

# OASIS



CernVM/CvmFS workshop  
5 - 6 March, CERN

# Contributors

Brian Bockelman: project manager

Dave Dykstra: FNAL Stratum-I operator, developer

Tony Tiradani: FNAL

Jose Caballero: project coordinator, lead developer

John DeStephano: BNL Stratum-I operator

John Hover: BluePrint liaison

Vince Neal: GOC

Rob Quick: GOC

Scott Teige: GOC operations manager

# What is OASIS?

- This talk covers the current version in production (1) and the coming version (2.0.x)
- OASIS stands for “OSG Application Software Installation Service”.
- Service offered by OSG to its community to facilitate software distribution, in an easy and reliable way.
- Provides a end-to-end set of services to OSG VOs: repository service, Stratum-0 service, Stratum-1 service, and client.
- Provides for a command line suite that encapsulates the internal technical details of the actual underlying technology used to distribute files: currently CvmFS. Allows for changes in the technology in a transparent way.

# OASIS architecture

The word OASIS can mean different things, depending on the purpose being served:

<b>OASIS as a Service (OaaS)</b>	is the central service provided by OSG Operations
<b>OASIS as a Repository (OaaSR)</b>	is the actual <i>oasis.opensciencegrid.org</i> CvmFS repository
<b>OASIS as a Federation (OaaSF)</b>	is the set of CvmFS repositories being distributed in OSG and the Stratum-Is used for that.
<b>OASIS as a Product (OaaSP)</b>	it is the package containing the code

# OASIS as a Service

- OASIS has been in production for OSG as a central service at the GOC since March 2013. And became very popular service since then.
- GOC infrastructure includes an interactive login host, a CvmFS Repository, one Stratum-0, and one Stratum-1.
- Two other Stratum-1s (at FNAL and BNL) completes the OASIS/OSG CvmFS architecture.
- There is in place a formal procedure for new VOs to register and request using it. Their information is recorded in the OSG Information Management System (OIM) by GOC staff.

# OASIS as a Service

- VO content managers log in interactively on a login host at the GOC, copy the new files, and run an oasis command to trigger the publishing. They log in with gsissh, and they are automatically given a UNIX ID based on their grid credentials and the information recorded in OIM.
- VO content managers just copy the files into their scratch area. A separate process running in the background (a cron job) is in charge of transferring those files to the CvmFS repository host and publish.
- It serializes the data transfer and publishing steps.

# OASIS as a Service

Currently based on CvmFS 2.0 Stratum-0 and CvmFS 2.1 Stratum-1s

Single repository: *[oasis.opensciencegrid.org](http://oasis.opensciencegrid.org)*

Serving 23 Virtual Organizations.

Distributes at the time this slide was written +11 Millions of files, +1.7 TB of data.

+3400 publishing operations.

For performance, it includes ~2000 nested catalogs.

(Some plots and more specs data in the backup slides)

# OASIS as a Service

The 3 replica host (GOC, BNL and FNAL) monitored via awstats:

Last Update: 09 Dec 2014 - 19:08

Reported period: Dec 2014 OK

Back to main page

Filter : Exclude filter : OK

**Hosts**

Total : 587 Known, 11 unknown (unresolved ip) - 597 Unique visitors

Hosts	Pages	Hits	Bandwidth	Last visit
lcgstratum01.gridpp.rl.ac.uk	354903	354903	29.67 GB	09 Dec 2014 - 18:37
cithep251.ultralight.org	71937	71937	22.16 GB	09 Dec 2014 - 19:07
admin.swin.edu	65422	65422	61.64 MB	09 Dec 2014 - 19:08
squid1.rcac.purdue.edu	56307	56307	22.28 GB	09 Dec 2014 - 19:08
red-gridftp1.unl.edu	45359	45359	15.59 GB	09 Dec 2014 - 19:08
cithep249.ultralight.org	41852	41852	15.50 GB	09 Dec 2014 - 19:07
red-gridftp3.unl.edu	41721	41721	14.61 GB	09 Dec 2014 - 19:07
red-squid1.unl.edu	40020	40020	10.72 GB	09 Dec 2014 - 19:08
red-gridftp2.unl.edu	38750	38750	12.42 GB	09 Dec 2014 - 19:08
cvmfs-1.t2.ucsd.edu	34329	34329	15.96 GB	09 Dec 2014 - 19:07
cvmfs1.fnal.gov	32290	32290	4.49 GB	09 Dec 2014 - 18:57
workflow.isi.edu	29741	29741	5.31 GB	09 Dec 2014 - 19:07
cayenne.hep.wisc.edu	28633	28633	9.48 GB	09 Dec 2014 - 19:07
squid2.cmsaf.mit.edu	28372	28372	18.55 GB	09 Dec 2014 - 19:07
capiscum.hep.wisc.edu	26926	26926	9.12 GB	09 Dec 2014 - 19:07
squid1.cmsaf.mit.edu	24209	24209	18.76 GB	09 Dec 2014 - 19:07
cvmfs08.racf.bnl.gov	24164	24164	4.48 GB	09 Dec 2014 - 19:06
cvmfs07.racf.bnl.gov	24157	24157	4.48 GB	09 Dec 2014 - 19:06

