



CCRC'08

Common Computing
Readiness Challenge

Jamie.Shiers@cern.ch

Introduction

- CCRC'08 is not at some point in the future
- ***It is now – it has started!***
- Many previously somewhat independent activities have been brought into sync by CCRC'08 planning
- There are still some open issues – but very little time
- This is an update on the current state of planning
- Plus a roadmap for the coming weeks...



Plan for yesterday

Done:

- ✓ Conclude on scaling factors
- ✓ Conclusions on scope / scale of February challenge
 - Not all 2008 resources available at all sites - minimum exercise at Tier1s is loop over all events with calibration DB look-up
- ✓ Monitoring, logging and reporting the (progress of the) challenge
 - Discussion at WLCG Service Reliability workshop with concrete proposal to CCRC'08 planning meeting & GDB this week

To do:

- Conclude on SRM v2.2 storage setup details
- Stop-gap + long-term solution for "storage token" on recall
- 'Walk-throughs' by experiments of 'blocks', emphasising "Critical Services" involved and appropriate scaling factors
- CDR challenge in December - splitting out 'temporary' (challenge) and permanent data
- Other tests that can be done prior to February??
- De-scoping - plan for agreeing this is required



LHCb as an example...

- Generated quite a few questions (in both directions)
- Some issues still not clear
- Fabio, Gonzalo & Luca (and others?) will prepare a list of questions based on this style of presentation & amount of detail
- This will then give a template for all experiments and for all sites
- Very little time to iterate - need feedback from sites as rapidly as possible (before Xmas!) to allow for an iteration in January
- **We also need to use the February challenge to further refine the information that sites require for planning and operations**

Planned tasks

- Raw data distribution from pit → T0 centre
 - Use of rfcpl into CASTOR from pit - T1D0
- Raw data distribution from T0 → T1 centres
 - Use of FTS - T1D0
- Recons of raw data at CERN & T1 centres
 - Production of rDST data - T1D0
 - Use of SRM 2.2
- Stripping of data at CERN & T1 centres
 - Input data: RAW & rDST - T1D0
 - Output data: DST - T1D1
 - Use SRM 2.2
- Distribution of DST data to all other centres
 - Use of FTS - T0D1 (except CERN T1D1)

All tasks envisaged during data taking in 2008

Useful Links

- LHCb planning meeting
 - <http://indico.cern.ch/conferenceDisplay.py?confId=23406>
- LHCb CCRC08 Twiki page
 - <https://twiki.cern.ch/twiki/bin/view/LHCb/CCRC08>

