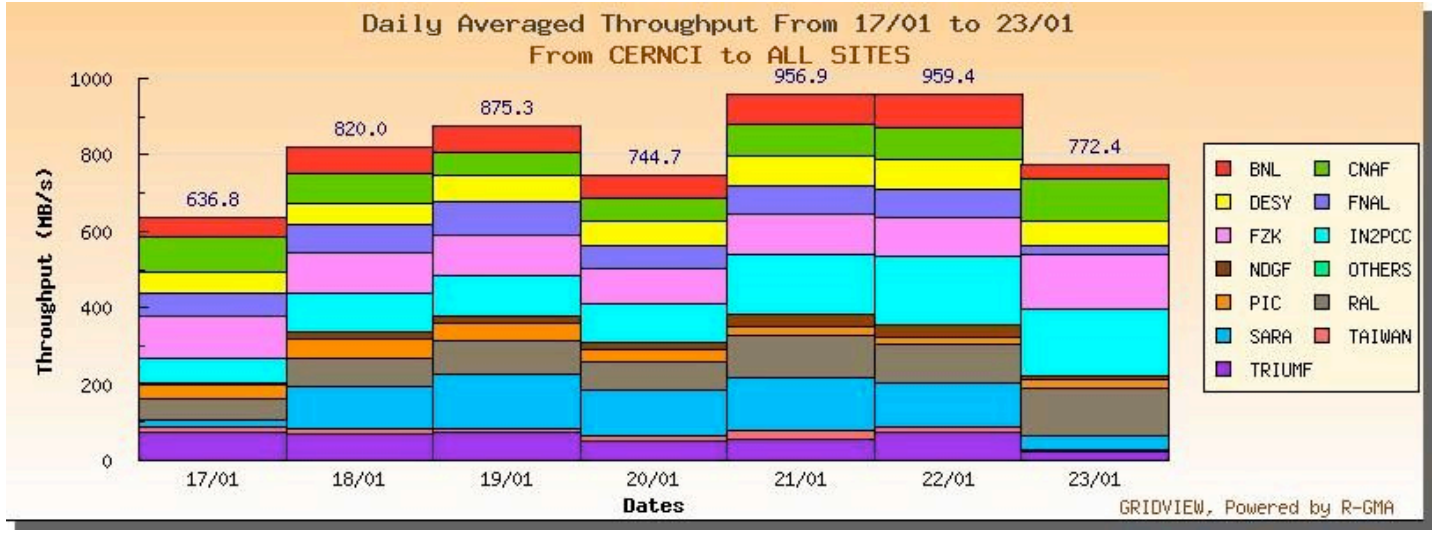

These rates must be sustained to tape 24 hours a day, ~100 days/yr. So: not appropriate to focus on ~% discrepancies**.** Still missing are: stability + extra capacity required to cater for outages / backlogs (namely running for 24-48h periods at ~2 x nominal). Still in early 2006, DM as a big WLCG challenge.

Other routes: T1↔T1 --- example: ATLAS

Take ATLAS 'average' T1 data flow as example

highest inter-T1 rates due to multiple ESD copies

Reprocessing ~1 month after data taking (better calibs) + at the end of yr with better calibs+algos

