# Transfer service review

***Gavin McCance***

***CCRC'08 F2F meeting, CERN***

***4 March 2008***

# Summary

- **Not had much time to do a major analysis**
  - Which is a lesson in itself…
  - I'm restricting to the T0-export
- **We still see critical problems (only) in production**
  - The "corrupted proxy" issue
- **Various issues uncovered:**
  - A few Globus-related issues
  - The way we use the information system
  - SRM copy
  - Channel management
- **Operations and problem resolution is still expensive**
  - Monitoring! Streamlined procedures!

# Status

- **Fairly good?**

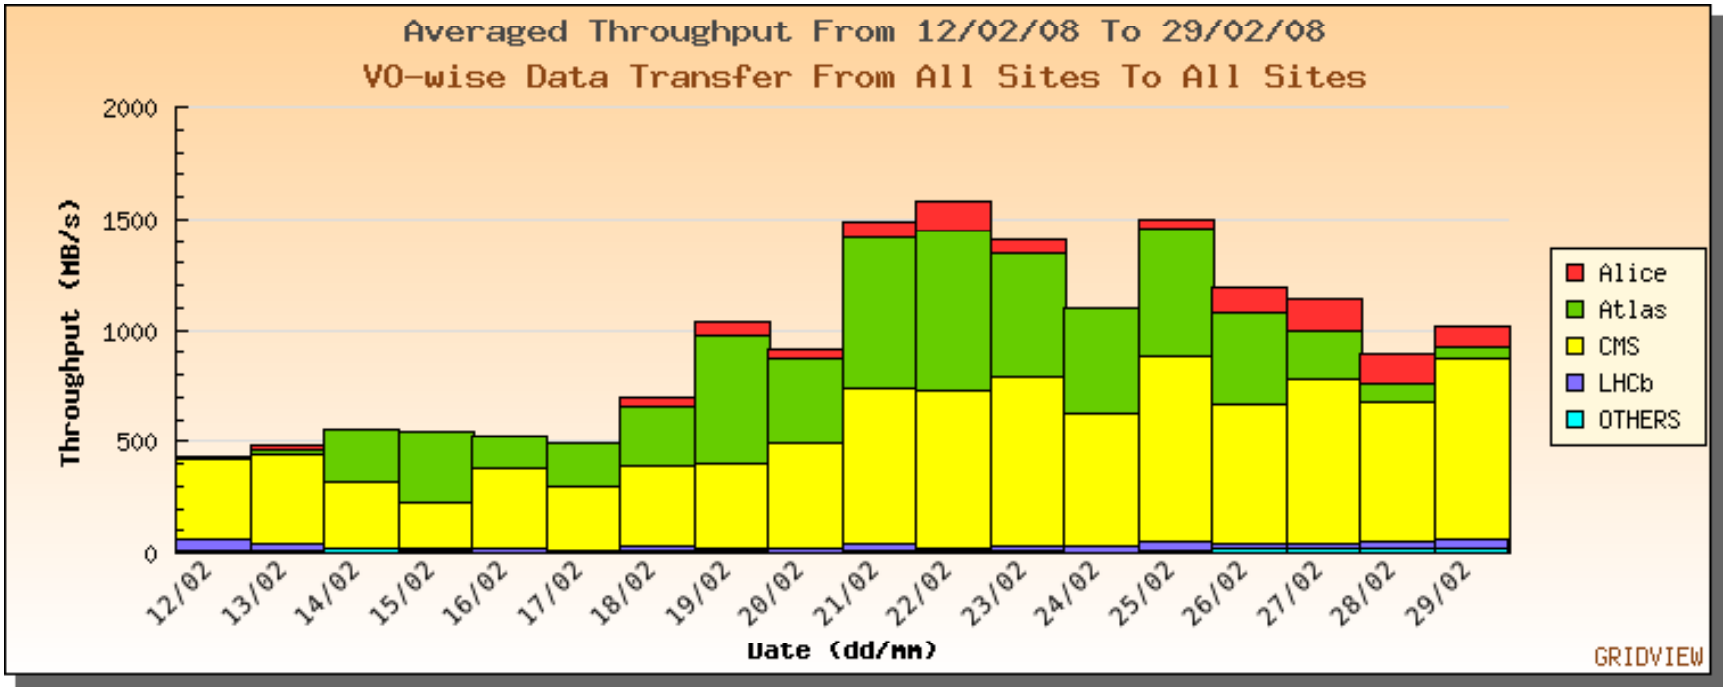

# "Corrupted" proxies

- **Intermittent downtime around Valentine's day**
  - https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortemFeb14
  - Some Castor related issues – resolved quickly
  - FTS "corrupted proxy" issue caused intermittent periods of downtime over a few days [Alice, Atlas, CMS]
- **Race condition in the delegation protocol**
  - Sometimes you get a bad proxy for 8 hours
- **Was the FTS delegation code well enough tested?**
  - The answer is clearly "no"
    - (3 months testing, 3 months experiment beta, 6 month production)
  - Success! Hey… we found it now, not in May ☺
  - Some things only get seem tested in production-like exercises

