



The Data Bridge

Cloud storage access for Grid and volunteer computing

Fabrizio Furano (speaker)
Laurence Field
IT/SDC

Pre-GDB 9 March 2015



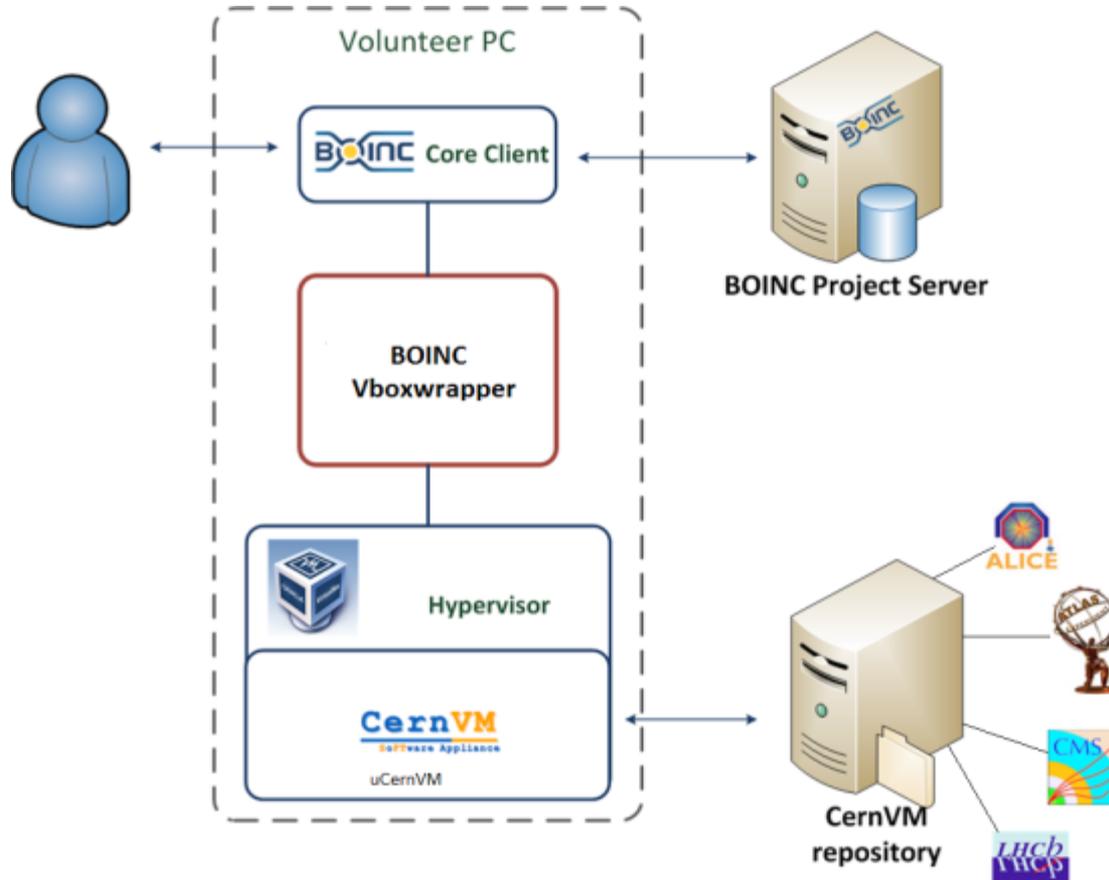
Intro

- The Data Bridge is a component that bridges authentication domains for storage access
- Context: volunteer computing (BOINC) and Grid environment
- It's built on Apache plus the Dynamic Federation technology with S3 buckets as backends
 - S3 was key to simplicity, provides an effective delegation mechanism
 - Dynafeds exploit that, plus giving scalability, uniform authorization and flexibility
 - Apache can host any number of other 'standard' authentication plugins, it will work
- It's a generic idea to harmonize Cloud storage and multiple authentication domains, including Grid/X509/VOMS

BOINC workflow: main points

- A volunteer PC gets a job agent running on a VM
- Agent gets work to do from a shared queue
- Agent writes its output to a known place
- FTS moves data (once validated) to some prod storage