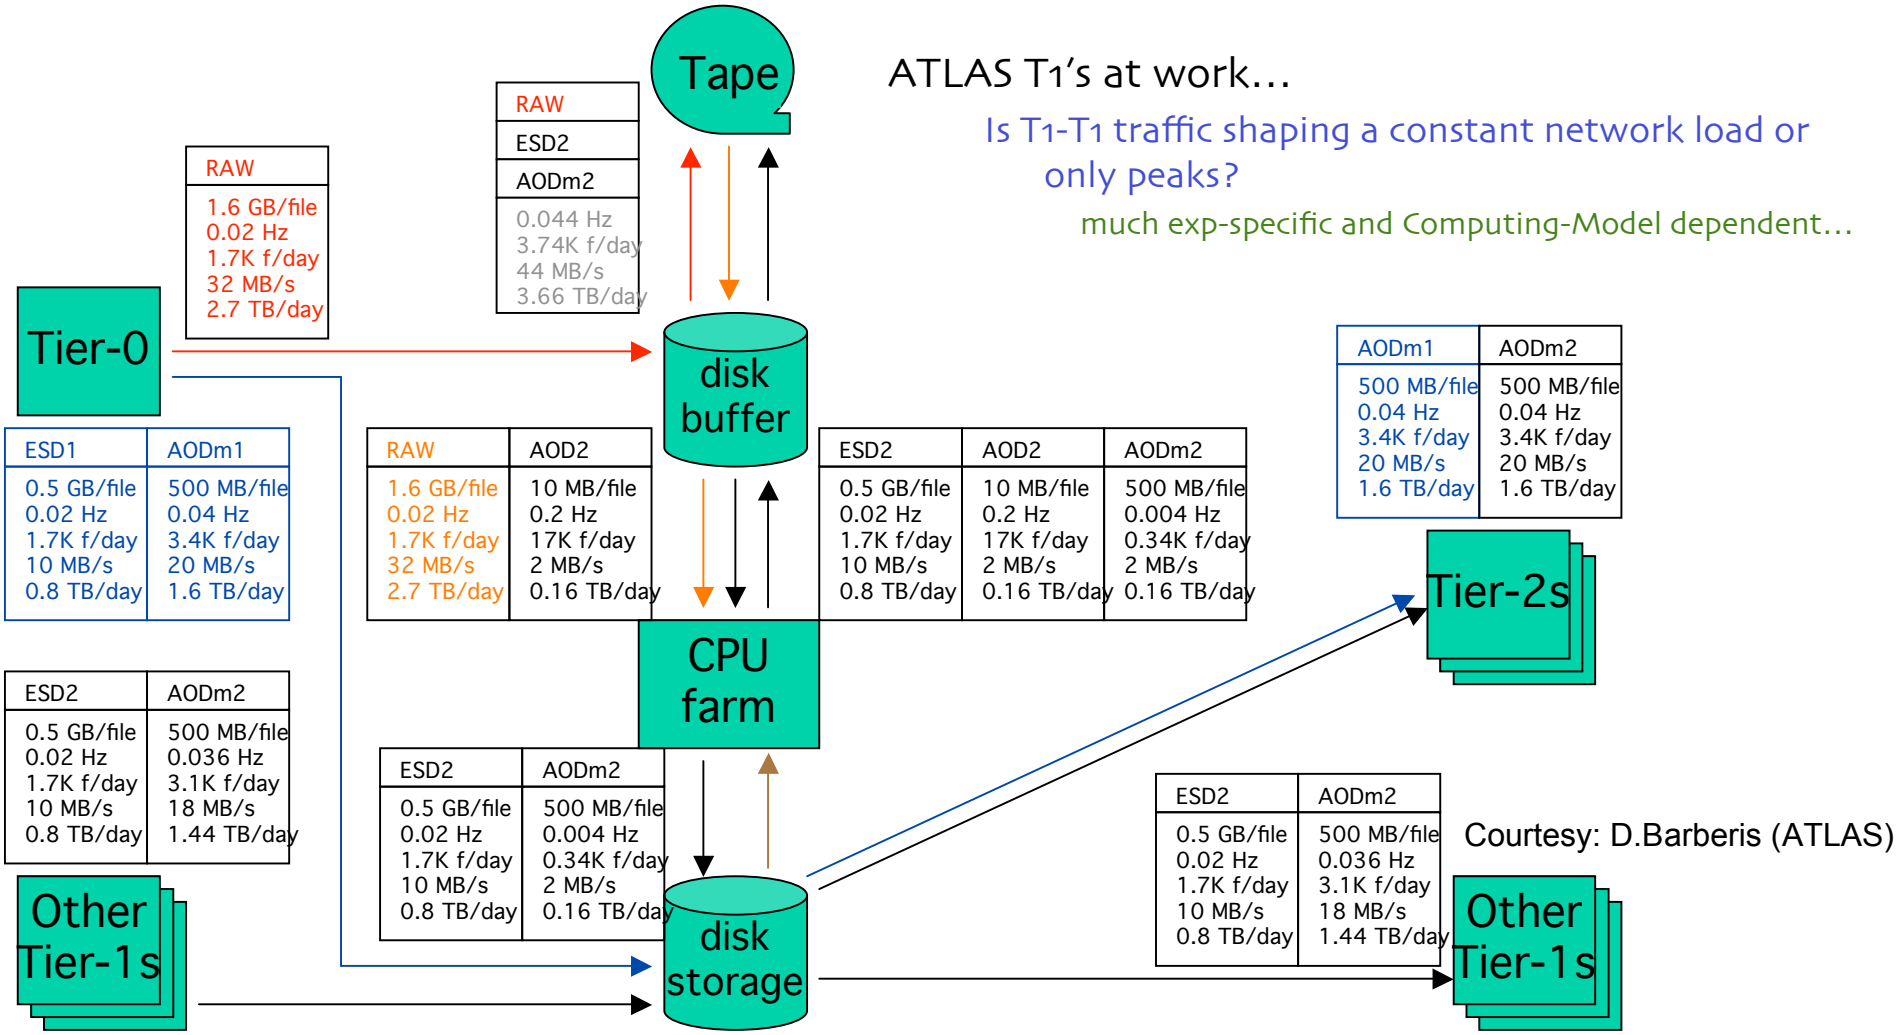

ATLAS T1's at work…

Is T1-T1 traffic shaping a constant network load or only peaks?

much exp-specific and Computing-Model dependent…

Courtesy: D.Barberis (ATLAS)

Other routes: T1↔T2 --- example: CMS

Take CMS data flow as example

- CMS T1→T2 transfers likely to be bursty and driven by analysis demands
  - ~10/100 MB/s for worst/best connected T2's (2006)
- CMS T2→T1 transfers almost entirely fairly continuous simulation transfers
  - ~ 10MB/s (NOTE: aggregate input rate into T1's comparable to rate from To!)

