FTS channel configuration jamboree

Gavin McCance

CCRC'08 F2F meeting, CERN

1 April 2008

Summary

- **Derek just presented what RAL have done**
- **The proposal was to do a similar thing elsewhere**
  - Initially on CERN-PROD
  - Then other T1s
- **Some discussion and comments on various mailing lists**
  - Mostly from CMS and LHCb

What we do now

- **The channel share is set as (e.g.):**
  - ATLAS 30, CMS 30, LHCB 10, ALICE 10
  - 40 files concurrently running for all VOs
- **Units are relative only:**
  - When everyone is running:
    - ALICE gets 10 / (30+30+10+10) * 40 files = 5 files
    - ATLAS get 30 / (30+30+10+10) * 40 files = 15 files
- **Shares are magically ‘elastic’:**
  - When only LHCb and ALICE are running:
    - ALICE gets 10 / (10 + 10) * 40 files = 20 files

The issues

- **Overload: ALICE's disk pool might only be sized for 5 files (say, max 10).**
  - If we give them 20, it'll kill the performance of the disk pool(s), or just plain kill the disk pool(s)

- **Stability and inter-dependence: the rate you get depends on the other VOs, so your rate is very variable**
  - The model doesn't well represent the actual reality of separate disk pools (and sometimes separate SRM instances).
  - **It increases the 'mystery' factor**
    - "it was going fast, now it going real slow"
  - This is a pain for the operations

Advantages

- **Main one quoted:**

- **The elastic share is good**
  - I get to use someone else's share if they're not using it

"The RAL approach"

- **Just outlined by Derek**

- **Essentially you use the new VO limit setting on the channel to set separately how many concurrent files each VO will run at any one time**
  - and you set the aggregate channel limit high enough so it never bites

Advantages

- File bandwidth guaranteed per channel/VO combination
  - PhEDEx limitation
- File bandwidth limited per channel/VO combination
  - Lopsided resources
- Independence
  - VO files can be adjusted without affecting file bandwidth of other VOs.

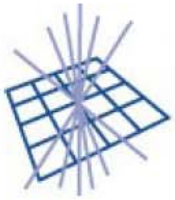

GridPP
UK Computing for Particle Physics

Science & Technology Facilities Council
e-Science

Objections to this

- **“I can’t use someone else’s share now”**
- **Sure.. Why don’t you just set your per-VO limit higher?**
  - (up to what your disk pool will reasonably manage)
    - Noticed going slow recently on CERN-CNAF with LHCb – we should just bump up the limit to something higher!
- **Risk? DOSsing a shared SRM server (as distinct from overloading the disk pool of any one VO)**
  - When all VOs run with a higher limit
  - It’s not clear this problem even exists (??)
    - If the bottleneck does exist you probably want to know about it!
  - Site negotiates with VOs – reduces all rates as agreed until you’re under the bottleneck limit
    - (Or… set the aggregate channel limit to the maximum the SRM can sustain, and that will bite in the limit of overloading the SRM)

Discussion

- **Can we try this?**
  - It should help a lot with the understanding of the service
- **I think the control it gives us far outweighs the downside of loss of the ‘elastic shares’**