BOINC and Virtualization

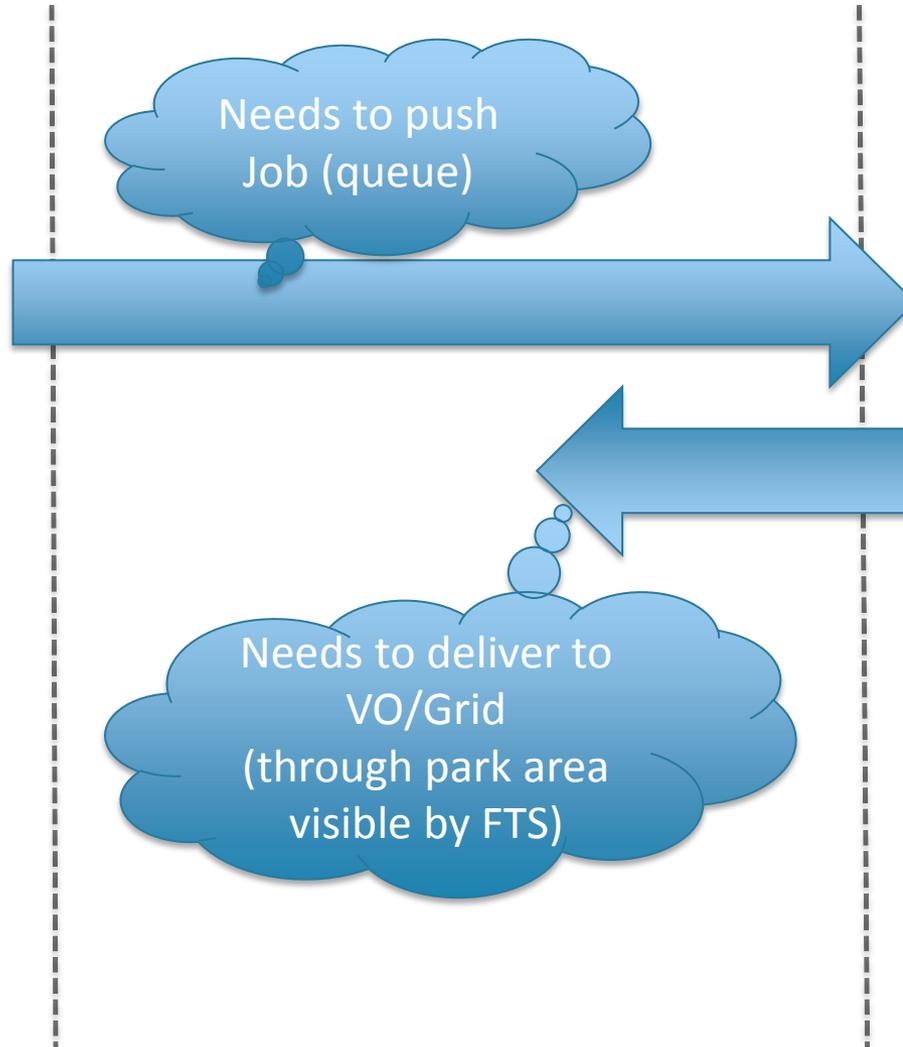
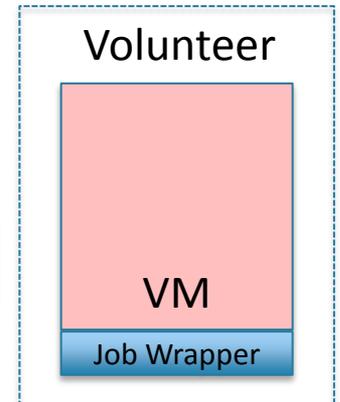