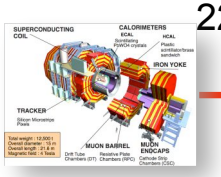

tape

280 MB/s
(RAW, RECO, AOD)

225 MB/s
(RAW)

Tier-0
4.6 M-SI2K
0.4 PB disk
4.9 PB tape
5 Gbps WAN

280 MB/s
(RAW, RECO, AOD)

Tier-1s

225 MB/s
(RAW)

WNs

40 MB/s
(RAW, RECO, AOD)

Tier-0

Tier-1s

AOD

AOD

Tier-1
2.5 MSI2K
1.0 PB disk
2.4 PB tape
~10 Gbps WAN

48 MB/s
(MC)

Tier-2s

240 MB/s
(skimmed AOD,
Some RAW+RECO)

~800 MB/s
(AOD skimming,
data reprocessing)

WNs

60 MB/s
(skimmed AOD,
Some RAW+RECO)

Tier-1

12 MB/s
(MC)

Tier-2
0.9 M-SI2K
0.2 PB disk
1 Gbps WAN

Up to 1 GB/s
(AOD analysis,
calibration)

WNs

NB: averaged sustained throughputs.

Courtesy: J.Hernandez (CMS)

Service Challenge 4 (SC4)

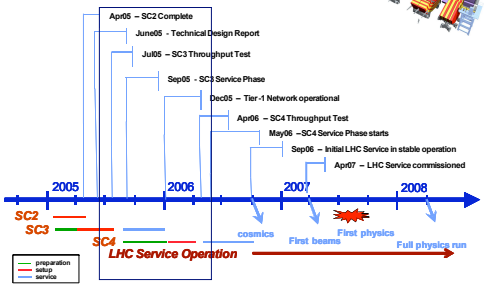

Aims to demonstrate that all of the offline data processing requirements expressed in the exps' Computing Models, from raw data taking through to data access, can be handled within Grid at the full nominal data rate of the LHC

when: Apr 06 - Sep 06

who: To, all T1's, majority of T2's

It will become the initial production service for LHC and made available to the exps for final testing, commissioning and processing of cosmic ray data

Set-up ("throughput") phase (1 month: Apr 06)

Throughput demonstration sustaining for 3 weeks the target data rates at each site

Target is a stable, reliable data transfer to T1's at target rates to any supported SRM implementation (dCache, Castor, …) + factor 2 for backlogs/peaks.

Service phase (5 months: May-Sep 06)

get the basic sw components required for the initial LHC data processing service into the loop

Target is to show capability to support full Computing Models of each LHC exp, from simulation to end-user batch analysis at Tier-2's

Example: SC4-phase1 disk-tape rates

Centre	*ALICE*	*ATLAS*	*CMS*	*LHCb*	*Target Data Rate MB/s*
Canada, TRIUMF		X			50
France, CC-IN2P3	X	X	X	X	75
Germany, GridKA	X	X	X	X	75
Italy, CNAF	X	X	X	X	75
Netherlands, NIKHEF/SARA	X	X		X	75
Nordic Data Grid Facility	X	X	X		50
Spain, PIC Barcelona		X	X	X	75
Taipei, ASGC		X	X		75
UK, RAL	X	X	X	X	75
USA, BNL		X			75
USA, FNAL			X		75

Apr 06: still using SRM 1.1 & current tape technology…

[courtesy of L.Tuura, D.Bonacorsi (CMS)]

