

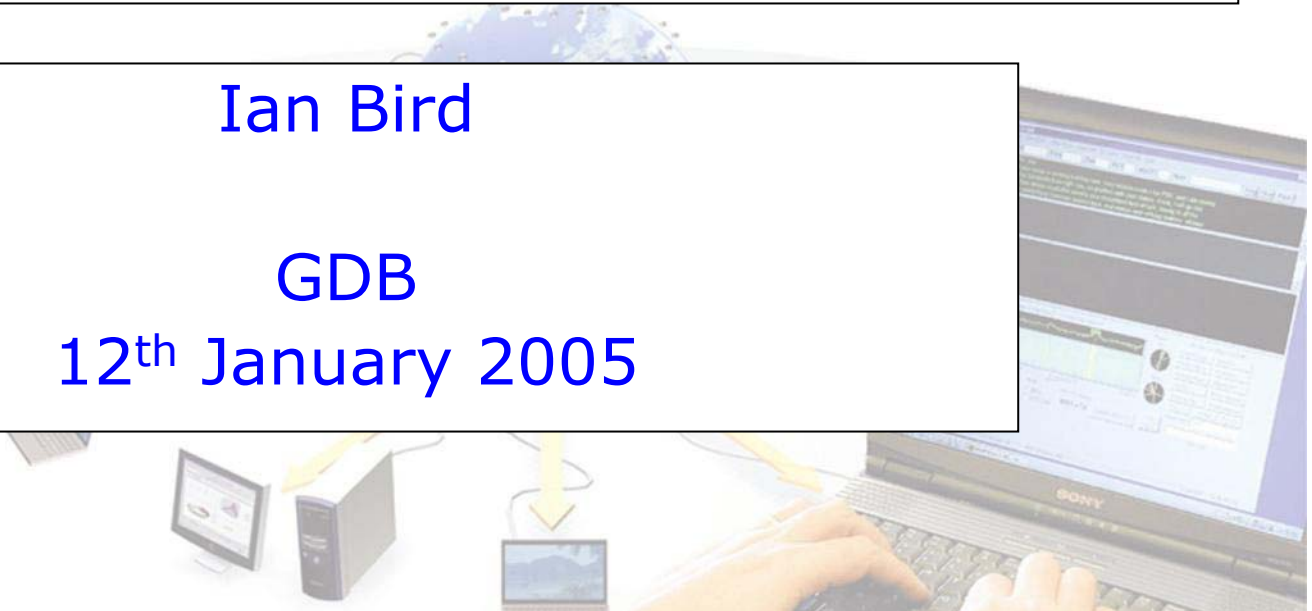


Data Transfer Service Challenge Infrastructure

Ian Bird

GDB

12th January 2005





Introduction

- Clarify what is expected to be in place to support the service challenges
 - Where are the responsibilities
- Three components:
 - Mass storage systems and their grid interfaces
 - Data transfer infrastructure at a site
 - File transfer service



Mass storage systems

- Interface to storage systems is SRM
 - Tier 1: Local integrations of MSS with SRM (e.g. CASTOR, dCache, etc)
 - Large Tier 2: Either as Tier 1, or large disk pools managed by dCache (or other SRM-enabled cache manager)
- Entirely the responsibility of each site to arrange for this integration



Infrastructure at a site

- What is needed to support the data transfer rates anticipated
 - Scalable transfer service infrastructure – disk servers, gridftp servers etc
- Again – clearly the responsibility of each Tier 1 to ensure this is in place in a timely fashion
- Examples:
 - CERN infrastructure
 - Ideas for SAN-based infrastructure – being prototyped with Taipei




Tier-0 Configuration

- 10 dual CPU Itanium machines
 - Each has 1Gb connectivity that gets aggregated into a 10Gb switch
 - All nodes run CERN SLC3
 - ~200GB local fast disk – 50MB/s single stream per disk
- Configured to run the following services:
 - 5 gridftp servers, non-load balanced
 - A 4-node load-balanced SRM/Gridftp system
 - 1 control node, for running the transfer management software
- Direct connections to external network
 - 10 Gb connection to GEANT
 - 10 Gb link to Chicago (via Starlight)
 - 10 Gb test link to SARA/NIKHEF (via Surfnet)



Prototypical SAN at a Tier 1





Site infrastructure – 2

- **Must foresee 2 sets of data transfer services**
 - Service challenge → next step up in performance
 - Ongoing service → as each step is demonstrated in the service challenge it should be moved into production
 - The production service continually improves in performance
- **Typically at present**
 - Expect ~10 transfer nodes to achieve 500 MB/s
 - Multiple requests, multiple files, multiple streams, ...

