

FTS channel configuration discussion

FTS admin workshop, SARA
18 October 2006

Gavin McCance



Scope and overview

- Scope is the WLCG service
 - "Full FTS services"
- Basic channel concept review
- Current WLCG recommendation
- Discussion...
 - Model changes if necessary
 - Understand discovery implications



Management concept: channel

- For management ease, the service supports splitting jobs into multiple “channels”
 - Once a job is submitted to the FTS it is assigned to a suitable channel for serving
- A channel may be:
 - A point to point network link
 - (e.g. we manage all the TO-export links in WLCG on a separate channel)
 - Various “catch-all” channels
 - (e.g. everything else coming to me, or everything to one of my tier-2 sites)
 - More flexible channel definitions are on the way (but not there yet)
- Channels are uni-directional
 - e.g. at CERN we have one set for the export and one set for the import



Management concept: channel

- "Channel": it's not a great name
 - It isn't tied to a physical network path
 - It's just a management concept
 - "Queue" might be better 😊
- All file transfer jobs on the same channel are served as part of the same queue
 - Intra-VO priorities for the queue (Atlas gets 75%, CMS gets the rest)
 - Internal-VO priorities within a VO
- Each channel has its own set of transfer parameters
 - Number of concurrent files running, number streams, TCP buffer, etc
- Given the transfers your FTS server is required to support (as defined by experiment computing models and WLCG), channels allow you to split up the management of these as you see fit
 - For WLCG, we do however make recommendations...

Channels



- We try to split up the two questions...
- **Policy:** What transfers are a site required to run?
 - i.e. from where to where
 - Should there be any redundancy in the setup for WLCG?
- **Technical:** What channels should a site deploy in order to service that responsibility?
 - Given that some of them can be "catch-all" STAR channels
- Related (for the client) is then the technical question of discovery
 - Given this network of sites all with differing (and possibly overlapping) transfer responsibilities, who do I send my job to?



Policy: Current WLCG policy

- Current policy is fairly simple: 3 rules
 1. **Tier-0:** If it involves CERN, then CERN's FTS does the transfer
 - This covers tier-0 to tier-1 and tier-1 to tier-0.
 2. **Tier-1 sites:** if you are the destination, your FTS is responsible for running the transfer
 - (if you're the source, the other end is responsible)
 3. **Tier-2 sites:** if you are the destination, your associated tier-1's FTS is responsible for running the transfer
 - Regardless of who is the source

- We prioritise control of writing over control of reading

<https://uimon.cern.ch/twiki/pub/LCG/FtsServerInstall15/SC4FTSsetupplan.doc>

<https://twiki.cern.ch/twiki/bin/view/LCG/FtsServerDeployExampleTier1>



Policy: Mapping to use-cases

- Primary use-cases from experiment computing models
 - tier-0 export to tier-1: CERN
 - tier-1 import into tier-0: CERN
 - tier-1 to tier-1 data transfer: the destination tier-1
 - tier-2 upload of data to associated tier-1: the associated tier-1
 - tier-1 push of data to associated tier-2: the associated tier-1
- Secondary(?) use-cases
 - non-associated tier-2 to tier-1: the destination tier-1
 - non-associated tier-1 to tier-2: the tier-1 associated to the destination tier-2
 - tier-2 to tier-2: the tier-1 associated to the destination tier-2



Policy: Motivation

- Keep the model as simple as possible
 - There is one and only one 'correct' place to send any job
- We don't want to run FTS servers at tier-2



Technical: what channels?

- CERN sets up 1 export and 1 import channel per T1
- T1 sets up one channel *pulling* data from each other T1
- T1 has an option for its associated T2s
 - Catch all import channel
 - Specific tier-2 import channels
 - Catch-all export channel
 - Specific tier-2 export channels
- T1 should also setup
 - STAR to each T2 channel (e.g. anywhere to my tier-2)
 - STAR to me channel (e.g. non-associated tier-2 to me)

Discussion



- Is the current model good enough?
- More redundancy?
- Better channel definitions?
 - "Clouds"



"Clouds"

- Possible extension to current FTS channel definition
 - This is at the "technical" level of this discussion
- Current channel definitions don't map so well to the way we model the expt. computing
 - Fixed point-to-point channels are very specific
 - STAR channels catch way too much on the "*"
- Propose "cloud" model
 - Specific sub-groupings? Define:
 - "me to all my tier-2s"
 - "everything in Italy to me"
 - Groupings based on other metadata? Define:
 - "All small US CMS sites to me"
 - The more complex it is the harder it is to verify that you cover all transfers that the policy requires you to cover
- Is this useful?



Questions: redundancy

- We could provide more redundancy
 - The current model has none
 - Allow more than one site to service a transfer (policy)
- "My tier-1 is down, so all ex-cloud transfers to its associated tier-2s have stopped"
 - Do your models care?
 - Are T1s really going to be down for that long?
- Complicating the discovery model
 - It would make the model more complex
 - Debugging gets harder... "where did I put my job?"



Tier-2: Where will the transfer requests come from?

- We prioritise control of writing over control of reading in the current WLCG model
- All **inbound transfers** are controlled by your associated tier-1 site
 - You can control this on the channel setup for you
- The majority of your **outbound transfer** will be upload to your associated tier-1 site
 - You can control this - if your T1 set up an explicit channel for you 😊
- You should expect some **outbound traffic** transferring to other tier-1 sites or to other tier-2 sites
 - These will be controlled by the other tier-1 sites
 - We offer no easy way to shut this off because of the way the channels are defined