



CCRC08-1 report

WLCG Workshop, April 22 2008

KorsBos, ATLAS/NIKHEF/CERN

Use simulated data with physics mix of 10^{31} Luminosity

run 10 hrs @200 HZ from SFO's

Full T0 operation:

Calibration, calibration database may take 24 hrs

Data quality, sign off end of the next day

First pass reconstruction may take 24 hrs

If necessary re-calibration may take another 24 hrs

Subscribe data for transfers to T1&2's

We had planned to concentrate on T0 operations first

but it took longer than was expected

WEEK 1 what we also did

Setup of ATLAS site services

- Storage spaces definitions in TiersOfATLAS

- usage of FTS delegation

Testing of SRM2.2 endpoints

- Space tokens, ACLs

- BDII publication

- FTS channels configurations

Several configuration issues spotted and cured.

WEEK 5

switched to other data sample

QDR sample too small for TT

Used generated files: nominal size

Ramped up to nominal rates

Full Computing Model, MoU shares

Target Rate: ~900 MB/s

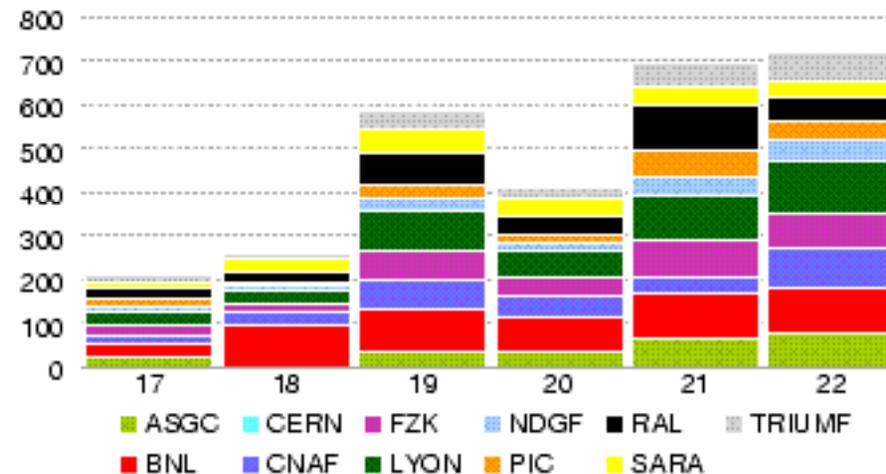
Full steam for 24h

Relatively good throughput achieved

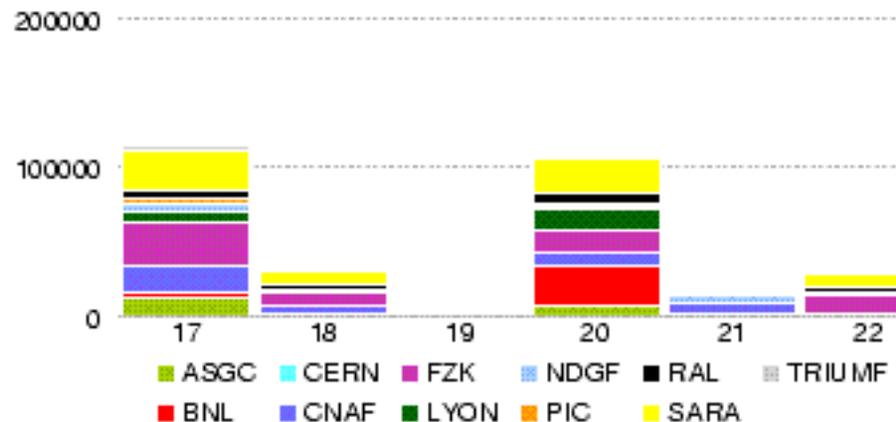
Sustained 700 MB/s for 2 days

Peaks above 1.1GB/s for several hours

Throughput



Errors

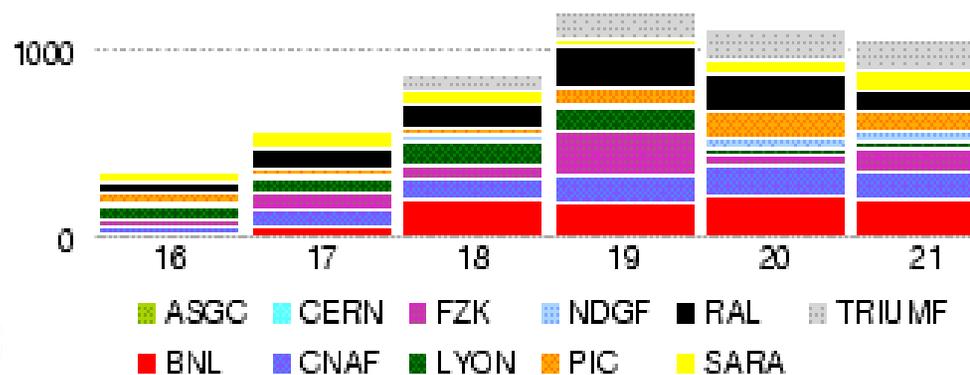


Work comparison

What happened to THIS?

T0 exercise in September

Much more precise in data transfers



The T0 buffer is recycled after 1 day, now after 2.5 days

If data did not make it in 1 day ... tant pis!

In the past, datasets were cleaned daily from T0

➔ now we control the cleaning

In September we have been “oversubscribing” data : ESD everywhere

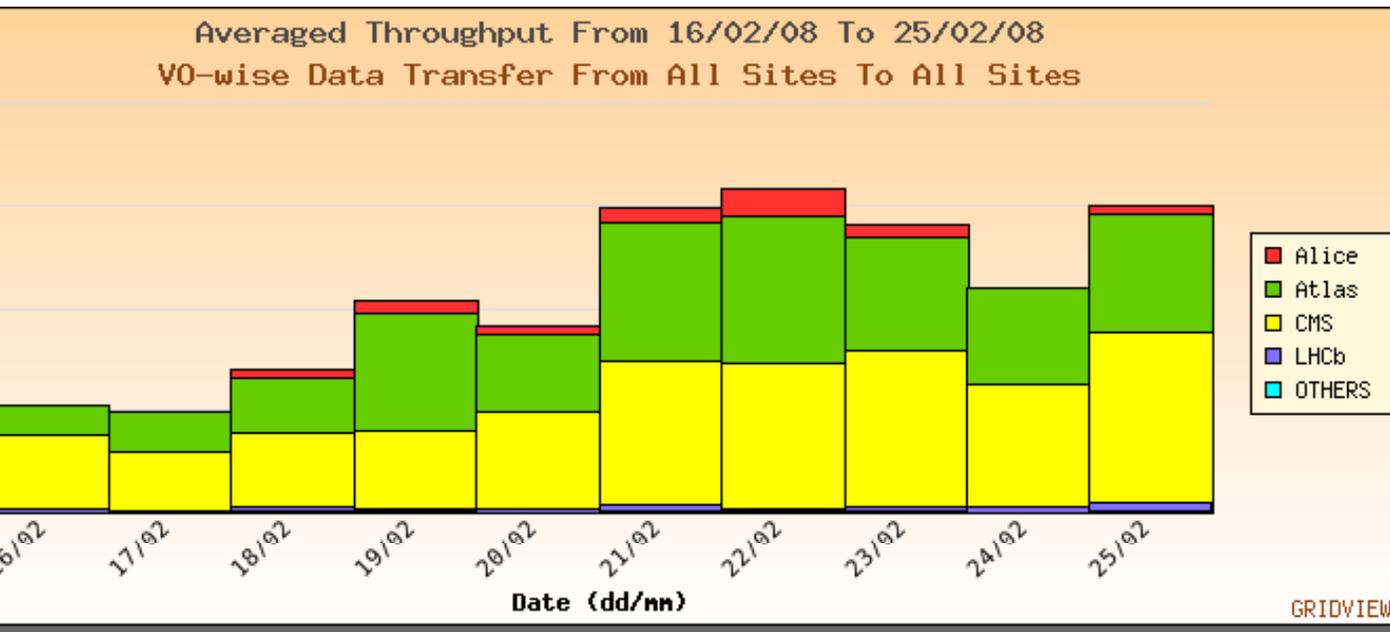
➔ Now we strictly followed the CM and the MoU ratio's