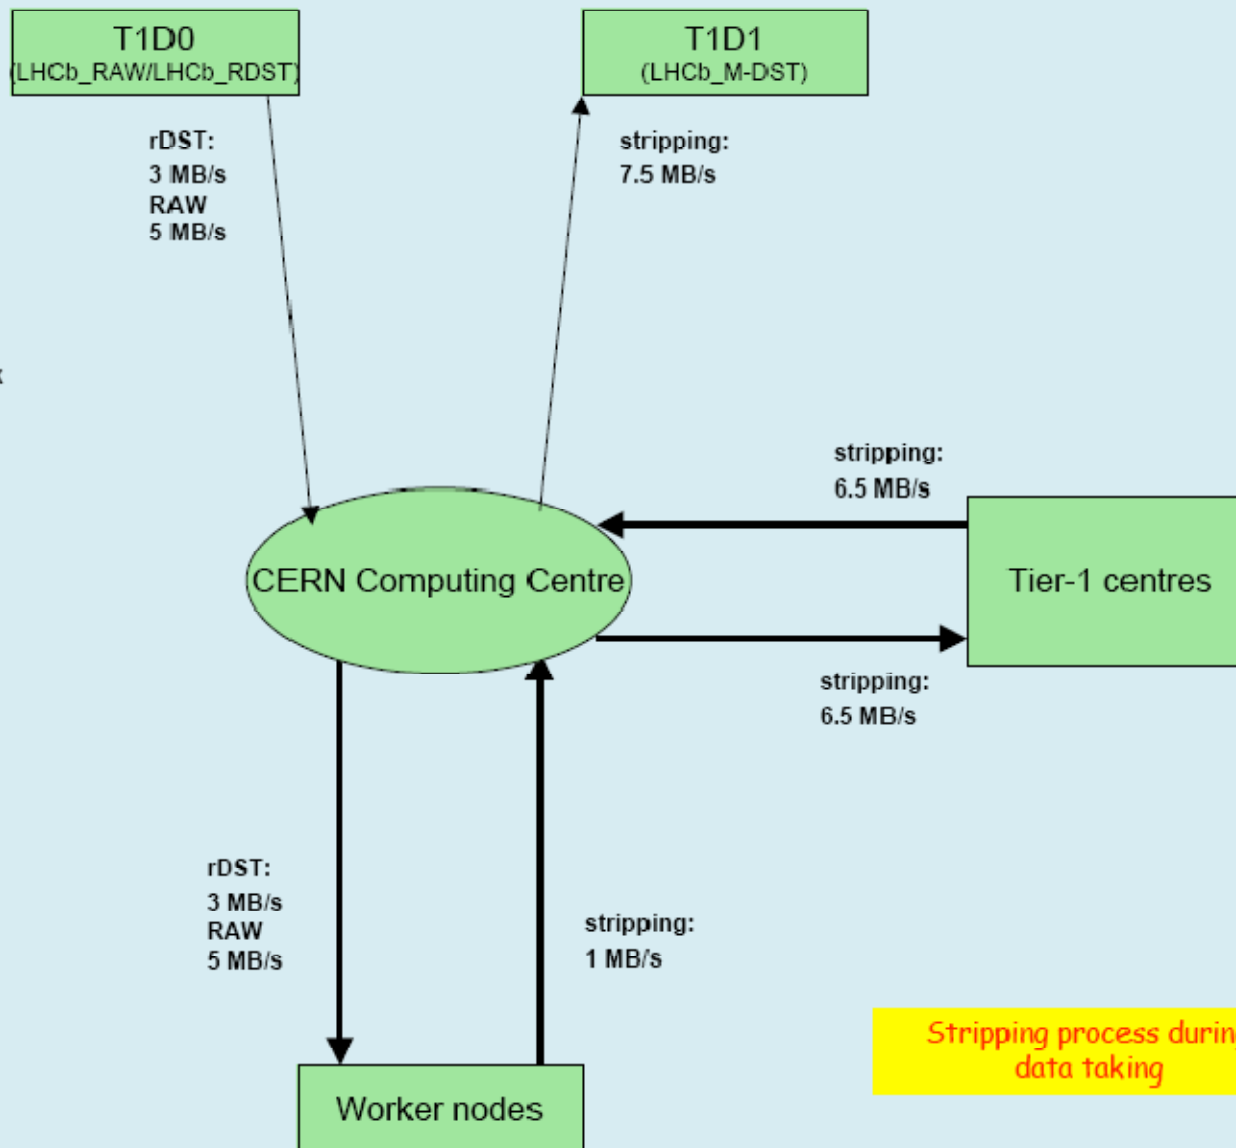
Nos of jobs/site

	<u>Total Jobs</u>			<u>Simultaneous jobs</u>		
	Recons	Strip	Total	Recons	Strip	Total
CERN	3300	400	3700	236	29	265
FZK	1700	200	1900	122	15	137
IN2P3	2700	300	3000	193	22	215
CNAF	1800	200	2000	129	15	144
NIKHEF	5700	700	6400	408	50	458
PIC	900	100	1000	65	8	73
RAL	6900	900	7800	493	65	558
Total	23000	3000	26000	1643	215	1858

Amount of data/site

	T0D1	T1D1	T1D0	Disk	Tape
CERN	0	8	45	8.0	53.0
FZK	7.4	0.6	5.2	8.0	5.8
IN2P3	7.1	0.9	8.1	8.0	9.1
CNAF	7.4	0.6	5.6	8.0	6.2
NIKHEF	6.0	2.0	17.2	8.0	19.2
PIC	7.7	0.3	2.8	8.0	3.2
RAL	5.6	2.4	21.0	8.0	23.4
Total	41.1	14.9	105.0	56.0	119.9