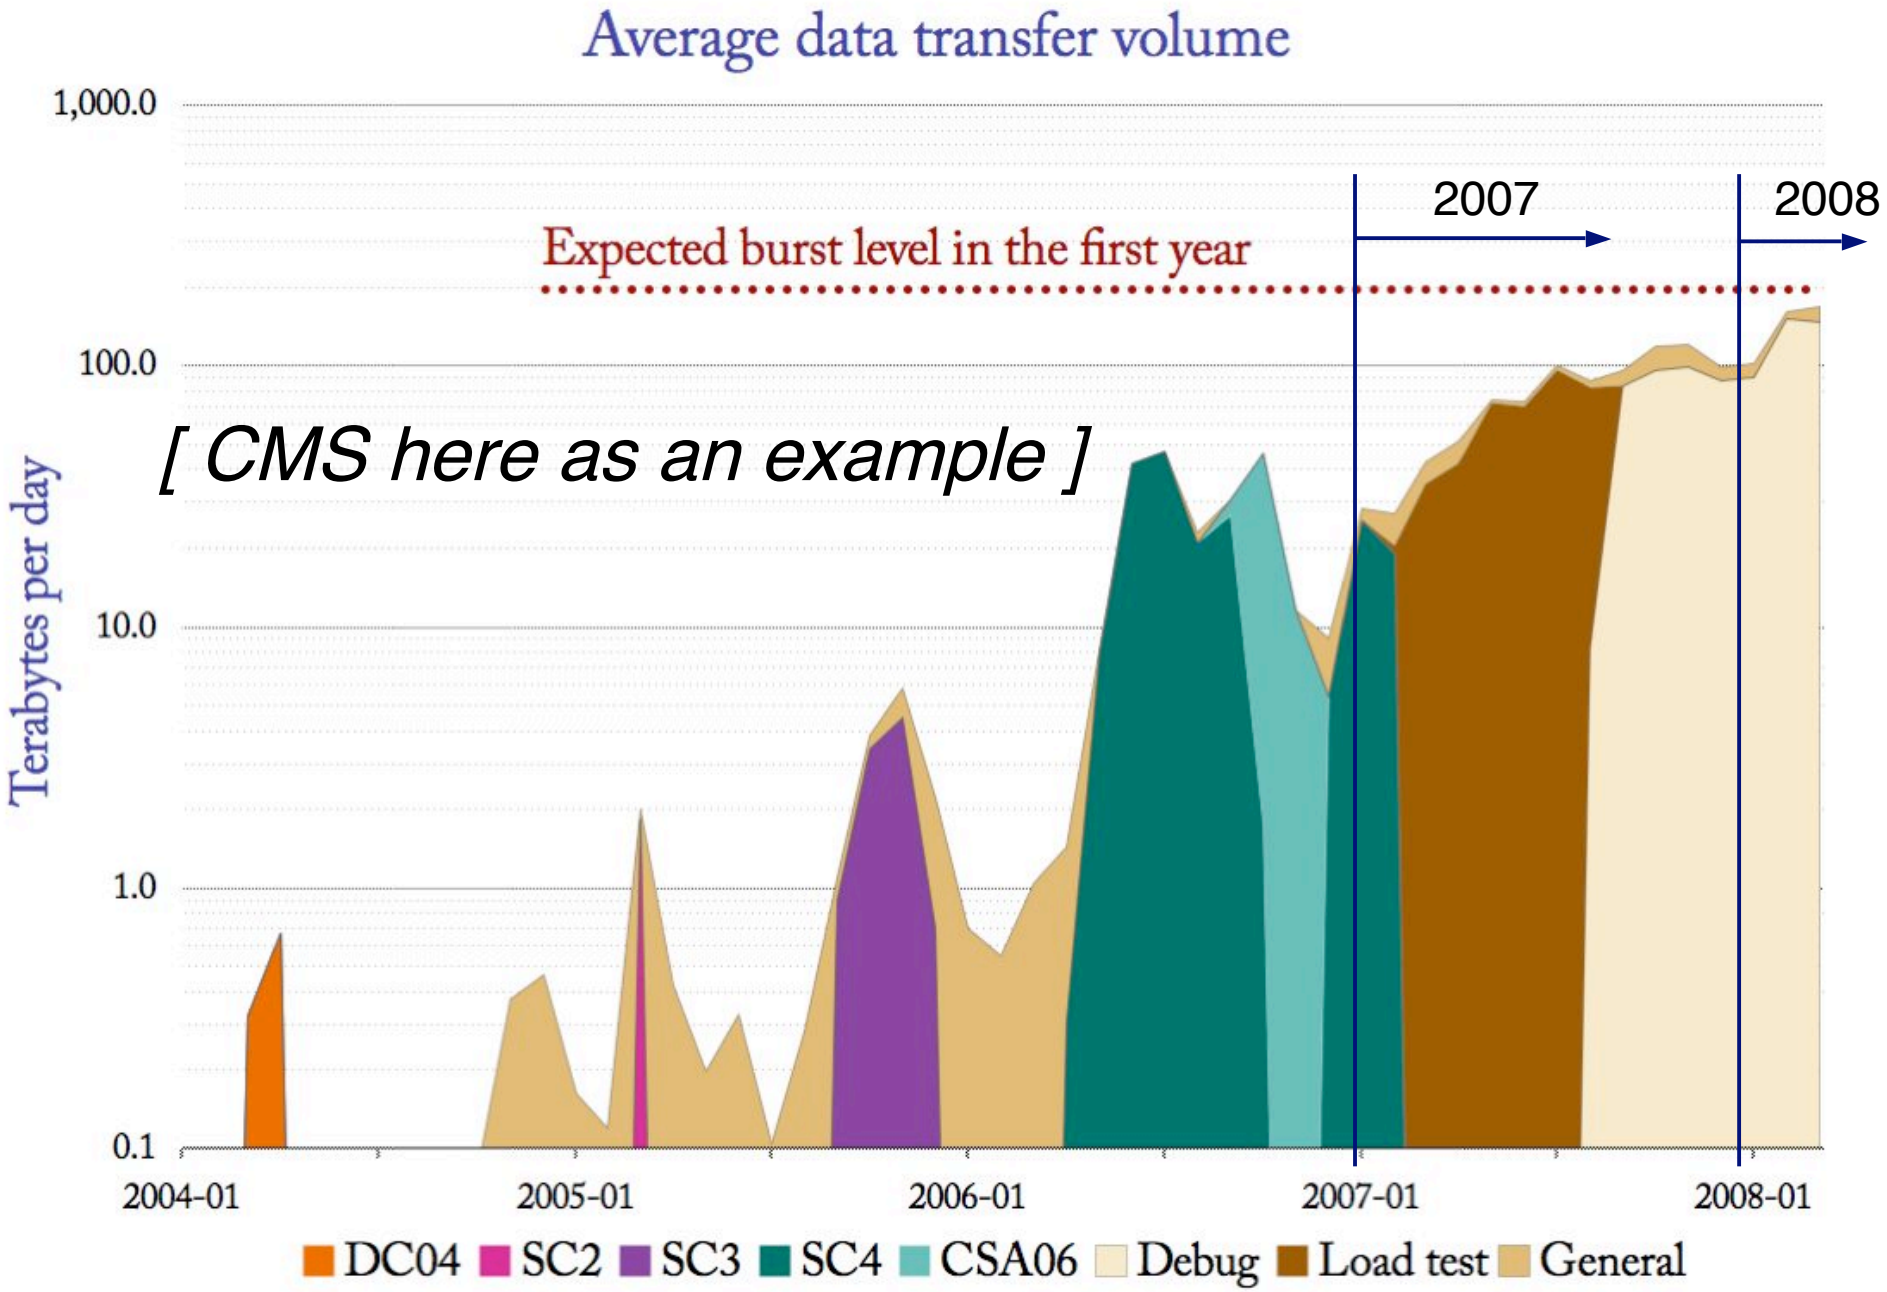

CCRC'08

WLCG Common-VO Computing Readiness Challenge

- all 4 LHC Vos
- 2 phases:
  - Phase-1: functionality and performance tests in Feb (4^{th} - 29^{th})
  - Phase-2: 4 week challenge in May (5^{th} - 30^{th})

High complexity!

- In itself
- Because of MANY other experiment-specific challenges going on simultaneously!
  - e.g. MD5/6 for ATLAS, CSA08 for CMS, …

Metrics pre-FDR

- **Week 4: Jan 21 – 27**
  - ✓ Setting up sites with SRMv2 production endpoints
  - ✓ Testing central bulk data deletion at sites
  - ✓ RDO file mixing into Byte Stream data
  - ✓ T0 – T1 transfers, storage space can be re-cycled
  - ✓ *Milestone 2: T0 and T1 space tokens tested*
- **Week 5: Jan 28 – Feb 3**
  - ✓ Setting up sites with SRMv2 production endpoints
  - ✓ Testing central bulk data deletion
  - ✓ RDO file mixing into Byte Stream data
  - ✓ T0 – T1 transfers, storage space can be re-cycled
  - ✓ Upload Byte Stream data into the SFO's

- **Week 4: Jan 21 – 27**
  - ✓ Many ToA changes
  - ✓ Not started yet
  - ✓ Not started yet
  - ✓ Wait for srmv2
  - ✓ *Milestone 2: failed*
- **Week 5: Jan 28 – Feb 3**
  - ✓ Many more ToA changes
  - ✓ Simple deletions
  - ✓ started
  - ✓ started
  - ✓ Not started yet

Ref. ATLAS Planning for FDR-1 and CCRC-1 vs.1 Jan.28
http://www.nikhef.nl/~bosk/documents/d2008/ATLAS-Document-for-FDR-and-CCRC.pdf

Metrics week 1&2

CCRC'08/1 ATLAS

- **Week 6: Feb 4 – 10**