ites request more TT tests

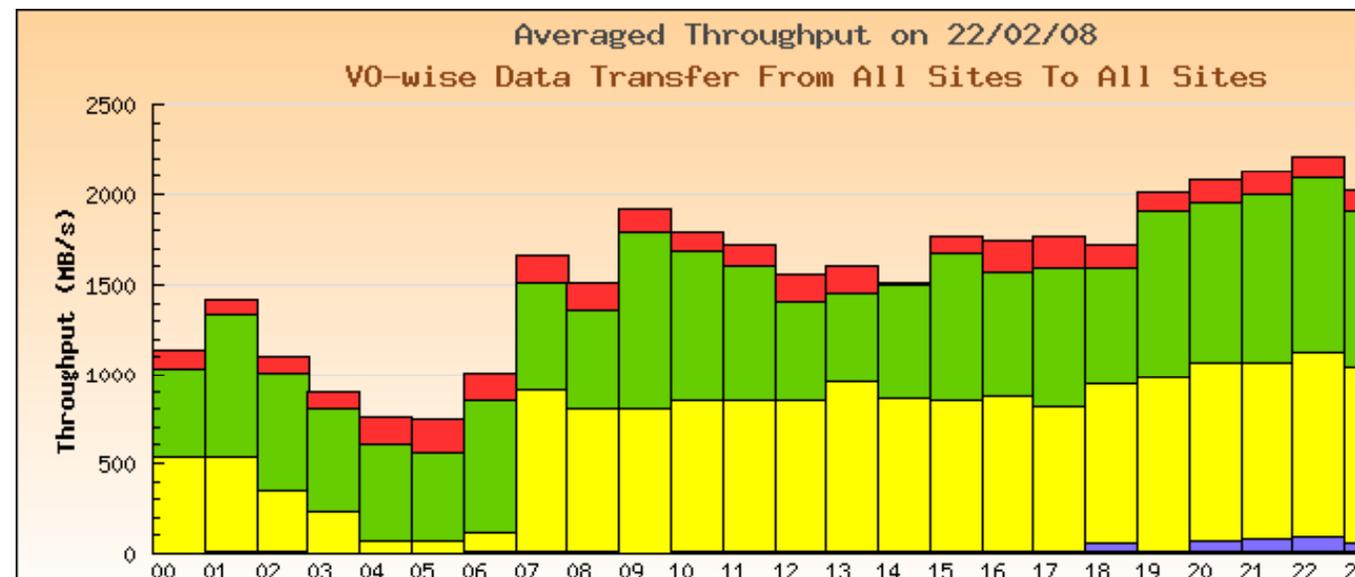
We want to do more FT tests

You'll never walk alone

Weekly Throughput



1 GB/s out
of CERN



FTS transfers

NL	FZK	INFN	IN2P3	NDGF	NIKHEF	PIC	RAL	TAIWAN	TRIUMF
X	100	100	20	100	25	15	20	100	100
X	X	100	30	100	40	50	30	100	100
100	90	X	0	100	25	30	X	100	100
X	100	100	X	100	35	20	30	100	100
100	100	100	0	X	40	50	30	100	100
0	0	0	0	0	X	0	0	0	0
5	100	95	3	100	15	X	5	90	95
00	100	100	10	0	30	10	10	100	100
100	100	100	25	0	40	70	80	X	100
100	100	30	20	100	100	30	30	X	100

Varying results, not all understood

✓ NIKHEF unable to export data (understood?)

Several issues of FTS channels configuration

✓ Many still being investigated (e.g. NDGF exporting via SRM1)

