



GridPP

UK Computing for Particle Physics

AliBaBa: Running BaBar jobs on the grid using gsub

Mike Jones

The University of Manchester



THE UNIVERSITY
of MANCHESTER



- aims
- BaBar computing in the UK
- gsub interface
- technology behind gsub interface
- web interface
- technology behind web interface
- job and fabric monitoring





- to get BaBar Scientists to run code on grids today
 - bridge gap between Globus and BaBar
- must be simple to use
 - use what the scientist knows
- must be simple to install
 - client commands are all in AFS
- must be lightweight
 - all Software is either staged
 - ...or available via AFS





UK BaBar computing

- BaBarUK compute farms distributed throughout GB
 - Bham, Bristol, Edin, IC , Liv, Man, RAL, RHul, QMUL
- large datasets located only at specific farms
 - All data at RAL, replica of data at SLAC
 - (overlapping) subsets at other sites
- executables with client
 - Local BaBar Releases
 - Personal Code
- results wanted back on Scientist's PC





Alibaba what-not

- what it is:
 - command line job submission to a Grid
 - web based monitoring of those jobs and the grid
 - stop-gap
 - re-use of existing tested software
 - a simple “input/output sandbox”
- what it is not:
 - web-based/portal job submission
 - load balancer
 - central resource broker





