



CSCS

Swiss National Supercomputing Centre

WLCG GridKa+T2s Workshop

Site Report

Peter Kunszt, CSCS, Switzerland

Overview

CSCS resources and plans for acquisition of resources

Grid services offered by CSCS

Participation to date in SC4

Participation in remainder of 2006 and beyond

Issues & Concerns



CSCS Size and Ramp-up Plans

Today

- 30CPU + service nodes + VO boxes
- 8TB

Tomorrow -> end of Year

- 50CPU cores -> 200 CPU cores
- 32TB -> 80TB

Level of resources planned at LHC startup

- CPU cores : 2000+
- Disk : 800+ TB
- Network (external) : 10 Gb/s
- MSS not planned but possibility exists

FTE

- Technical and Operation Services (hardware operations and basic operating system) : 1FTE + backup
- Grid Support : 4 FTE
- Not dedicated to LCG but total size of Grid Team, working also for other projects



Grid services offered by CSCS

Running on glite-3.0.2

Grid Services

- BDII
- CE
- DPM

Application Services

- LFC
- VO-box CMS: Phedex, Frontier (squid)
- VO-box ATLAS: NorduGrid CE

Hardware and OS

- SUN
- Linux (scientific o.k. if compatible with RedHat AS)
- maybe Solaris



Participation in SC4 : CMS

CMS

- Installation of CMS services fairly easy:
 - Phedex
 - Central CMSSW installation, with some manual fixing for rfiio
- CSCS is among the sites with over 90% stability:
http://lxgate30.cern.ch/rrdtool/statistics_eff.php?Time=month
shows an efficiency of ~93%
- FTS transfers are very unstable, basically unusable
- srmcp is also anything but optimal
- we need a tool to control the exports from/to our site
- we are still limited by our own HW, not enough disk and CPU



Participation in SC4: ATLAS, LHCb

ATLAS

- VOBox does not run DQ yet, but rather the ARC middleware
- Requested an LFC, working properly at CSCS
- SC4: only FTS used, which is set up at FZK for CSCS
- has seen the same unstability of FTS as CMS

LHCb

- Constant levels of production, no complaints
- Tier2s are not part of SC4 exercise
- CSCS may be used also for analysis in the future



SC4 File Transfers for CMS

PhEDEx via srmcp

- ~90% success with FNAL, IN2P3, CERN
- only ~40% with FZK
- Maximal throughput 30MB/s peak
- CPU was most of the time waiting for IO
- 10MB/s stable should be doable.

FTS via the STAR-CSCS channel hosted by FZK

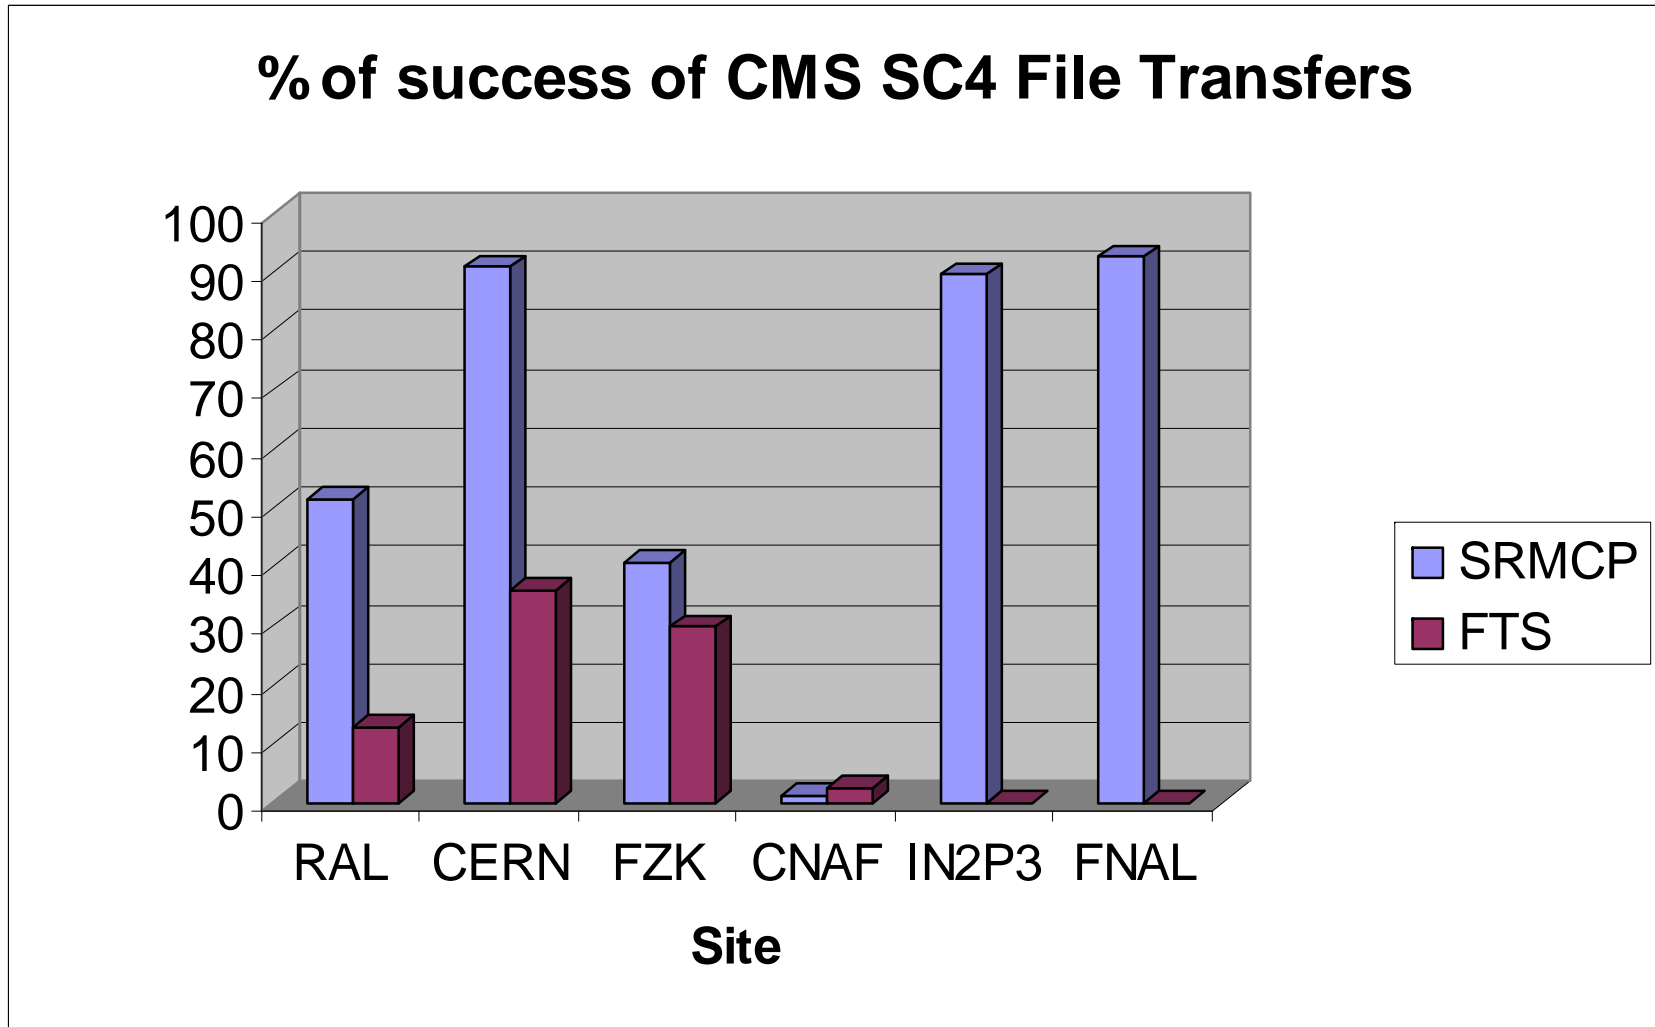
- very unstable
- poor success rates: see next slide