T1-T2 replication within cloud

Ongoing exercise also in production, see next slide

Replication of M5 data to T1 tapes

Needed for re-processing exercise

Re-processing at T1's: Triumf, NDGF, BNL

DR-2 simulation production validation started

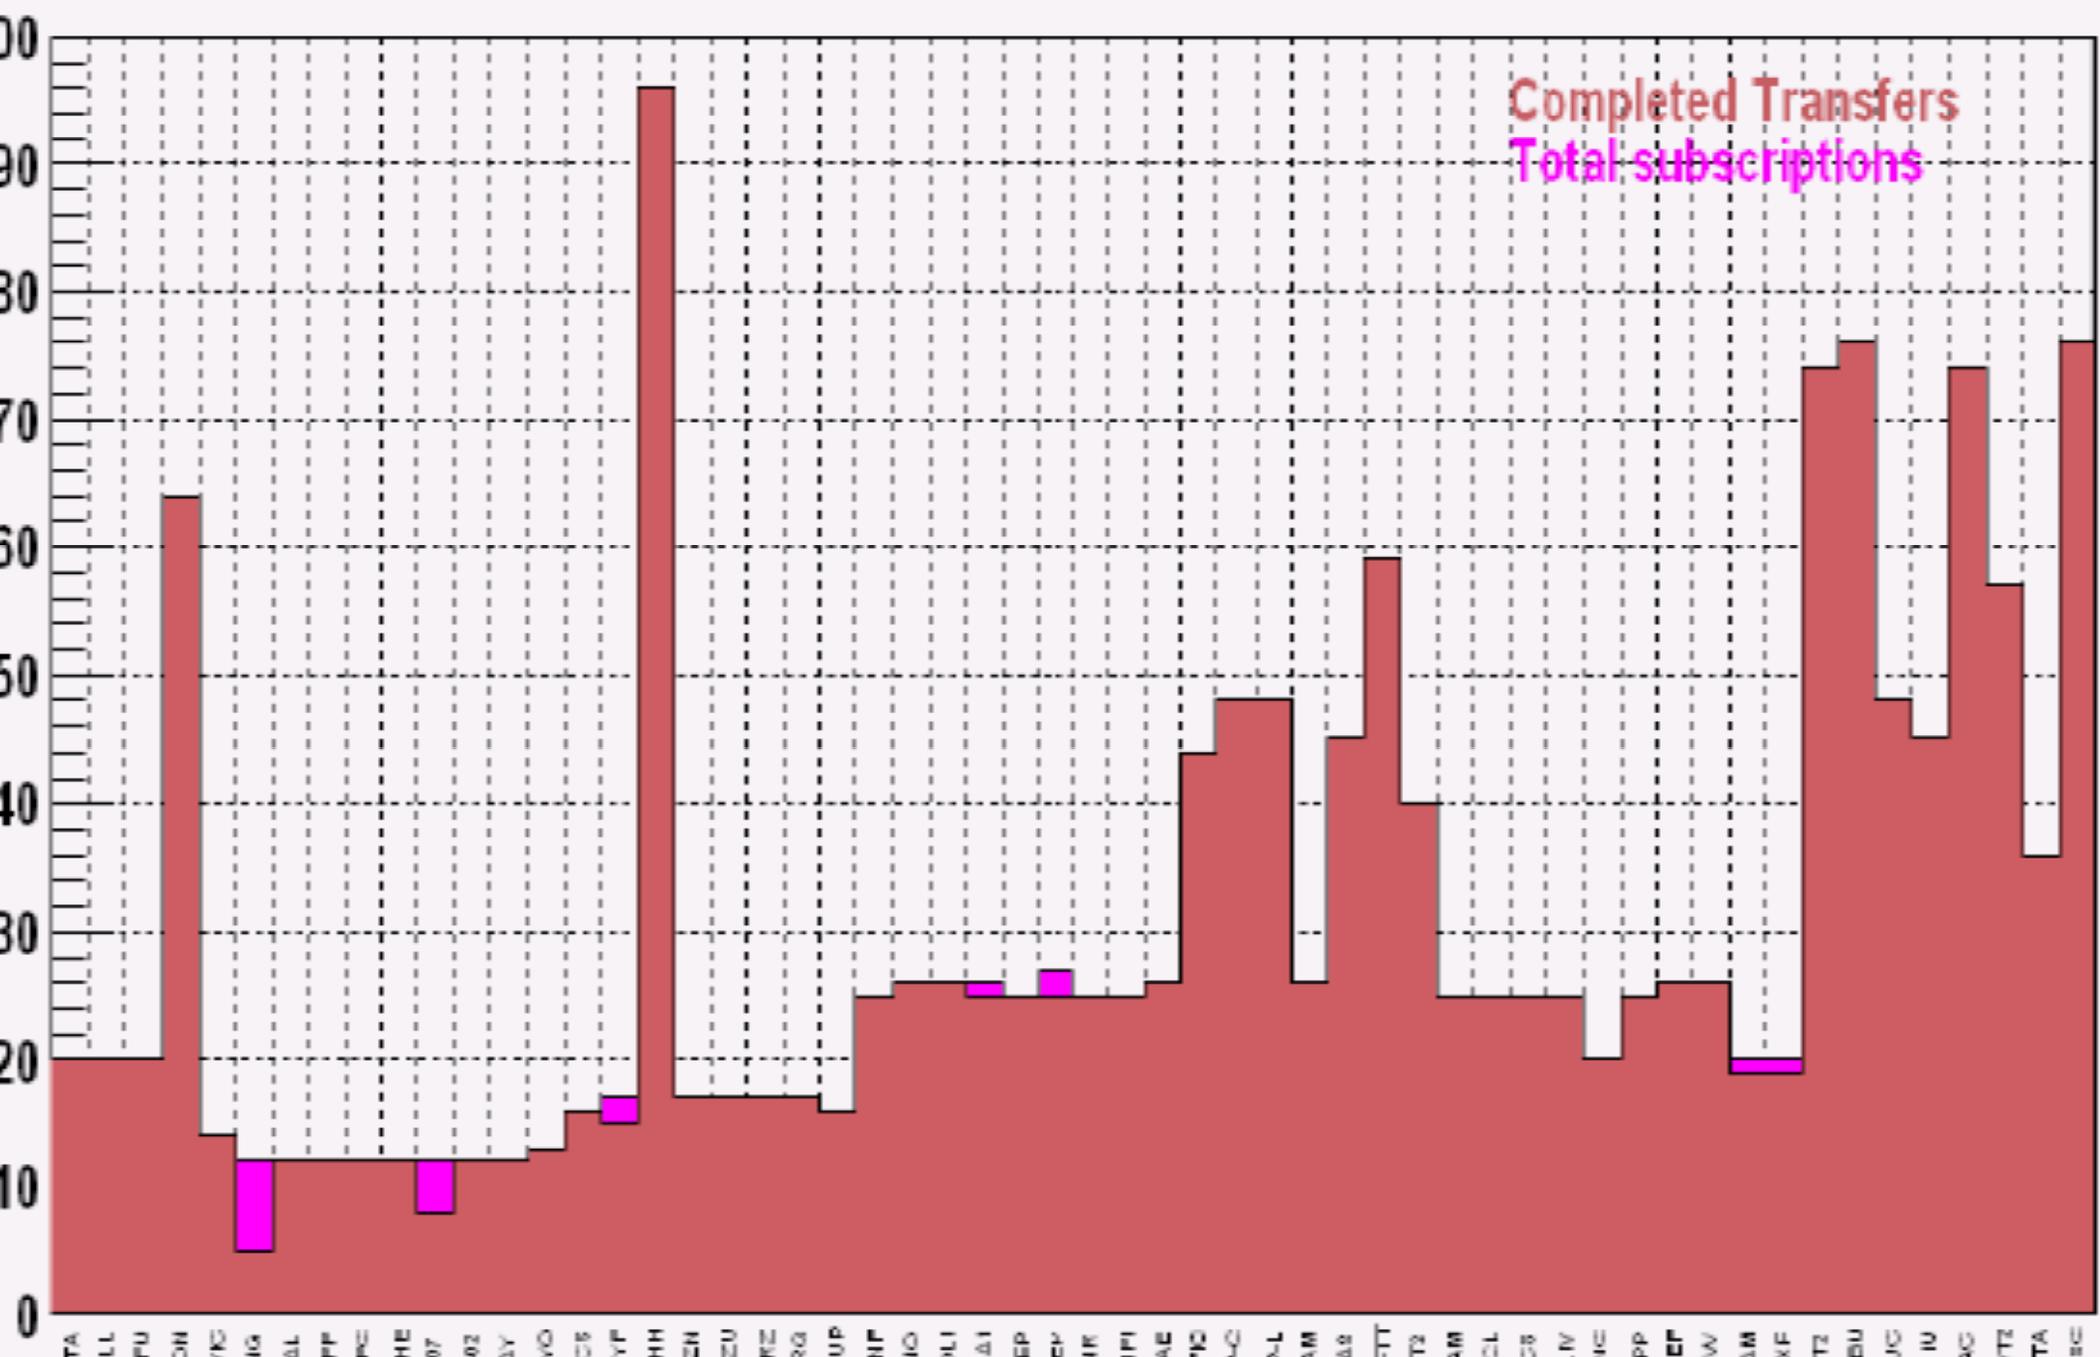
T0→T1 transfers stopped on Thursday because of M6 preparation