(Note: This should be pinned on disk cache in between rDST production & stripping)



Data Access

- Will need SRM 2.2 SE with correct space tokens
 - LHCb space tokens are:
 - LHCb_RAW (T1D0)
 - LHCb_RDST (T1D0)
 - LHCb_M-DST (T1D1)
 - LHCb_DST (T0D1)
 - LHCb_MC_M-DST (T1D1)
 - LHCb_MC_DST (T0D1)
 - LHCb_FAILOVER (T0D1) <- NEW!!!!
- Need access to lcg utils/GFAL on WN
 - Will be using lcg-gt & lcg-cp as a minimum as part of running applications

Databases

- Conditions DB at CERN & Tier-1 centres
 - For February will use static information replicated using "streaming" from CERN to Tier-1's
 - No plans to test replication of conditions DB Pit ↔ Tier-0 (and beyond) during February
 - Standalone tests at a later date
- LFC
 - For February will use local T1 instance if available & tested (probably RAL, IN2P3 & CNAF)
 - Use "streaming" to populate the read-only instance at T1 from CERN
 - Programme of testing already under discussion with above sites



Proposed blocks of tests

1. Data recording at CERN

- **Scope:** readout from P5, HLT, w. stream definition, incl. Storage Manager, transfer to T0, perform repacking, write to CASTOR
- **Performance goal:** 250Hz
- **Resources required:** CPU, T0 disk, Tape bandwidth, Tape storage

2. Processing at T0:

- **Scope:** from CASTOR, use CMSSW.x.x, write N out-streams to CASTOR
- **Performance goal:** 250Hz
- **Resources required:** CPU, T0 disk, Tape bandwidth, Tape storage

3. CERN data export to T1:

- **Scope:** export to all T1's to MSS
- **Performance goal:** 600MB/s aggregate
- **Resources required:** T0 and disk, network, Tape bandwidth and storage at T0 and T1

4. T1 data handling and processing:

- **Scope:** processing and skimming from tape
- **Performance goal:** full '08 scale (or hw limit)
- **Resources required:** tape bandwidth, disk, CPU



Proposed blocks of tests

5.1 T1 data export:

- **Scope:** Test T1 export to all seven CMS T1
- **Performance goal:** full '08 scale
- **Resources required:** commissioned links, tape bandwidth, disk, network

5.2 T1 data export:

- **Scope:** Test T1 export to T2
- **Performance goal:** full '08 scale, to > 3 T2 at 20 MB/s
- **Resources required:** commissioned links, disk, network

5.3 T1 data import:

- **Scope:** Test T1 import from T2 to tape
- **Performance goal:** full '08 scale, from > 3 T2
- **Resources required:** commissioned links, tape bandwidth, disk, network

T1 import export tests (5.1-3) should be done individually and then together

6 T2 MC production and Analysis:

- **Scope:** Test MC production and Analysis
- **Performance goal:** tbd
- **Resources required:** CPU, disk

Block 1: T0



- Data collection in CASTOR
- Archive to TAPE
- Calibration processing *)
- merging
- Processing → ESD, AOD *), DPD *)
- Subscriptions

*) only when we can use FDR data

Block 1 T0 storage req.s



- FDR data: 20 TB, M5 data: 20 TB
- Calibration pool *)
- Merge pool: 50 TB
- 5 day export buffer: 100 TB
- CAF: 300 TB
- Analysis pool
- *More accurate numbers soon*

Block 2: T1's



- Store RAW (share) on TAPE
- Store ESD (share) on DISK
- Store full AOD & DPD copy on DISK
- Export AOD & DPD to Tier-2's on request
- *Re-processing has to worked out in detail*
- *Production processing has to be worked out in detail*
- *More accurate numbers will follow*