The Challenge

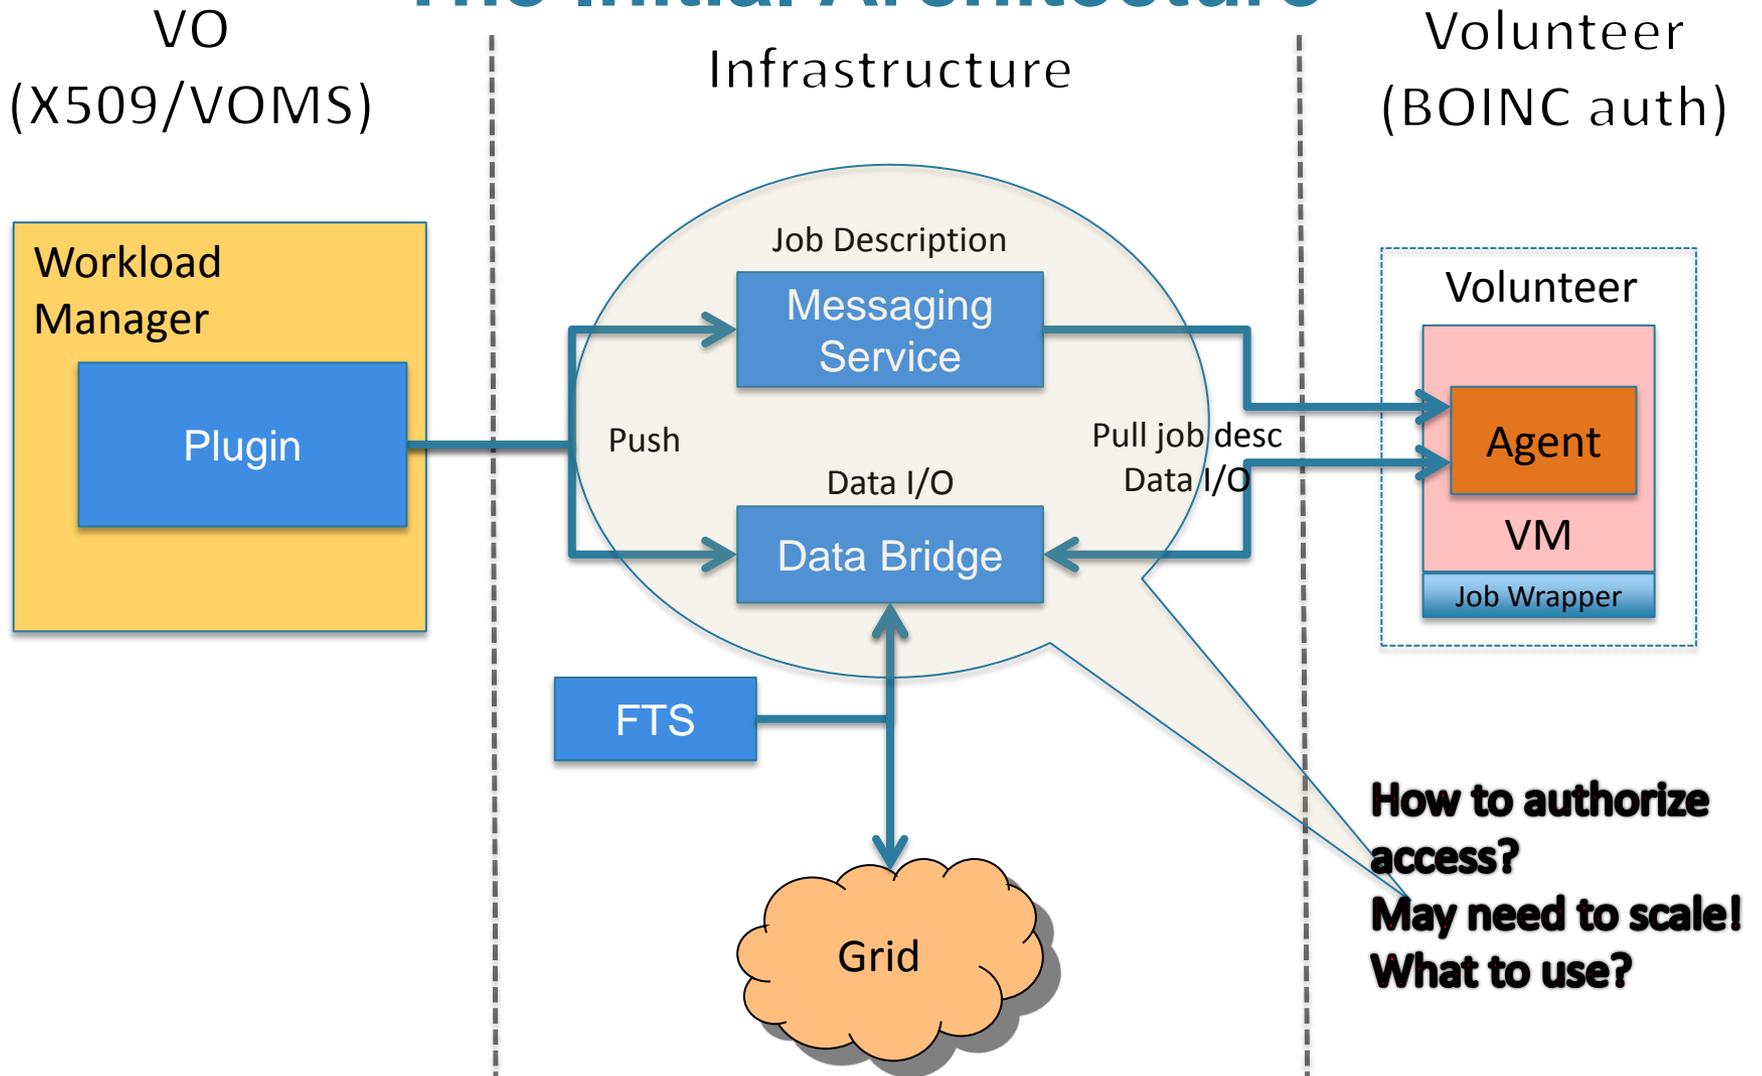
VO
(X509/VOMS)