requirements

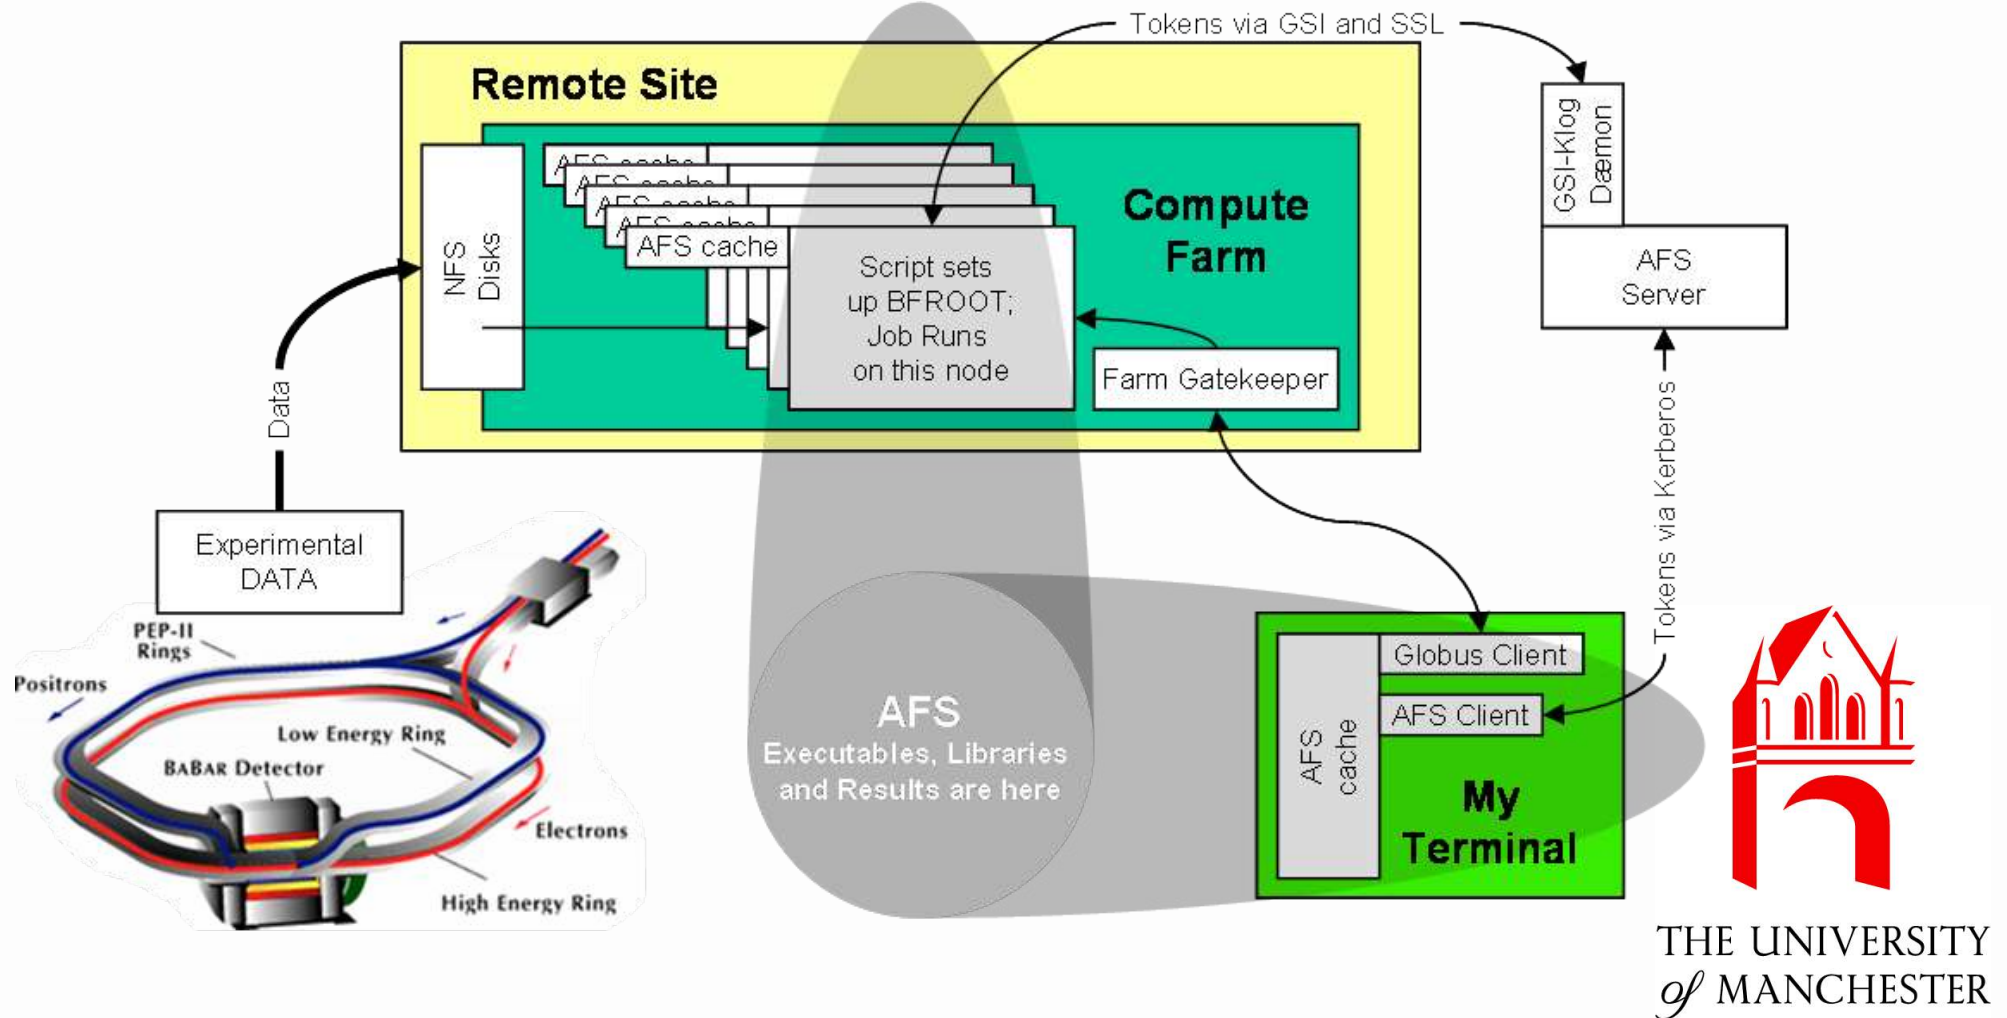
- client
 - bash, X509 cert, globusrun, grid-proxy-init, browser
 - optionally myproxy and/or voms-proxy-init
- BaBar farm
 - bash, Globus gatekeeper and functional job-manager
 - IP address (front and back end), AFS client software
- AFS server
 - gsklogd / gssklogd
- AliBaBar server (≥ 1 per grid)
 - perl, Gridsite
 - optionally globus





Cartoon Architecture

- AFS spans, Globus spawns & gsub holds it together





- what gsub does

- sanity checks
- gets the current list of gatekeepers, on-line
- creates a script (to wrap the executable on farm PC)
 - sets up a normal environment
 - uploads state
 - gets (pag separated) AFS credentials using gsi klog
 - creates BROOT - BaBar environment
 - changes to directory submitted from
 - starts a shepherd process
 - this will look after job's grid stuff and upload status changes
 - runs user's executable (script or binary)
 - unlogs
- Globus: execute the script, on a machine, on the grid
- uses curl to upload the status of the submission





GridPP

UK Computing for Particle Physics

usage

gsub [options] command args

AFS related:

[{-a|-afs} <user@cell>]

[{-a|+afs} <extra user@cell>]+

[{-c|-cell} <cell>]

[{-p|-principal} <principal>]

If not specified by one method above, gsub will try to guess principal and realm.

