


BeSTGRID DataFabric

BLUEFERN



BeSTGRID DataFabric
New Zealand's online storage service based on
iRODS

Vladimir Menci

BlueFern Supercomputing Services

University of Canterbury

Talk Outline: BeSTGRID DataFabric

BLUEFERN

- BeSTGRID DataFabric – service overview
 - Software stack – user point of view
 - Sites
- Authentication and authorization – user experience
 - Shibboleth, X509 certificates, internal accounts
- Software stack – under the hood
 - webDavis – web + webDAV front-end
 - iRODS
 - Griffin – gridFTP interface into iRODS

BeSTGRID DataFabric

BLUEFERN

- BeSTGRID DataFabric – what it is for users
 - a storage system “in the cloud” (but hosted by friendly universities they know)
 - local to the country (legal restrictions, bandwidth, bandwidth costs)
 - with a web and webDAV interface for ease of use
 - with iDrop interface using native iRODS protocol for better performance
 - integrated with computational grid
 - job directories get archived onto the DataFabric
- Key features
 - auto-provisioned accounts for NZ researchers (see authNZ bellow)
 - multiple entry points (Auckland, Canterbury, more coming)
 - multiple storage resources (associated with entry points...)
 - behind the scenes data replication

BeSTGRID DataFabric - screenshot

BeSTGRID Data Fabric - Start Page - BeSTGRID - Mozilla Firefox

df.bestgrid.org/BeSTGRID/home#/BeSTGRID/home/vladimir.mencil

You are logged in as <vladimir.mencil>

You are browsing: /BeSTGRID > home > vladimir.mencil

17 items listed, no items selected

[Upload File](#) [Create Directory](#)

Name	Last Modified	Size	QuickShare
... Parent Directory			
archived-jobs	Thu 15 Dec 2011 03:11:26 GMT		
BlueFern	Wed 10 Nov 2010 05:57:49 GMT		
comp-results	Wed 20 Jul 2011 02:03:06 GMT		
evo-recordings	Wed 19 Sep 2012 03:10:24 GMT		
grid	Fri 22 Oct 2010 02:08:56 GMT		
grid-sw-archive	Thu 16 Aug 2012 03:46:11 GMT		
link-test	Thu 01 Apr 2010 12:43:06 GMT		
public-test	Thu 08 Apr 2010 03:33:54 GMT		
r	Thu 24 May 2012 02:57:59 GMT		
tardir	Fri 09 Apr 2010 00:03:44 GMT		
test	Tue 30 Mar 2010 04:42:30 GMT		
test-from-grisu	Tue 01 Mar 2011 22:40:09 GMT		
ChChQuakeMap.png	Thu 13 Oct 2011 22:49:43 GMT	1.1M	https://df.bestgrid.org/quickshare
DOEGrids-2006-Nov-8.tar	Fri 19 Aug 2011 00:30:38 GMT	81K	https://df.auckland.ac.nz/quickshare
DucklingsBeware.jpg	Fri 19 Aug 2011 00:30:39 GMT	84K	https://df.bestgrid.org/quickshare
job-mpi-mrbayesdir-df.xml	Fri 19 Aug 2011 00:30:41 GMT	1.3K	https://df.bestgrid.org/quickshare
pbs.pm	Fri 19 Aug 2011 00:30:42 GMT	14K	https://df.bestgrid.org/quickshare

[Support](#)

[Select all](#)

- Access Control
- Metadata
- Replicas
- Rename
- Copy
- Move
- Delete
- Share/Unshare

BeSTGRID DataFabric – Sites

BLUEFERN

- University of Canterbury
 - <http://df.bestgrid.org/>
 - Primary server on x86_64 Linux
 - webDavis
 - Griffin
 - iRODS master + iCAT + PostgreSQL
 - but no resource
 - Slave data server on IBM p520
 - GPFS access to 37TB storage resource
- University of Auckland
 - secondary site: <http://df.auckland.ac.nz/>
 - Secondary webDavis, Griffin
 - iRODS slave server with 15TB

BeSTGRID DataFabric – User authNZ

BLUEFERN

- Web interface: two authentication methods
 - Shibboleth (via Tuakiri NZ federation)
 - Nice, ties with federated SSO
 - But does not work with webDAV
 - no redirects / forms supported by clients
 - iRODS Username + password (needed for webDAV)
 - Deprecated: MyProxy username+password
 - Produces an X509 certificate for GSI
- Accounts auto-provisioned on first access
- User accounts tracked by SharedToken
 - Hash value passed as a Shibboleth attribute
 - And in DN in GSI X509 certificates
 - Stored in user info in iCAT
 - **Different access methods with same sharedToken get mapped to same account.**

BeSTGRID

DataFabric – under the hood - webDavis

BLUEFERN

- WebDavis: runs as a Jetty web application
 - Apache with ShibbolethSP as front-end (AJP connector)
 - Accepts multiple authentication mechanisms
 - Requests on plain http use Shibboleth login
 - Requests on https use username + password
 - Talks to iRODS server through Jargon
 - iRODS server typically collocated on the same server
 - But Jargon redirects the communication directly to the target resource
 - Handles both web and webDAV requests
 - And supports files >2GB where supported by clients
 - not all browsers and webDAV clients are equal...

Login handling in webDavis

BLUEFERN

- iRODS username+password: direct login to iRODS
- MyProxy username+password: do a login to MyProxy, use the X509 certificate returned for a GSI login to iRODS
 - New DN triggers account creation via acGetUserByDN rule
- Shibboleth login:
 - Authenticate as rods (via GSI with host certificate) to look up the user by sharedToken (iCAT query)
 - If not found, create user by iexec cmd
 - Authenticates as user (newly implemented) – use a rods session with clientUserName
 - Modification to Jargon (Classic) was needed.
 - Cannot use GSI – must do a rods password-based login (clientUserName with GSI not supported server-side)
 - Older hack: iexec cmd reset password on internal iRODS account for the user
 - Prevented us from using internal iRODS password

BeSTGRID DataFabric – under the hood - Griffin

BLUEFERN

- Griffin
 - GridFTP interface for iRODS implemented in Java
 - Again, uses Jargon to talk to iRODS
 - Uses GSI to authenticate as the user
 - Runs as a standalone Java application
 - GridFTP endpoint registered in our computational grid infosystem (MDS replacement)
 - GridFTP endpoint also reachable via GlobusOnline
- => Seamless integration of iRODS into the NZ computational grid

References

BLUEFERN

- Software packages
 - webDavis – web + webDAV front-end
<http://code.google.com/p/webdavis/>
 - Griffin (GridFTP front-end for iRODS)
<http://code.google.com/p/datafabric-griffin/>
- Live system – main website
 - <http://df.bestgrid.org/>
 - Anonymous access at <https://df.bestgrid.org/BeSTGRID/home/NZDF>
- User documentation
 - http://technical.bestgrid.org/index.php/Using_the_DataFabric
 - http://technical.bestgrid.org/index.php/Administering_the_DataFabric
 - http://technical.bestgrid.org/index.php/Installing_an_iRODS_slave_server