Volunteer
(BOINC auth)



The Initial Architecture



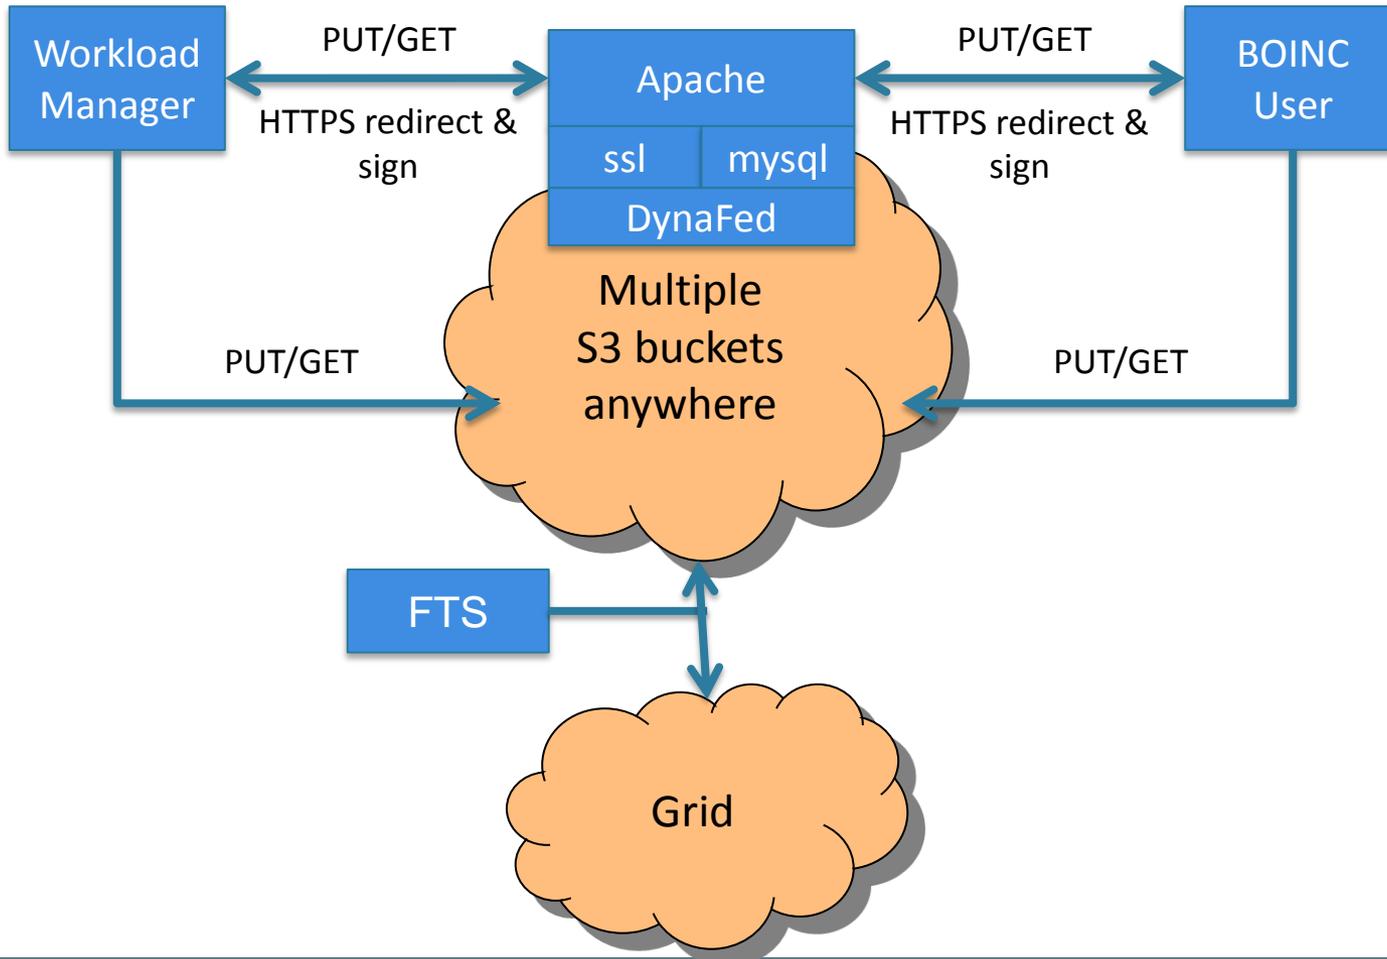
Authentications

- **BOINC users are authenticated by BOINC**
 - In the VM, BOINC credential provided via /dev/fd0
 - BOINC_ID
 - BOINC_AUTHENTICATOR
 - A BOINC Project DB provides the user Identity (IDP)
 - MySQL
 - User Table
- **The Grid (and FTS) speaks X509 with VOMS...** how to match that ?
- BOINC users need to receive the Job desc AND write the output
- Grid agents (FTS) need to read what the BOINC user wrote

Solution: the Data Bridge

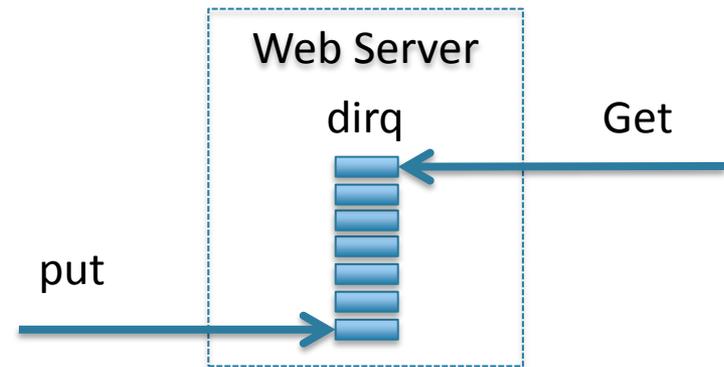
- The Data Bridge is an authentication bridge AND a shared authorized cloud-based storage
 - Uses **mod_auth_mysql** to recognize BOINC users
 - Maps username/password to DB table
 - AuthMySQLUserTable user
 - AuthMySQLNameField id
 - AuthMySQLPasswordField authenticator
 - Uses **mod_gridsite/mod_ssl** to recognize X509 with VOMS
 - That's normally the WebDAV entry point of the Dynamic Federation
 - Uses the **Dynamic Federation** on S3 to provide properly delegated redirections
 - Together with authorization, high performance data discovery and extreme scalability
- Based on Apache, reusing Web components and cloud storage

The Data Bridge



Data Bridge-based Message Queue

- Messaging services ActiveMQ-like do not support BOINC authentication
 - Not clear if it is possible or worthwhile to provide functionality
- Simple queuing implemented on the Data Bridge itself as a Web app
 - A shared storage that already authenticates clients properly
- Two simple cgi scripts
 - put-job.cgi
 - get-job.cgi
- Simple file-based queue
 - python-dirq
- Job descriptions from Workload mgmt
- Generic, supports arbitrary info in the queue
- No additional systems!