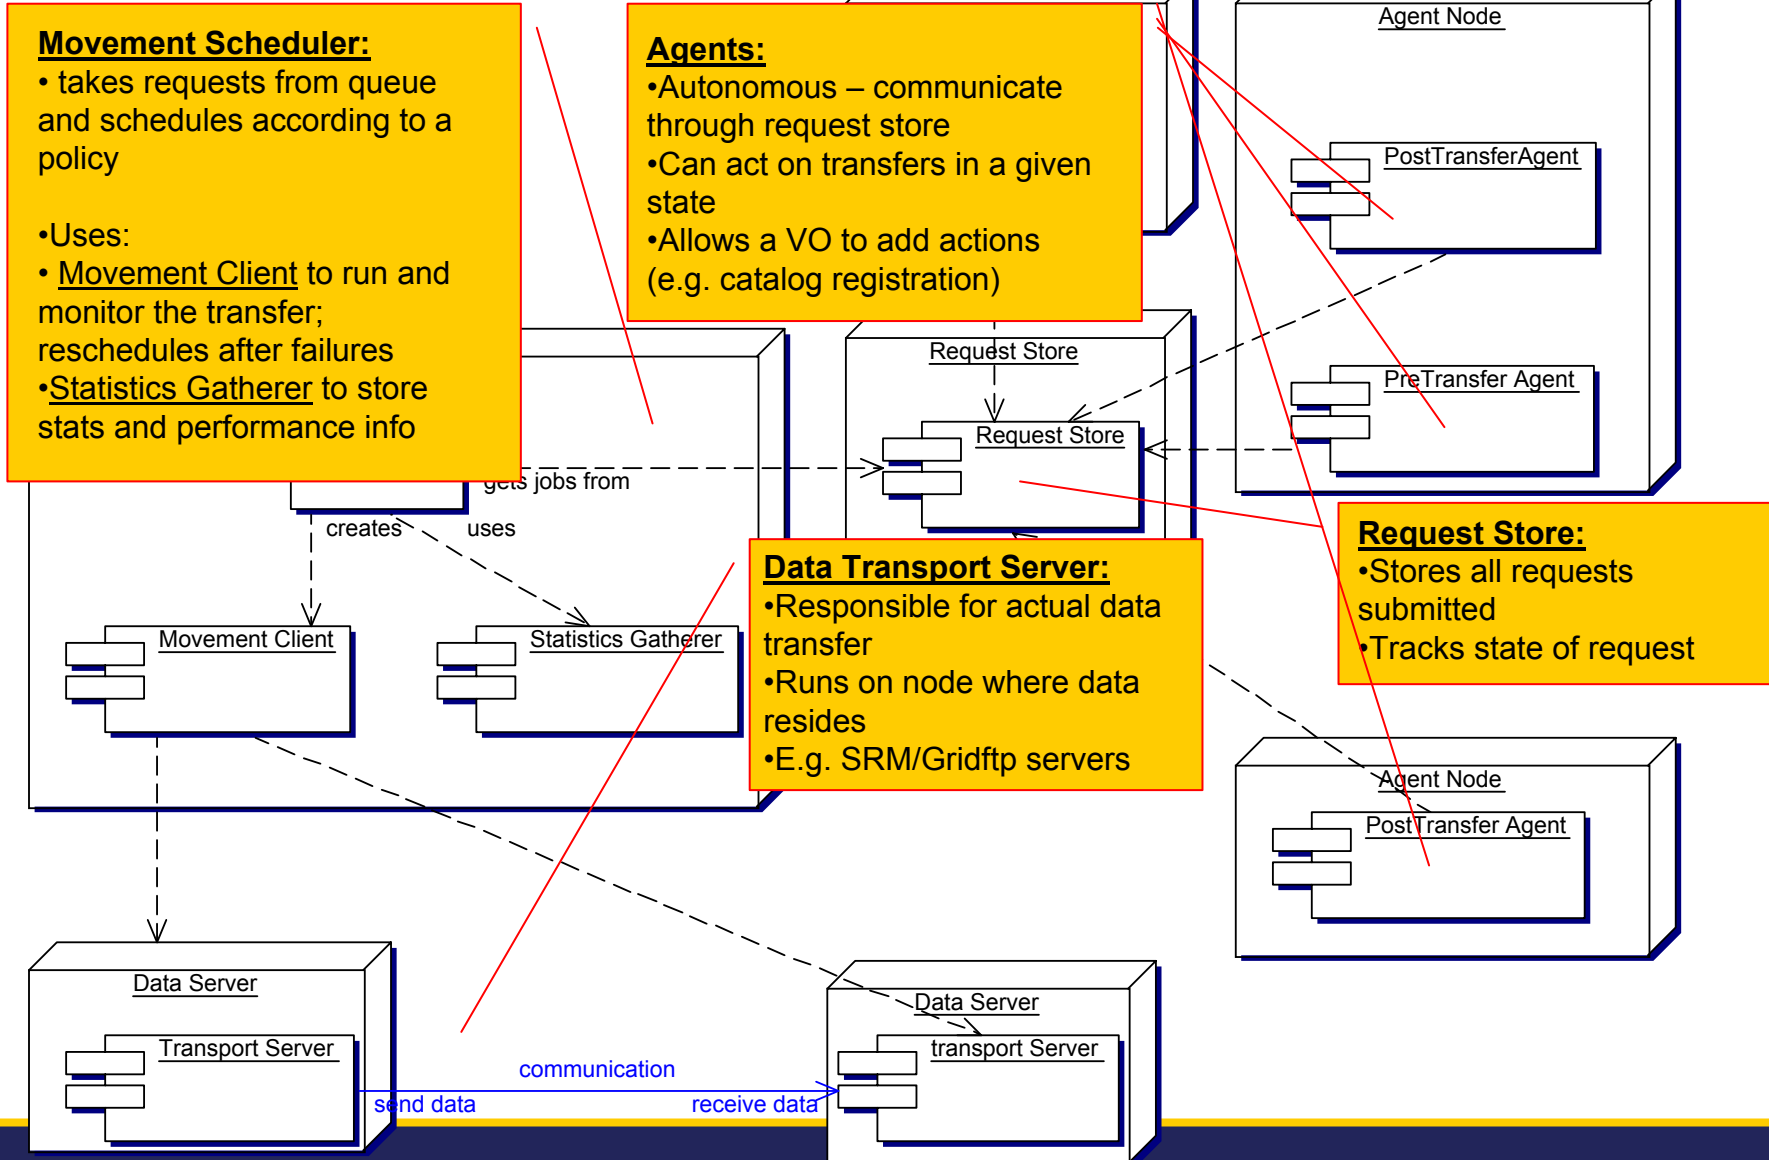


File transfer service software

- Architecture designed with gLite Data Management team
 - Proposal produced July 2004
 - Uses a loosely coupled agent-based model to allow for interaction of different components
 - Current Solution
 - Transfers carried out by simple perl-based multi-process daemon
 - Simply wraps **globus-url-copy**
 - UI/ Web Services submission interface from gLite
 - User interface already what gLite users will see
 - Integrated gLite FTS by Spring 2005
 - As the software becomes available, we will try it
 - We need it for the March 2005 Service Challenge
 - Good testbed for gLite – currently separate from rest of LCG-2 infrastructure
- ⇒ **This will be a joint development between GDA and gLite teams**



Transfer software (radiant





Project Management

- Personnel
 - Jamie Shiers joins Grid Deployment team at CERN with responsibility for Service Challenges
- Communication
 - Website created:
 - <http://service-radiant.web.cern.ch/service-radiant/>
 - Mailing list for CERN service team created
 - service-radiant@cern.ch
 - Single point of contact for any issues related to CERN side of service challenge
 - Hep-forum-wan-tier1@cern.ch
 - List used for dissemination of general information and discussion with Tier-1s about current results and technical issues
 - Does this need to split into two – technical and “info” ?



Current Site Progress (Dec 04)

	RAL	Fermilab	Brookhaven	Karlsruhe	IN2P3	CNAF	PIC
1. Network Configuration	Contact	√	√	√	√	Contact	Contact
2. Storage Configuration		√	√	√	√		
3. Site Tuning			
4. "Service Challenge"		Dec 04		Jan 05			

	Taipei	Nordugrid	TRIUMF	NL	DESY
1. Network Configuration	Contact		Contact	√	Contact
2. Storage Configuration				√	
3. Site Tuning				...	
4. "Service Challenge"				Dec 04	

KEY
 √ = Complete
 ... = In Progress




Summary

- Tier 1 sites must plan and set up this infrastructure rapidly now
 - Network and Storage setup each take ~ 2 weeks to complete
 - Tuning can take much longer
 - Difficult to work out which layer is causing the problem
 - Systems don't scale linear with number of nodes
- Service Challenges require scheduling of many resources
 - Personnel
 - Network allocation
 - Storage allocation



Ongoing Data Transfers



* Graphs produced
01/12/04