For details see

<https://twiki.cscs.ch/bin/view/LCGTier2/FTSChannelDebugging>



Comparison of FTS and srmcp at CSCS using Phedex in August



Difference due to direct (srmcp) vs 3rd party copy (FTS) ??



Interpretation of the Comparison

- Today, the percentages are much better, also with FTS, but they still should be much improved
- Too many things can go wrong (the stack is very large) and therefore actually do go wrong
- A lot of optimization and fine-tuning is necessary, the relationships between the optimizations is unclear
- SC4 fulfills its purpose: the abilities are thoroughly tested
- Recommendations
 - Focus on robustness and responsiveness of SRMs – too many connection timeouts observed
 - Make the stack thinner if possible: too many layers of abstraction – Gridftp, SRM, FTS, Phedex
 - A better error reporting is much needed



Participation in 2006 and beyond

We plan to participate in the CMS, ATLAS and LHCb challenges

The operations will continue for LHC for the foreseeable future.



Issues & Concerns

Unstable Central services?

- Most common reason for critical SFT failure on the data management test is an intermittent failure of connecting to CERN:
BDII ERROR: lcg-bdii.cern.ch:2170 Success

Poor middleware update management by CERN?

- Most of our issues were due to bad patches and updates that were wrong, and had to be fixed manually
- SFTs were not updated and site failed for no reason
- CRL and CA updates were sometimes manual and ill-prepared and caused a large number of problems
- Is the Preproduction Service used to test the component before rolling it out or are we the testers?

Meeting the 95% availability?

- We are at 93% as measured by CMS, but this is the best possible value we can get: the rest of the failures are mostly due to the factors above which we don't control directly



Issues & Concerns, cont.

Data management is too fragile

- Failures can happen at too many levels in the many layers
- Unclear error messages, hidden causes for error, hard to debug
- Better communication on best practices would be good

Currently CSCS is too small

- The funding is secured for the new hardware for this year
- There are some administrative problems to overcome

EGEE Support issues

- CSCS participates in the DECH ROC shifts, CIC and soon TPM.
- Synchronization between the DECH ROC tickets and the GGUS tickets?
- Escalation procedure could be much improved, not intuitive
- Not always clear responsibilities



Suggestions for Improvement

Information on the kind of jobs being run at the site could be better

- Granularity of info is only on VO level
- Are the VO jobs production jobs? analysis? tests?
- This would help to fix problems as we have seen very strange behavior sometimes, i.e. jobs sitting on the WN waiting on the network for hours, wasting the resource

Information for Tier2s could be better concentrated

- Known bugs, Best practices – current wiki lists only a subset of possible topics
- Hardware procurements, sites running similar hardware, similar configurations could exchange knowledge
- DPM and dcache configurations, FTS/srmcp transfer tuning, maximizing the network throughput