Globus related:

[{-g|-gate} <gatekeeper>]

[{-j|-jobman} <jobmanager>]

[{-x|-proxy} <non-standard proxy location>]

local machines related:

[{-bf|-bfroot} <local BFROOT>]

[{-d|-display} <DISPLAY>]

remote machine related:

[{-S|-site} <BABAR-SITE>]

[{-s|-source} <RemoteSourceFile1> [{-s|-source} <File2>] ...]

[{-rb|-rbfroot} <Path to Remote BFROOT on Remote Machine>]

[-nb]

[-t|-tmp]

[{-CA|-capath} <path to CA's>]

[{-queue|-q} <queuename>]

user interaction related:

[-i|-int [-e|-err <errorfile>] [-o|-out <outfile>]]

[-l] [-v|-verbose] [-vv|-vverbose] [-D|-dump] [-T|-dry] [-C|-cat]

[-h|-?|-help] [-u|-usage] [-V|-version]

etc.



THE UNIVERSITY
of MANCHESTER



Web based monitoring

- is a CGI perl script
- is hosted by a Gridsite
- takes several variables in get method
 - Default returns a web page with overall status map
 - Links to specific sites' statuses
 - Methods for running jobs to upload their statuses securely
 - Methods for using the server to retrieve globus status and output
- records job statuses
- draws pretty pictures





Alibaba Front page

- site queue status
- jobs submitted
- jobs running
- jobs finished
- links

Alibaba: UK BaBar Farm Grid at a glance - Mozilla

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop Search Print

Home Bookmarks Red Hat Network Support Shop Products Training

Alibaba: UK BaBar Farm Grid at a glance SVE Staff - Mike Jones Frk: The world's first Grid Enabled ...

Alibaba: UK BaBar Farm Grid at a glance

The map below shows the status of the Grid of BaBar Farms in the UK.
Version: 0.98

Site Status on Wed Mar 31 18:32:19 BST 2004

Summary for last 7 days
Average queue time
Weighted by time

- < 1 min
- < 5 mins
- < 20 mins
- < 1 hour
- < 12 hours
- > 12 hours
- Status unknown

Number of jobs.

- Submitted
- Confirmed
- Started
- Running
- Finished

Morgiana.pl, © 2004 University of Manchester

Select a site

Query the status of

Further Information

[Click here for further information about alibaba.pl, gsub and morgiana.pl](#)

Done





Querying job status

Contact URL	Stages					Exit Status	Action
	Submitted	Confirmed	Started	Running	Finished		
https://bfb.hep.man.ac.uk:50103/14701/1080753881/	Wed Mar 31 17:24:29 2004	+0:00:15	+0:00:53	+0:00:55	+0:01:27	0	
https://bfb.hep.man.ac.uk:50100/14637/1080753863/	Wed Mar 31 17:24:14 2004	+0:00:14	+0:00:31	+0:00:35	+0:01:05	0	
https://bfb.hep.man.ac.uk:50095/14453/1080753840/	Wed Mar 31 17:23:51 2004	+0:00:22	+0:00:52	+0:00:55	+0:01:25	0	
https://bfb.hep.man.ac.uk:50074/14325/1080753826/	Wed Mar 31 17:23:37 2004	+0:00:12	+0:01:05	+0:01:09	+0:01:39	0	