# "Corrupted" proxies

- **Workarounds in place once we understood the problem**
  - Suggested workaround on client side
    - Not very nice to have to implement this on short timescales
  - Dirty server workaround to limit damage caused by bug
  - Code fix underway: https://savannah.cern.ch/bugs/?33641
- **How did the (reporting, fixing) process go?**
  - From our perspective, not badly
    - we got initial reports quite fast (GGUS, mailing lists) and development team were able to get a 'live bad proxy' as soon as one surfaced again
  - Understanding took a while
    - A few days until a we got a valid workaround
  - Operational workaround publication
    - Daily meeting / Twiki with follow-ups. Not sure everyone saw it?
- **Q. How can we improve the process? (I'm sure we can…!)**

# Globus VDT 1.* related

- **Couple of issues found:**
- **[Alice] Delegation library that the FTS uses doesn't support version 3 proxies properly**
  - Not critical while we remain on VDT 1.2 / 1.6
  - But the clock is ticking…
- **[CMS, FZK] Globus / VDT1.6 makes more stringent checks on reverse DNS lookup**
  - Reverse DNS entry needs to be correctly configured
  - Requires care when using load-balancing
  - Action: issue relevant advice on this

# BDII and services.xml

- **[CMS] FTS relies on the information system where sites appear and disappear from**
  - BDII doesn't (yet) split static and dynamic information. The assumption was that if a site is down, its OK for the information not to be there. This is ~OK for the WMS, but not OK for the FTS or Phedex, where we care about the existence of a site, as distinct from the state of the site
- **So we cache the information in a sticky cache**
  - Created from the BDII using *glitesd2cache* regularly
  - We recommend sites install this for now
- **Action: publicise and document this better!**
  - We could have prepared the sites better for this
- **Action: ongoing BDII service improvements [Markus]**

# SRM copy

- **Why?**
  - Performance. dCache's gridFTP door architecture makes SRM copy a better option.
- **Issues:**
  - Recently: Some known issues regarding SRM copy and v1 / v2 interaction. [Flavia, this afternoon]
  - Operational: FTS is a client of SRM-copy. If it goes wrong, we can't help much! (i.e. you get the ticket back)
- **Changes on horizon:**
  - NDGF have a VDT patch to allow us to use the gridFTP2 client within VDT. This will allow us to talk the gridFTP2 protocol to dCache 1.8, which by-passes the door problem. This should remove the URLCOPY performance "bottleneck".
  - Status: NDGF talking to VDT and Globus for patch inclusion

# Channel management

- **[ATLAS] Would like to change the way we do channel management**
  - Using approach prototyped by RAL with some success
- **A new feature of FTS 2.0 allows you to manage each VO share separately**
  - (Old way: #streams, and per-VO shares of this)
  - The SRM resources are ~dedicated anyway
  - Better control and stability
  - Less risk of interference between experiments
  - Less risk of overloading disk-pools with the ‘elastic share’
- **We would like to implement this on CERN-PROD soon**

# Operations

- **Investigating problems remains expensive**
  - I didn't get much else done!
  - It's still log-file intensive [FTS, Castor, dCache, gridFTP]
  - We're still tending to be re-active rather than pro-active

- **Monitoring monitoring monitoring!**
  - The strategy and plans are fairly clear
    - [ GDB monitoring session tomorrow ]
    - We know what data we have in the FTS
    - We just need to get it done: extract the information from the services, get it out there into our daily operations flow

# Ticket flow

- **Issues are typically reported to GGUS .and. eLogger, .and. mailing list (Remedy feed or otherwise)**
  - Most get picked up first from lists or eLogger
  - GGUS usually comes late
  - eLogger RSS feed is good
  - Short daily meeting gives good direction
    - top issue of the day slot per-VO is useful
- **GGUS is probably still the best placeholder we have for 'tracking the issue'**
- **Some synchronization overhead**
- **Some noise – out-of-band follow-up of already ticketed issues**
  - this is costly and should be avoided

# Summary of lessons learned

- **We could have had better preparation**
  - Installation / dissemination of various tools
  - Better testing (always…)
- **Reaction to issues is ~reasonable**
  - Processes worked quite well
  - Some overhead in the procedures (room for improvement)
- **Operations remains fairly costly**
  - Monitoring!
- **Even with these issues, we achieved a lot**
  - And (importantly) we learned quite a lot too…