  - ✓ Streaming Data read from SFO's into CASTOR
  - ✓ Full T0 operation:
    - Data stored onto tape
    - Calibration
    - First pass reconstruction
    - Data subscriptions

    T0 – T1 transfers, storage space re-cycled
- *Milestone 4: Data streaming from SFO's tested*
- *Milestone 5: Full T0 operation tested*
- **Week 7: Feb 11 – 17**
  - ✓ Without SFO's from now on
  - ✓ Full T0 operation from streaming data on CASTOR
  - ✓ Data subscriptions to Tier-1's and Tier-2's
  - ✓ Test SRMv2.2 bring-online functionality in at 2 sites
  - ✓ Central bulk data deletion on the last day
- *Milestone 6: Data subscription to all Tiers and storage classes tested*

- **Week 6: Feb 4 – 10**
  - ✓ Done
  - ✓ Full T0 operation: Done
    - ✓ Done
    - ✓ Done
    - ✓ Done
    - ✓ Done
    - ✓ Done
- *Milestone 4: Done*
- *Milestone 5: Done*
- **Week 7: Feb 11 – 17**
  - ✓ True
  - ✓ Done
  - ✓ Done
  - ✓ Done
  - ✓ Requested to leave the data
- *Milestone 6: Done*

Metrics week 3&4

- **Week 8: Feb 18 – 24**
  - ✓ Full T0 operation from streaming data on CASTOR
  - ✓ Data subscription to all Tiers and storage classes
  - ✓ Test SRMv2.2 bring-online functionality
  - ✓ Group Analysis in at least two Tier-2's
  - ✓ DPD production in at least two Tier-1's
  - ✓ Central bulk data deletion on the last day
- *Milestone 7: staging of data from tape tested in 2 Tier-1's tested*
- *Milestone 8: DPD production in Tier-1 tested*
- *Milestone 9: Group analysis in Tier-2 tested*