Job Details

- Basic details for all

- Times, URL

- Advanced Details for Owner of job

- Times, URL, Locations, Proxy life

Contact URL	Stages					Exit Status	Action
	Submitted	Confirmed	Started	Running	Finished		
https://bfb.hep.man.ac.uk:50103/14701/1080753881/	Wed Mar 31 17:24:29 2004	+0:00:15	+0:00:53	+0:00:55	+0:01:27	0	Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs Started as masj@194.36.3.78 The proxy certificate 673 minutes
https://bfb.hep.man.ac.uk:50100/14637/1080753863/	Wed Mar 31 17:24:14 2004	+0:00:14	+0:00:31	+0:00:35	+0:01:05	0	Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs Started as masj@194.36.3.18 The proxy certificate 673 minutes
https://bfb.hep.man.ac.uk:50095/14453/1080753840/	Wed Mar 31 17:23:51 2004	+0:00:22	+0:00:52	+0:00:55	+0:01:25	0	Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs Started as masj@194.36.3.16 The proxy certificate 673 minutes
https://bfb.hep.man.ac.uk:50074/14325/1080753826/	Wed Mar 31 17:23:37 2004	+0:00:12	+0:01:05	+0:01:09	+0:01:39	0	Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs Started as masj@194.36.3.78 The proxy certificate 673 minutes
https://bfb.hep.man.ac.uk:50073/13908/1080753763/	Wed Mar 31 17:22:31 2004	+0:00:31	+0:01:11	+0:01:13	+0:01:43	0	Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs Started as masj@194.36.3.78 The proxy certificate 673 minutes
https://bfb.hep.man.ac.uk:50100/14117/1080753794/	Wed Mar 31 17:23:04 2004	+0:00:32					Submitted by masj@130.88.1.130 Submitted to x-gram://bfb.hep.man.ac.uk:2119/jobmanager-pbs The proxy certificate 673 minutes





Alibaba

Behind the scenes


- **gsub at the command line**
 - https with local proxy send XML files to Gridsite CGI
 - job ready to submit
 - job successfully submitted
- **gsub script on the remote machine**
 - https with delegated limited proxy
 - send gsub script started
 - send real job started
 - send job finished/terminated





Status Map

- image updated on server every time state changes
- site blob colour
 - time jobs spend in queue
 - weighted by age of result
- extremely easy to add a new sites
 - add directory on server
 - create xml file with xy position of site!





- is AFS slow? - not really
- BaBar jobs run (if they're proxy is good)
- what does AFS do?
 - cache and transfers files
 - list, read, write, create, delete, lock, dir admin
- time consuming components
 - token creation
 - actual transfer
 - obtaining locks
 - cache limits
- script to test AFS speed using gsub
 - read ~ 250-500KB/s small files ~ 2-10MB/s large
 - write ~ 50-100 KB/s small files ~ 1-3MB/s large
 - append ~ 1-3 KB/s small files ~ 1-3MB/s large





Token expiry

- DESY, SLAC, RAL ... have AFS home directories
 - Current jobs fall over when AFS token expires
- A job submitted with gsub has its own shepherd
 - shepherd can catch a new proxy
 - Globus' refreshproxy
 - shepherd can ask for a new proxy
 - myproxy
 - shepherd can obtain new AFS token
 - gsiklog





- Janusz Martyniak - Ganga Interface

File View Job Actions Help

Jobs

- New
 - globus_generic
 - GLOBUS_babar
 - BaBar_testEDG
 - alibabar_097a
- Configured
- Submitted
- Running
- Completed

Job

Batch System: GLOBUS State: New

GLOBUS Job Status: Gatekeeper: bfb.hep.man.ac.uk:2119

Comments:

AliBaBar Application

AFS username: janusz AFS cell: hep.man.ac.uk

AFS user workdir: /afs/hep.man.ac.uk/u/janusz/analysis/beta14a/workdir

Executable File: BetaApp

Input Tcl File: /afs/hep.man.ac.uk/u/janusz/analysis/beta14a/BetaUser/kangaFilterMicro.tcl

Use Index

Index File/Dirname: /afs/hep.man.ac.uk/u/janusz/analysis/tclfile_short.idx

Parameters

Variable: BetaKanga yes Add

```
1 Welcome to Python Shell (PyCrust 0.9.2)
2 Python 2.2.2 (#1, Jan 30 2003, 21:26:22)
3 [GCC 2.96 20000731 (Red Hat Linux 7.3 2.96-112)] on linux2
4 Type "help", "copyright", "credits" or "license" for more information.
5 >>>
```



- based on the UK eScience GITS
 - which are based on Teragrid's original tests
 - bash (or ksh) cf perl - for job control reasons
- index-centric
- contains extra test for gsub
- writes results in text, html and xml
- XML files are compatible with UK eScience GITS database

- Is wrapped in a script: bftests
- uses gatekeepers.xml rather than GITS
- writes results directly to webserver





- Grid Fabric
 - test results
 - descriptions

Grid Resource Interoperability Results - Mozilla

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop http://bfhome.hep.man.ac.uk/BFtests.html Search Print