Last Update:

Reported period:

Back to main page

Filter :

Total : 8754 K

206-167-180.67.cloud.com	732419	732419	109.95 GB	28 Feb 2015 - 23:59
192.153.161.25	717338	717338	117.57 GB	28 Feb 2015 - 23:59
192.153.161.201	713399	713399	116.83 GB	17 Feb 2015 - 20:12
206-167-180.68.cloud.com	679119	679119	97.60 GB	28 Feb 2015 - 23:59
brux4.hep.brown.edu	626721	626721	110.26 GB	28 Feb 2015 - 23:59
nova-qc.dair-atir.canarie.c	613061	613061	106.97 GB	28 Feb 2015 - 23:59
net2.rc.fas.harvard.edu	608880	608880	107.19 GB	28 Feb 2015 - 23:59
nat105.farm.particle.cz	606528	606528	107.39 GB	28 Feb 2015 - 23:59
ec2-54-162-119-105.compu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
zam1205.zam.kfa-juelich.k	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-80-65-112.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-144-55-233.compu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-146-39-234.compu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
zam1206.zam.kfa-juelich.k	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-81-173-96.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
130.199.185.42	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-211-33-59.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
130.199.185.118	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-144-45-19.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-198-23-83.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-91-234-68.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-184-73-76-102.compu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-23-20-192-115.compu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-145-89-91.comput	587136	587136	100.05 GB	28 Feb 2015 - 23:59
atl-prod09.slac.stanford.edu	587136	587136	100.05 GB	28 Feb 2015 - 23:59
kraken01.westgrid.ca	587136	587136	100.05 GB	28 Feb 2015 - 23:59
t1-squid5.triumf.ca	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-54-145-82-49.compute-1.amazonaws.com	587136	587136	100.05 GB	28 Feb 2015 - 23:59
ec2-107-22-30-153.compute-1.amazonaws.com	587136	587136	100.05 GB	28 Feb 2015 - 23:59
cache2.aglt2.org	587136	587136	100.05 GB	28 Feb 2015 - 23:59
cache0.aglt2.org	587136	587136	100.05 GB	28 Feb 2015 - 23:59
cache1.aglt2.org	587136	587136	100.05 GB	28 Feb 2015 - 23:59
cache3.aglt2.org	587136	587136	100.05 GB	28 Feb 2015 - 23:59
lcg-admin4.scinet.utoronto.ca	587136	587136	100.05 GB	28 Feb 2015 - 23:59

Bandwidth	Last visit
372.68 GB	28 Feb 2015 - 23:57
640.49 GB	28 Feb 2015 - 23:59
147.01 GB	28 Feb 2015 - 23:52
169.20 GB	28 Feb 2015 - 22:21
444.50 GB	28 Feb 2015 - 20:20
245.09 GB	18 Feb 2015 - 04:24
209.42 GB	18 Feb 2015 - 04:29
73.14 GB	28 Feb 2015 - 23:59
116.31 GB	28 Feb 2015 - 23:59
967.46 MB	28 Feb 2015 - 23:59
234.18 GB	28 Feb 2015 - 21:30
107.84 GB	28 Feb 2015 - 23:59
70.46 GB	18 Feb 2015 - 05:38
69.17 GB	28 Feb 2015 - 22:28
105.12 GB	28 Feb 2015 - 23:59
170.46 GB	28 Feb 2015 - 23:59
97.87 GB	28 Feb 2015 - 23:59
85.52 GB	18 Feb 2015 - 04:33
94.09 GB	28 Feb 2015 - 23:59
80.89 GB	28 Feb 2015 - 23:56
81.04 GB	28 Feb 2015 - 23:59
80.01 GB	28 Feb 2015 - 23:59
80.20 GB	28 Feb 2015 - 23:59
86.14 GB	28 Feb 2015 - 23:59
80.18 GB	28 Feb 2015 - 23:59
84.45 GB	28 Feb 2015 - 23:59
67.58 GB	28 Feb 2015 - 23:59
66.00 GB	28 Feb 2015 - 23:59
64.52 GB	25 Feb 2015 - 22:55
67.07 GB	28 Feb 2015 - 23:59
56.75 GB	28 Feb 2015 - 23:59
104.36 GB	28 Feb 2015 - 23:59
105.30 GB	17 Feb 2015 - 22:03