- **Week 9: Feb 25 – 29**
  - ✓ Full T0 operation from streaming data on CASTOR
  - ✓ Data subscription to all Tiers and storage classes
  - ✓ Re-processing in 2 Tier-1 sites
  - ✓ Group Analysis in Tier-2's
  - ✓ DPD production in Tier-1's
- *Milestone 10: Re-processing in two Tier-1's tested*

- **Week 8: Feb 18 – 24**
  - ✓ Done. but with other data
  - ✓ Done
  - ✓ Done
  - ✓ Started
  - ✓ Not done
  - ✓ Done but done dayly
- *Milestone 7: done*
- *Milestone 8: not done*
- *Milestone 9: started*

- **Week 9: Feb 25 – 29**
  - ✓ Done but with different data
  - ✓ done
  - ✓ Done, even at 3
  - ✓ started
  - ✓ Not done
- *Milestone 10: done*

Ref. ATLAS Planning for FDR-1 and CCRC-1 vs.1 Jan.28
http://www.nikhef.nl/~bosk/documents/d2008/ATLAS-Document-for-FDR-and-CCRC.pdf

Lessons learned

- CCRC1 was a very useful exercise for ATLAS
- Achieved most milestones in spite of external dependencies
- The small size of the FDR data sample was a complication
- It's difficult to serve the Detector and Physics and IT community
- It is good that FDR2 does not coincide with CCRC2
- A T2 workshop in Moscow was a major disruption
- The same holds for S&C workshop: we don't have enough people
- We did not want to stop but carried on March + April with CCRC 1.5
- Should be better prepared for CCRC2

Summary

- For ALICE, the February/March phase of CCRC'08 exercise was very useful
  - Focus on data management
  - Brings all experiments together
  - Controlled tests, organization
- The focus of the May phase and beyond will continue to be data management
  - At the T1s - mostly deployment of additional capacity
  - At the T2s – more complex, involves a period of tuning for every new T2

WLCG Collaboration Workshop

CCRC'08/1
ALICE

Summary (2)

- In general, the data registration, transfer and processing tools worked well
  - Many thanks to the experts at CERN, T1s and T2s for the close support, feedback and suggestions
- ALICE is continuing to re-process the RAW data taken in Feb/March
  - In addition to the standard MC production and user analysis
- Plans and resources requirements for the May exercise will be presented tomorrow

WLCG Collaboration Workshop 25

CCRC'08/1 LHCb

- Initial phase of CCRC'08 was dedicated to development and testing of DIRAC3 + interaction with SMR2
- CCRC'08 ran smoothly after initial tests
  - Online->T0 and T0-T1 transfers on the whole a success
  - Some minor issues with reconstruction activity and data upload from the WNs
  - Quick turnaround for reported problems
    - Very good responsiveness from site admins
  - Quick development and deployment of new brand of middleware or bug fixes
    - Very good responsiveness from service coordinators and software providers
- Preparation underway for 4 weeks steady running at nominal rate in May
    - Missing components (Stager/Failover/GANGA-DIRAC3 integration/ TimeLeft, Stripping workflow

LHCb

R.Santinelli WLCG T0/T1/T2 Workshop

CCRC'08/1
CMS

Outlook

CMS is in continuous production mode

Phase-1 was a productive (set of) exercise(s) for CMS

More is needed, and we look fwd to phase-2

The 'functional blocks' tests in CCRC/phase-1 were scheduled to overlap with other VOs

Need to push more on this for the May phase

Since then, CMS continued with production and tests

The CCRC08/phase-2 exercise starts now with the pre-production using all available resources, and running this week

[see also the CMS talk tomorrow]