Home Bookmarks Red Hat Network Support Shop Products Training

Grid Resource Interoperability Results

Grid Resource Interoperability

Run on 2004-01-02 at 16:43:12 by /C=UK/O=eScience/OU=Manchester/L=MC/CN=michael.jones/CN=proxy from pc53.hep.man.ac.uk

Test Hosts obtained from http://bfhome.hep.man.ac.uk/gatekeepers.xml

PING	GPING	RSL	GJR	GJRS	GJSUB	GJST	GJCL	GJRET	GASS	GFTP	GSSH	GSUB
Timeouts (in seconds) set to:												
5	10	60	60	60	60	60	120	90	90	40	30	600
gppce04.gridpp.rl.ac.uk:2119/jobmanager-pbs:/C=UK/O=eScience/OU=CLRC/L=RAL/CN=gppce04.gridpp.rl.ac.uk/Email=s.traylen@rl.ac.uk												
PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	FAIL	FAIL	FAIL	PASS
bottom.phy.bris.ac.uk:2119/jobmanager-pbs:												
PASS	PASS	PASS	PASS	PASS	PASS	PASS	FAIL	PASS	FAIL	PASS	FAIL	PASS
bfh.hep.man.ac.uk:2119/jobmanager-pbs:												
PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	FAIL	PASS

Keys

- PASS** Passed!
- FAIL** Failed
- N.T.** Not Tested
- unAuthz** Remote user not Authorised
- Unknown** Remote Authentication failed
- Bad Gate** Gatekeeper's certificate from http://bfhome.hep.man.ac.uk/gatekeepers.xml is different to certificate supplied by gatekeeper
- No 3rd party** GSI Authentication to ftp server ok but failed port or passive command.
- T.O.** Time Out

PING # ping <host>

GPING # globusrun -a -r <gatekeeper>

RSL # globusrun -o -r <gatekeeper> '&(executable="/bin/echo")(arguments="Hello World")'

GJR # globus-job-run <gatekeeper> /bin/echo Hello World

GJRS # globus-job-run <gatekeeper> -s \$TEMPFILE

GJSUB # globus-job-submit <gatekeeper> /bin/sleep 600

GJST # globus-job-status <contact>

GJCL # globus-job-clean -f <contact>

GJRET # globus-job-submit <gatekeeper> /bin/echo HelloWorld
globus-job-status <contact>
globus-job-get-output <contact>

GASS # globusrun -s -r <gatekeeper> "&(executable=\\$(GLOBUS_LOCATION)/bin/globus-url-copy)(arguments=\\$(GLOBUSRUN_GASS_URL)\\${TEMPFILE} \ \ "file:\\${TEMPFILE}.g1")(environment=(LD_LIBRARY_PATH \\$(GLOBUS_LOCATION)/lib))"
globusrun -s -r <gatekeeper> "&(executable=\\$(GLOBUS_LOCATION)/bin/globus-url-copy)(arguments=\ "file:\\${TEMPFILE}.g1" \ \ \\$(GLOBUSRUN_GASS_URL)\\${TEMPFILE}.g2)(environment=(LD_LIBRARY_PATH \\$(GLOBUS_LOCATION)/lib))"
diff -q \\${TEMPFILE}.g1 \\${TEMPFILE}.g2

GFTP # globus-url-copy file:\\${TEMPFILE} gsiftp://<host>\\${TEMPFILE}.1
globus-url-copy gsiftp://<host>\\${TEMPFILE}.1 file:\\${TEMPFILE}.2
diff -q \\${TEMPFILE}.1 \\${TEMPFILE}.2
globus-url-copy gsiftp://<host>\\${TEMPFILE}.1 gsiftp://<host>\\${TEMPFILE}.2

GSSH # (gsi)ssh -o "BatchMode=yes" -p \$PORT <host> /bin/echo HelloWorld

GSUB # gsub -site <BABAR-SITE> "/bin/echo Hello World > afstempfile" ; echo afstempfile | grep "Hello World;"

Done



Conclusions

- Easy to use
 - job submission similar to non-grid used today
 - job status and tracking in browser
- Easy to install
- Uses basic globus (can work along-side LCG...)
- Uses AFS fabric already required for non grid
- Uses whichever VO Management is deployed





- `gsiklog` -> `gssklog`
- expand `gssklogd` take-up
- more automated data discovery
 - skimData (Index for raw data) web service interface
- Better resource discovery
 - static list of resources -> dynamic list
- Data movement
 - AFS stuff fine for small (<1GB) transactions
 - want to run at any grid enabled farm
 - Data must be present or moved