Block 2 T1 Storage req.s



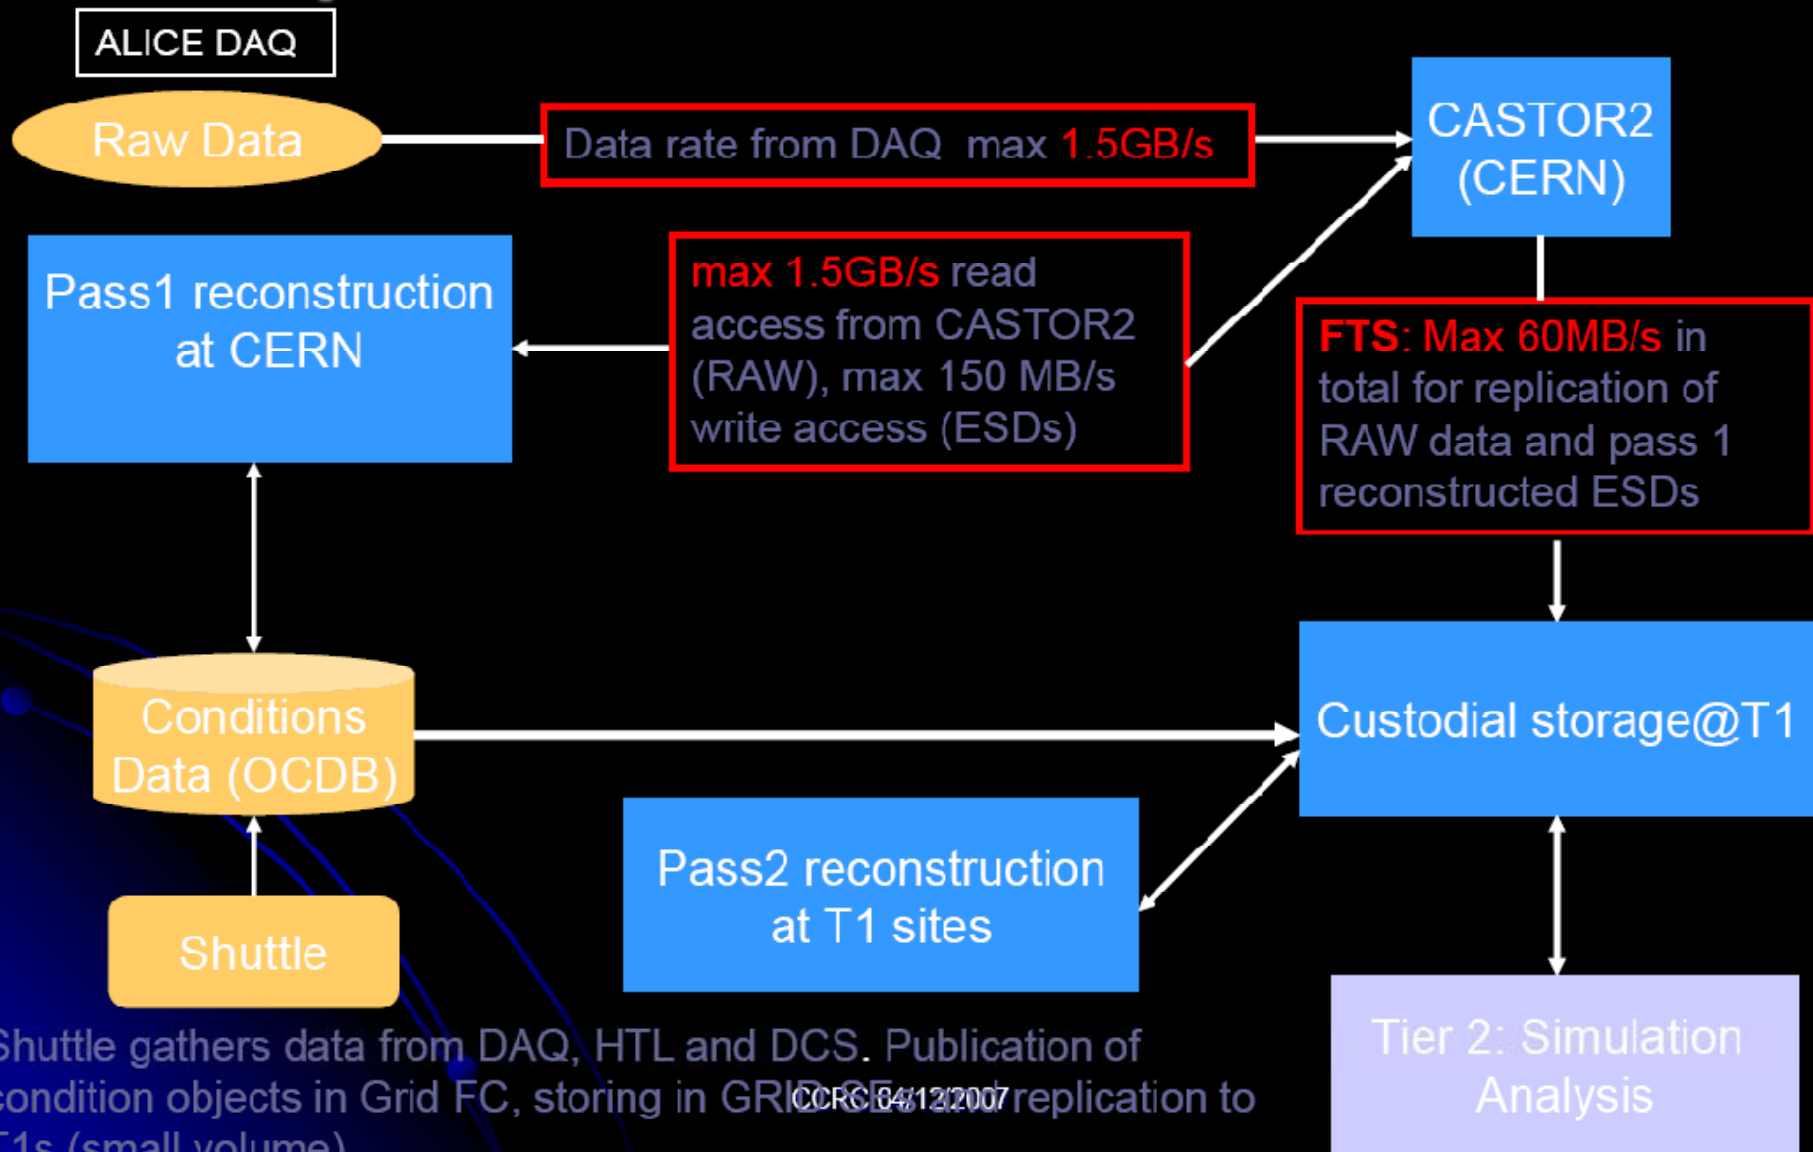
- Want to use real (not test) endpoints
- ~20 TB/day → ~600 TB/month
- Roughly 50% Tape and 50% Disk
- ~300 TB/month for 10 Tier-1's
- So a 10% Tier-1 should have ~30 TB disk
- Data can be removed shortly after Feb.
- Will provide more precise numbers
- *This accounts for primary RAW data only*
- *Re-processing & production will follow*