Clear CCRC data from all sites

Set up and test T0-T2 channels for Muon calibration

Re-arrange CERN Castor pools

FDR AOD datasets replication within clouds (56 Tiers)



RM2.2. general considerations

RM2.2 testing was (one of the) main goals of ATLAS CC

Delivered “just in time”

➔ Most of the sites delivered during last pre-FDR week

But, still, in time

MHO, it has been working very well

No major problems observed in all implementations

➔ Various bugs fixed promptly.

Generally good storage setup at sites

Various issues in the interaction between different components (FTS-SRM) have been identified and overcome

space tokens at T1s

ATLASDATADISK and ATLASDATATAPE correctly configured in all T1s

➔ ATLASMCDISK and ATLASMCTAPE begins being tested

Should start thinking about T1D1

➔ ATLASDATADISKTAPE and ATLASMCDISKTAPE

➔ It is clear how this will be implemented at every T1?

Put more pressure on group analysis at T1

➔ ATLASGRM<group> need to be setup for a (very) few groups

space tokens at T2s

Many T2s are still on SRMv1 (50%)

SRMv2 T2s set up ATLASDATADISK and ATLASMCDISK

Space Tokens and Disk Pools

CRC08 Storage Solution Working Group

Technical discussion about prioritized evolution of SRM

➔ <http://indico.cern.ch/conferenceDisplay.py?confId=29191>

On top of the list every experiment put “protection of space tokens from generic user access, based on VOMS groups and roles”

➔ CASTOR and dCache

This is true for both “READ” and “WRITE”

timescale of implementation is by CCRC09

How to protect in the meanwhile?

➔ Write+Read: @T0, static DN(now Group)-to-pool mapping. @T1s?

➔ Read: at T0, interesting “service class fallback” capability

Communication Channels

LOG: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>

Found very useful.

➔ Need an ATLAS specific one also?

But not a ticketing system

Email your favorite T1 (or T0) contact

Many times very effective. But your favorite contact might be skiing on the Alps

GUS is a ticketing system

But the latency is too high

Need to define a “hotline”

At CERN: atlas-grid-alarm@cern.ch is seen by the 24 hour operator who may escalate

➔ Very restricted list of people can post (need to revisit)

Asked Jamie to also discuss this with T1s

Metrics pre-DR

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

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*

MILESTONES WEEK 3Q4

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

✓ **Dones**

✓ **Done**

✓ **Started**

✓ **Not done**

✓ **Done** but done daily

- *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*

Lessons learned

CRC1 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 communities

it is good that FDR2 does not coincide with CCRC2

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

we should be better prepared for CCRC2

Throughput Test

FT=Functional Test

T1-T1 FTS channels

March 03 -07 M6

March 10 - 14 FT T1-T1 Lyon, PIC

March 17 - 20 TT

March 25 deadline SRMv2 @ CERN

March 21 -24 Easter

March 25 - 28 FT T1-T1 BNL, SARA

April 02 deadline for re-proc @ all T1's

April 02 deadline for SRMv2 @ all T2's

April

• **April 31 - 04 TT**

• **April 07 - 11 FT T1-T1 RAL, ND**

• **April 14 - 16 TT**

• **April 21 - 25 FT T1-T1 CNAF, FZ**

• **April 28 - 30 TT**

• **April 01 - 04 Ascension**

May 02 deadline for re-proc FDR d

May 02 deadline for all T1-T1 tested

May 05 - 30 CCRC08-2