WLCG Collaboration Workshop - CERN, 21-25 Apr 08 D. Bonacorsi 70

CCRC'08

Phase-1 | Phase-2

week-1

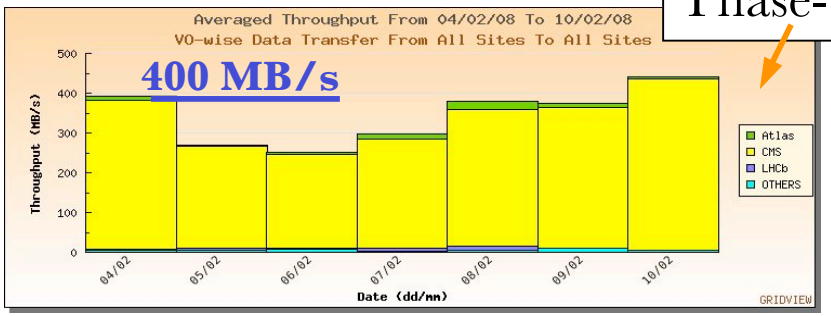

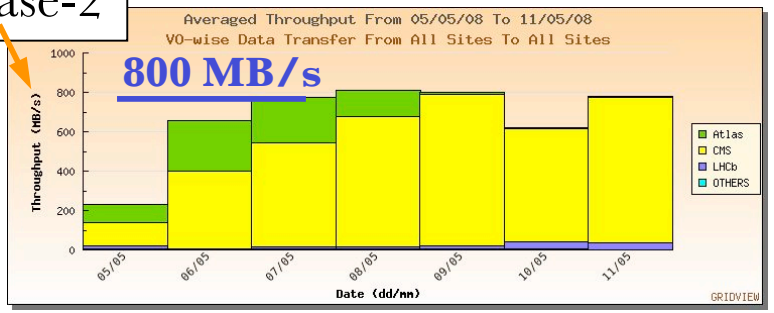

week-2

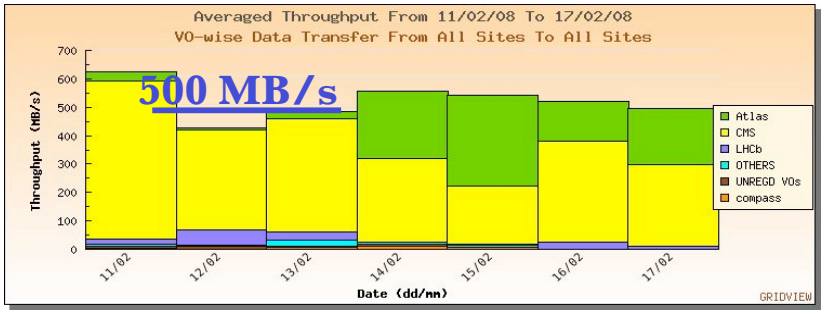

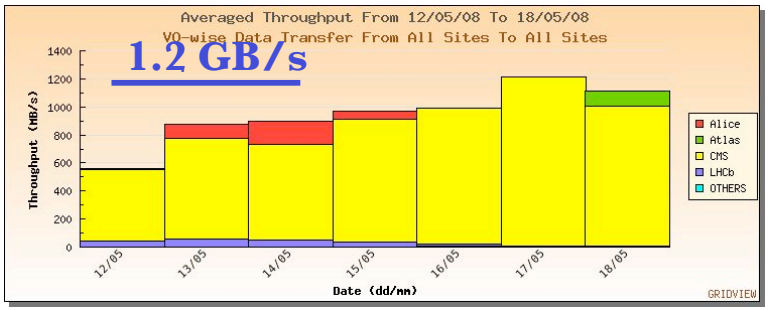

week-3

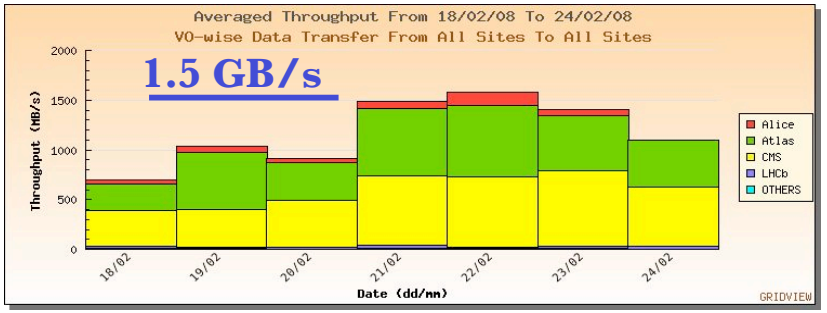

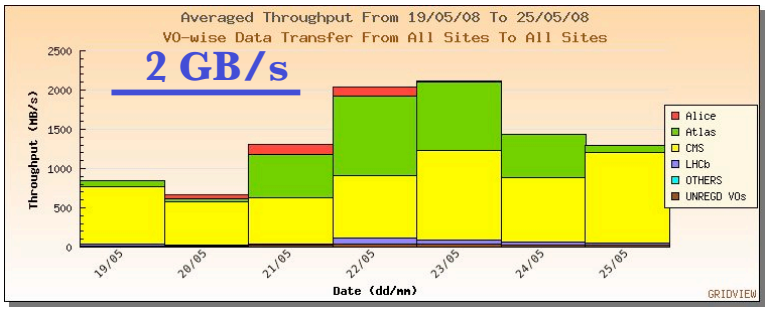

week-4

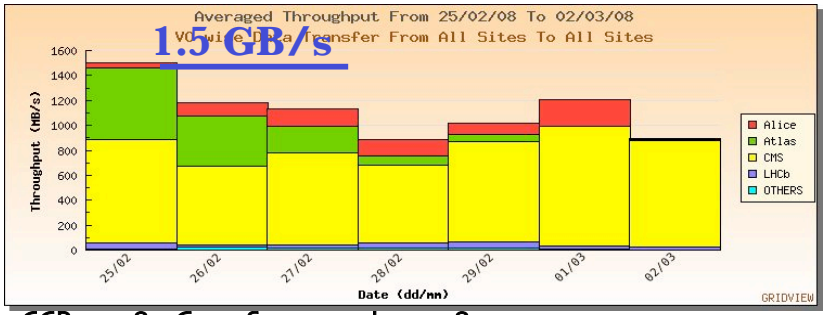

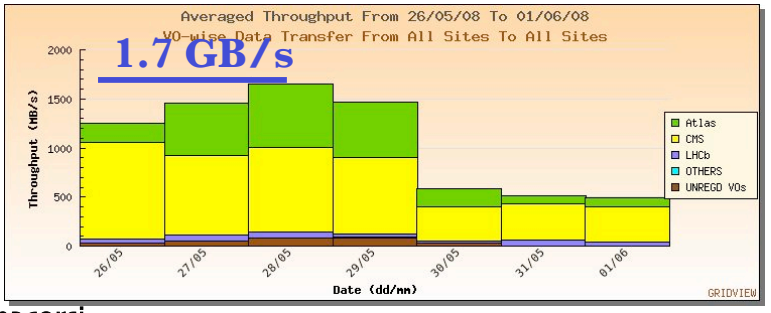

CCRC '08 – How Does it Work?