ATLAS Tier-1 tokens



DATADISK	T0D1	Real data (ESD,AOD,DPD)
DATATAPE	T1D0	Real data (RAW)
DATADISKTAPE	T1D1	Real data (REPROCESSING)
MCDISK	T0D1	Simulation production activities
MCTAPE	T1D0	Simulation production activities
MCDISKTAPE	T1D1	Simulation production activities

Data paths and rates



Shuttle gathers data from DAQ, HTL and DCS. Publication of condition objects in Grid FC, storing in GRID CE and replication to T1s (small volume)

FTS/SRM (Dec 2007)

- SRM Space Tokens and service classes
 - Ready at CERN (default, castoralice) - single pool and service class currently
 - In preparation at GridKA - three cases (related to dCache v.1-8)
 - With space reservation (SRM compliant) used by FTS (r/o for xrootd)
 - In addition a 'recuperable' tape pool need to be created
 - T1D0 and T0D1 without space reservation (r/w for xrootd)
 - Various cross-pool copying mechanism are discussed (if at all needed)
 - May also be managed through specific directories
 - In contact with dCache experts at CCIN2P3
 - Will follow with CNAF, RAL (CASTOR2), SARA, NDGF (dCache) using the experience from the discussion with GridKA, CCIN2P3 and CERN



Situation at the end of this week(?) (See also ATLAS Tier1 Jamboree)

Need also results from next week's CMS week...