Dynamic HTTP Federations

- *Dynafed* implements federated storage over HTTP/WebDAV
 - In testing in LHCb and Canada (ATLAS)
 - Federates WebDAV or S3 enabled storage systems
 - Apache front end, no catalogues needed
- Can be used as a data bridge, in conjunction with S3's delegation mechanism
 - S3 storage backend(s)
 - Acts as a security gateway between X509 or *BOINC Auth*
 - Authorized clients then redirected directly to the storage
 - Quick response time (low latency, thousands reqs/s)
 - Robustness (automatic RT checking of endpoints status)
 - Great scalability potential
 - Global system, smart replica selection (availability, proximity)
- <http://svnweb.cern.ch/trac/lcgdm/wiki/Dynafeds>
- <http://federation.desy.de/>

Aggregation of S3 buckets

This is what we want to see as users

Sites remain independent and participate to a global view

All the metadata interactions are hidden and cached on the fly

For S3, clients are redirected using S3 signatures

The result looks like a normal WebDAV big repo

`/dir1`
`/dir1/file1`
`/dir1/file2`
`/dir1/file3`

With 2 replicas

S3 bucket 1

`.../dir1/file1`
`.../dir1/file2`

S3 bucket 2

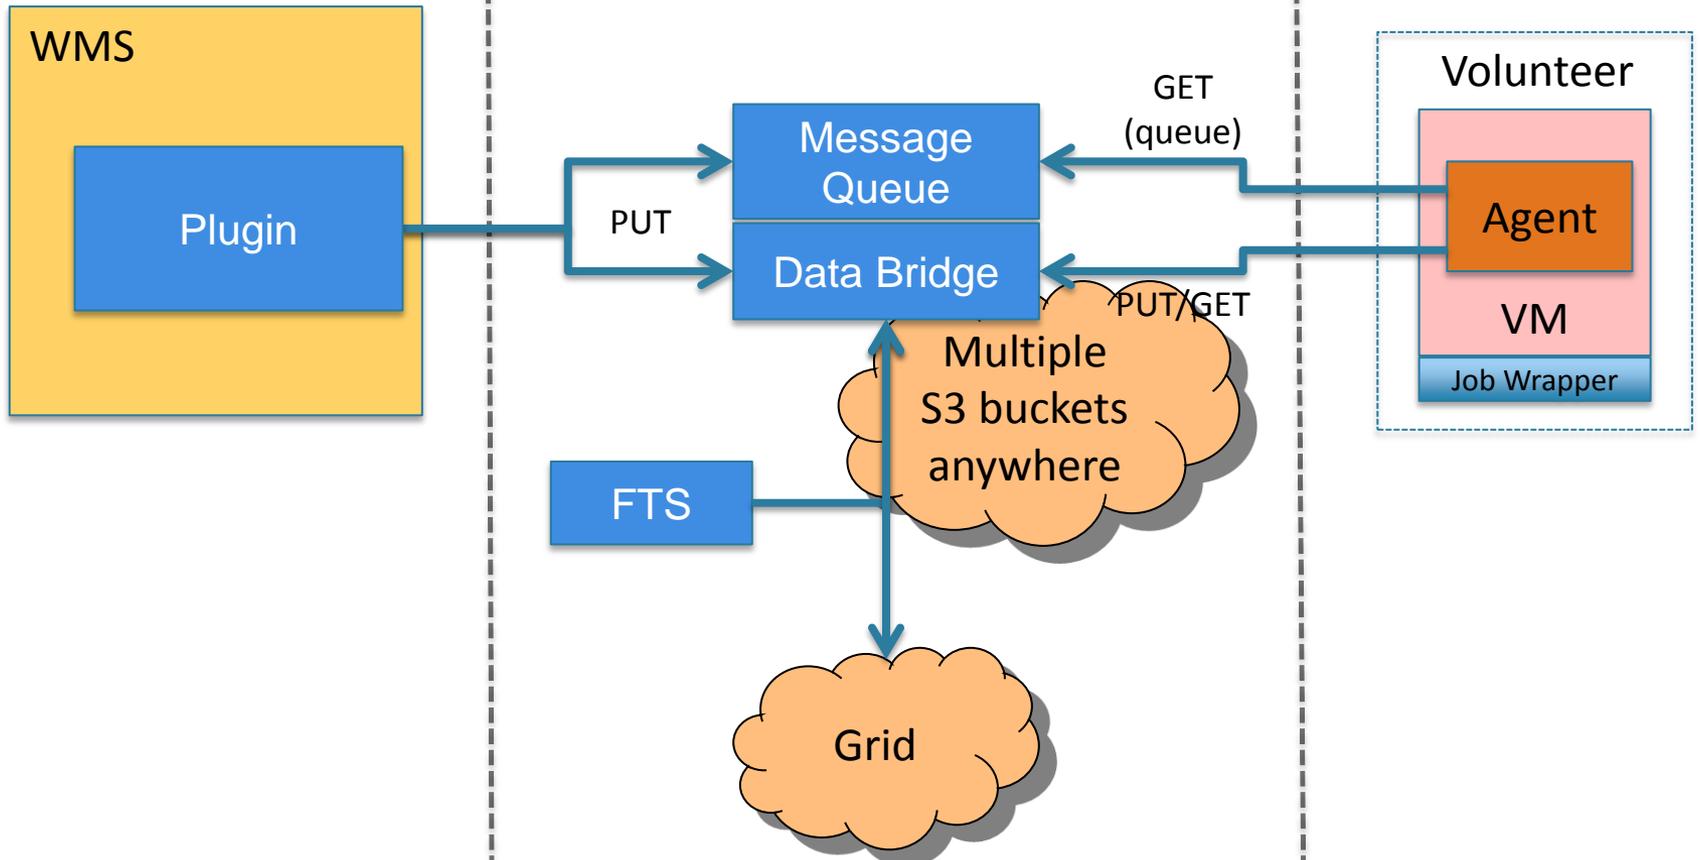
`.../dir1/file2`
`.../dir1/file3`

Final Architecture

VO
(X509/VOMS)

Infrastructure

Volunteer
(BOINC auth)



Adoption

- Can be adopted by other BOINC-based projects
- Build upon vLHC@home and the data bridge as a platform
- Require:
 - WM to POST job description to the data bridge message queue
 - Stage input data to the data bridge's input *bucket* (if needed)
 - CernVM3 Image
 - Contextualized including CVMFS configuration
 - Credentials read from /dev/fd0
 - BOINC_ID and BOINC_AUTHENTICATOR
 - Job Agent
 - Method to get job description
 - GET input from the data bridge
 - Run job
 - PUT output on the data bridge
 - Read data from the data bridge's output path (stored in a *bucket*)

Rollout Plans

- CMS@Home
 - Pioneered the adoption of the data-bridge
 - Will hopefully enter the beta testing phase soon
- Test4Theory
 - Plan to migrate to the data bridge within the next 2 months
 - To address co-pilot support issues
- Beauty@Home
 - Currently integrating the data-bridge
 - This will solve the x509 credential distribution issue
 - Allowing to open the project up to the public

Summary

- The Data Bridge spans auth domains
- Reuses HTTP federation component with S3
 - Friendly feel, quick and scalable
- Honours X509/VOMS and BOINC auth at the same time
 - Can honour any other apache-based authentication
- Secure S3 delegation mechanism
- Web tech used for the job descriptions queue
- Infrastructure provides a VM image along with a job agent
- Adopted by CMS@Home, Test4Theory, Beauty@Home
- Generic, towards a platform for volunteer computing
- Next: moving to production
- Next: experience with Single Sign On auth (CERN, Google, Facebook...) and bridge with the Grid/VOMS