- Experiment "shifters" use Dashboards, experiment-specific SAM-tests (+ other monitoring, e.g. PhEDEx) to monitor the various production activities
- Problems spotted are reported through the agreed channels (ticket + elog entry)
- Response is usually rather rapid – many problems are fixed in (<)< 1 hour!
- A small number of problems are raised at the daily (15:00) WLCG operations meeting

Basically, this works!

- We review on a weekly basis if problems were not spotted by the above → fix **[+ MB report]**

With time, increase automation, decrease eye-balling

(by J.Shiers, end of May 2008)

CCRC '08 – Areas of Opportunity

- Tier2s: MC well run in, distributed analysis still to be scaled up to (much) larger numbers of users
- Tier1s: data transfers (T0-T1, T1-T1, T1-T2, T2-T1) now well debugged and working sufficiently well (most of the time…); reprocessing still needs to be fully demonstrated for ATLAS (includes conditions!!!)

➢ **Tier0: best reviewed in terms of the experiments' "Critical Services" lists**

- These **strongly emphasize** data/storage management and database services!
- **We know how to run stable, reliable services**
- IMHO – these take **less** effort to run than 'unreliable' ones…

➢ **But they require some minimum amount of discipline…**

(by J.Shiers, end of May 2008)

CCRC'08 – Conclusions

- The **WLCG** service is running (reasonably) **smoothly**
- The functionality **matches:** what has been tested so far – and what is (known to be) required **&** the experiments are **happy**!
- **We have a good baseline on which to build**
- (Big) **improvements** over the past year are a good indication of what can be expected over the next!
- (Very) detailed **analysis** of results compared to up-front metrics – in particular from experiments!

(by J.Shiers, end of May 2008)

Summary

WLCG:

- more and more ready to provide infrastructure/components to LHC exps

Tiers:

- more and more ready to operate WLCG-enabled services to realize the Computing Models of the LHC exps

WLCG is approaching the 'regime' through several challenges involving Tiers

- We indeed needed to practice repeatedly, and expand the scope of tests, 'SC after SC'… so, *thanks* to WLCG!

The infrastructure is basically prepared

Work is still needed on a robust services' delivery

Tests could now change perspective

- analyse the still-unsuccessful stories…

The WLCG Service

- Is a combination of services, procedures, documentation, monitoring, alarms, accounting, … and is based on the "WLCG Service Model".
- It includes both "experiment-specific" as well as "baseline" components
- No-one (that I've managed to find…) has a complete overview of all aspects
- ➢ **It is essential that we work together and feel joint ownership – and pride – in the complete service**
- It is on this that we will be judged – and not individual components…

(by J.Shiers, end of May 2008)

On WLCG Readiness

- The service runs smoothly – most of the time
- Problems are typically handled rather rapidly, with a decreasing number that require escalation
- Most importantly, we have a well-proved "Service Model" that allows us to handle anything from "Steady State" to "Crisis" situations
- We have repeatedly proven that we can – typically rather rapidly – work through even the most challenging "Crisis Situation"
- Typically, this involves short-term work-arounds with longer term solutions
- It is essential that we all follow the "rules" (rather soft…) of this service model which has proven so effective…

(by J.Shiers, end of May 2008)

WLCG Operations

I am sure that most – if not all – of you are familiar with:

- Daily "CERN" operations meetings at 09:00 in 513 R-068
- Daily "CCRC" operations meetings at 15:00 in 28 R-006 (will move to 513 R-068 if the phone is ever connected)

- In addition, the experiments have (twice-)daily operations meetings → continuous operations in control rooms

- These are pretty light-weight and a (long) proven way of ensuring information flow / problem dispatching

(by J.Shiers, end of May 2008)

Are We Ready?

- Based on the experience in the February and May runs of the Common Computing Readiness Challenge, an obvious question is

¿ "Are we ready for LHC Data Taking?"

  - any time mid-July 2008 on...

- The honest answer is probably:

- "We are ready to **face** LHC Data Taking"

- There are subtle, but important, differences between the two...

(by J.Shiers, end of May 2008)

Acknowledgements

This talk reflects the hard work of a large nb of people across many sites and within LHC exps Collaborations, so ALL credit goes not to speaker but to:

all people in WLCG Collaboration

all people in network, storage, farming, … at WLCG sites

all people (from exps and sites) specifically working on SCs/CCRC

people coordinating WLCG efforts