Done:

- ✓ Conclude on scaling factors
- ✓ Conclusions on scope / scale of February challenge
- ✓ Monitoring, logging and reporting the (progress of the) challenge
- ✓ Conclude on SRM v2.2 storage setup details
- ✓ Stop-gap + long-term solution for "storage token" on recall
- ✓ 'Walk-throughs' by experiments of 'blocks', emphasising "Critical Services" involved and appropriate scaling factors

- CDR challenge in December - splitting out 'temporary' (challenge) and permanent data
- Other tests that can be done prior to February??

- **De-scoping - plan for agreeing this is required**



Did we succeed?

- Good exposure of plans from, in particular LHCb (pictures!)
 - They had some questions which were answered
 - And the presentation generated quite a few questions / comments from the sites
- **Ideally, have same level of detail (and presentation) from all experiments**
 - The fact that the LHCb presentation generated questions suggests (to me at least) that there would also be questions for the other experiments
 - "unknown unknowns"
- ATLAS Tier1 Jamboree this week and CMS week next week will help to clarify things
 - But there's no time for another F2F this year!
 - We will we have one in January - looking at Thursday 10th in 160-1-009
 - **Probably only 1 phone call < Xmas (and even this might be hard!)**
- Space tokens - still need input from CMS (after next week's CMS week)
- **SRM v2.2 issue: - live with what we have - revisit after concrete experience from February run**
- Monitoring / logging & reporting - deferred to GDB - but very important!

Questions (& Answers?)

- Will a weekly con-call plus a monthly F2F be enough to ensure good information flow in the coming months?
- How do we see 'the big picture' amongst all the detail?
- How do we make sure that the necessary detail flows up and triggers the required actions?
- There is no time-slot that suits everybody – regular calls will start with Asia-Pacific [tomorrow](#) (wider than CCRC'08)
- Associated follow-up: SRM production, reliable services etc.
- ✓ Better use of current meetings: issues for sites will be prepared by Thursday COB so sites can investigate and prepare response; issues that are not resolved can be escalated via WLCG Service Report to GDB; things on a shorter timescale → daily meeting or SCOD

Some Meetings...

- Weekdays at 09:00 – operations meeting – phone available (not used)
 - Mondays at 17:00 – CCRC'08 planning call – 14, 21, 28 January
 - Mondays at 16:00 – joint operations meeting
 - Mondays at 15:30 – SRM v2.2 Production Deployment call
 - Mon / Wed / Fri – EMT
 - Wednesdays at 10:00 – LCG SCM
 - Thursdays at 10:30 – Physics Services Meeting (GLs + PL)
 - Fridays at 09:00 – “C5”
- Not including experiment meetings, technical meetings, workshops, workshops, workshops...
- **Have to make best use of these – rationalize where possible – and ensure there is the right information flow and identification / follow-up / resolution of problems**
 - **Escalation does work – but we can't escalate all issues!**



Next steps...

COMMUNICATION

COMMUNICATION

COMMUNICATION

Summary

- CCRC'08 has already started!
- It is a joint effort – bringing many different aspects of WLCG service & experiments' production together
- Planning is still incomplete – more detail is needed by the sites, plus transparent logging / visualization / communication
- The clear focus now has to be on production
- We will learn many things in February and will need to analyze these rapidly for the full-scale challenge in May
- We need to work closely together and look forward to first data – and first results!
- May you live in interesting times!

- ***May you live in interesting times*** is reputed to be the English translation of an ancient Chinese proverb and curse. However, many people have searched for the original Chinese language version and have not been able to determine its origin. This has led to some doubting its authenticity. It has been attributed to various American speakers, though research is continuing into the phrase's first appearance. It dates back to at least 1950, when the April issue of *Astounding Science Fiction* included the saying in one of the magazine's stories.
- It is reported that it was the first of three curses of increasing severity, the other two being:
 1. ***May you come to the attention of those in authority***
 2. ***May you find what you are looking for***

Questions?



Competition...

- The name "CCRC'08" is somewhat cumbersome...
- Aka "SSCL"
- A local radio station had a competition for a nickname
- Candidates:
 - Ronald Reagan Center for High Energy Fisix
 - CLYDE
- Any (polite) suggestions?