# OASIS as a Federation

The Federation is:

- the set of repositories serving content to OSG sites
- the set of replica hosts distributing (and filtering) that content

Includes all repositories with domain name “.opensciencegrid.org” (not originated at GOC):

<i>darkside.opensciencegrid.org</i>	<i>fermilab.opensciencegrid.org</i>
<i>icecube.opensciencegrid.org</i>	<i>lsst.opensciencegrid.org</i>
<i>mu2e.opensciencegrid.org</i>	<i>usatlast3.opensciencegrid.org</i>

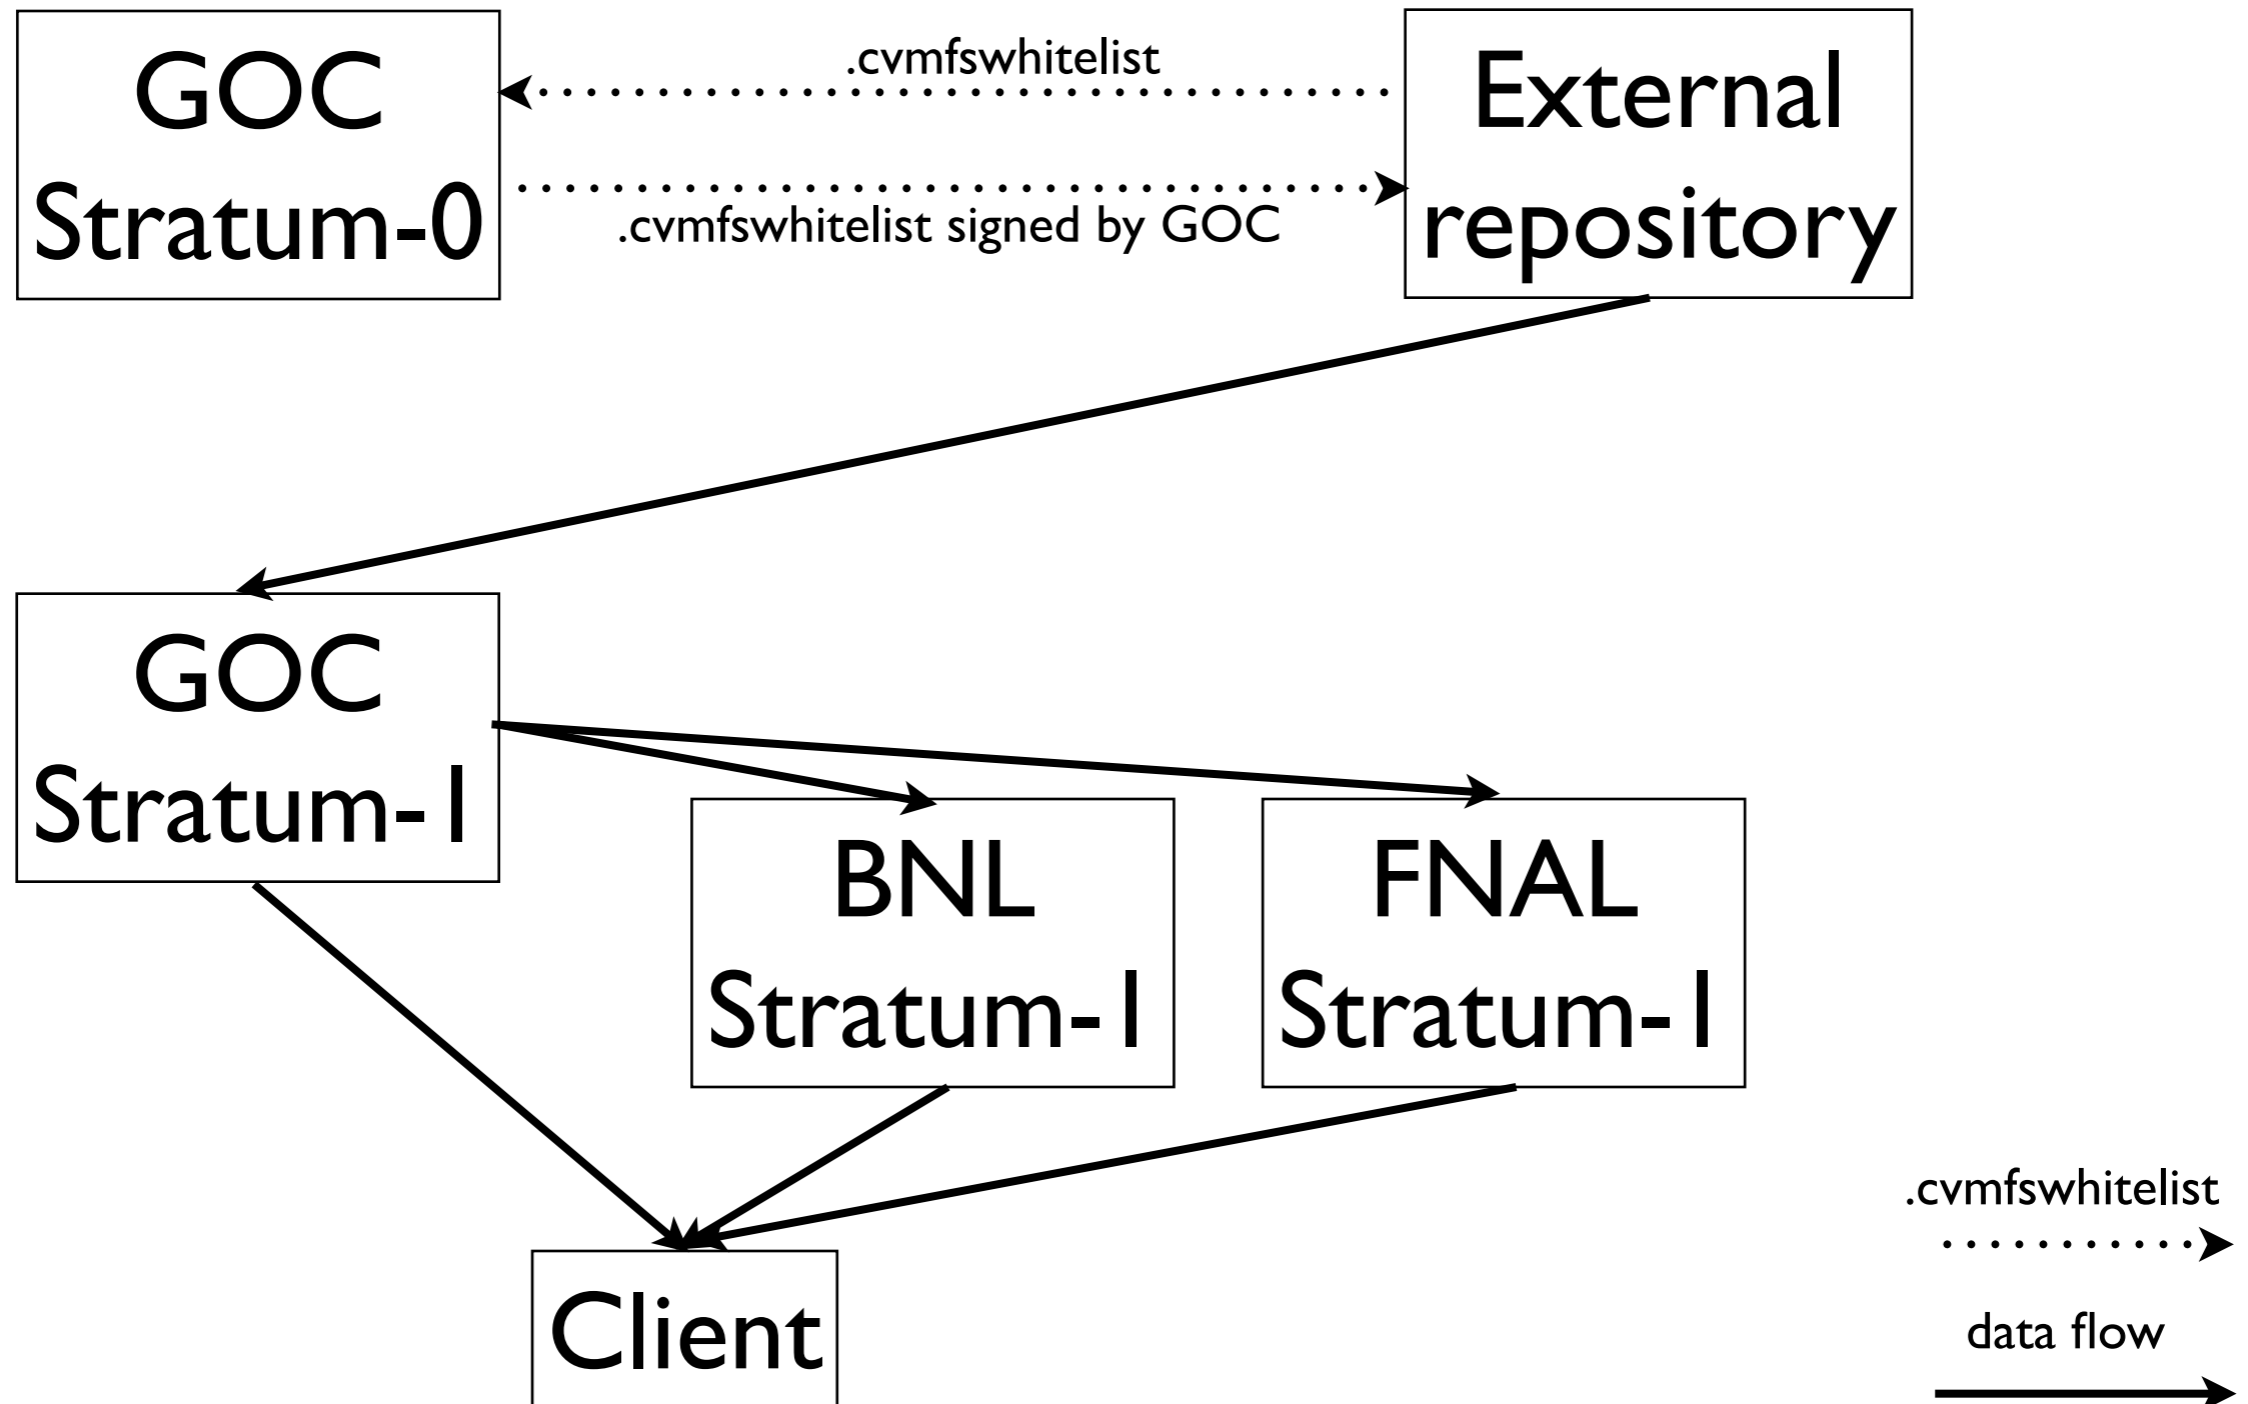
Other foreign repositories can be added to the federation if an OSG stakeholder requests it:

<i>auger.egi.eu</i>	<i>glast.egi.eu</i>
---------------------	---------------------

# OASIS as a Federation

- The GOC provides a signing service for the remote repositories' `.cvmfswhitelist`, so only one public key is needed on the client for all the `“.opensciencegrid.org”` domain.
- The GOC replicates the repositories on its Stratum-I, and then the BNL and FNAL Stratum-I's replicate from there. This data flow architecture allows for emergency content blanking.
- Repositories are automatically added to the Stratum-I's when their URLs are registered by GOC staff in OIM.
- Similar architecture is currently in place to distribute `“.egi.eu”` repositories content (without the signing step).

# OASIS as a Federation



# OASIS as a Federation

Since the non-GOC repositories are outside of the GOC's control, OSG Security required the ability for the GOC to blank them in case of a security emergency where dangerous content is distributed and the repository owner is not available to respond quickly.

This is done by replacing the replica on the GOC's Stratum-I with a newly created blank CvmFS repository, signed by the OSG key and with a repository publication number slightly higher than the original replica.

Clients are configured so repositories in the ".egi.eu" domain accept both the original EGI signing key and the OSG signing key, and they read only from the OSG Stratum-I's.

# OASIS as a Product

Starting with version 2.0.x (under testing), OASIS code is being distributed as an RPM.

Improvements that OASIS 2 will include:

- Python. Plugins architecture. Behavior driven by configuration files. Background process is a daemon instead of a cronjob.
- Improved the command line suite, including several commands to make sys-admins work easier.
- The ability to pass optional probes before the new content is actually published. Probes can just raise a warning, stop the publication, or even fix the content. Probes can enforce maximum filesize, create new cvmfs subcatalogs, look for non-relocatable software, etc. Probes can be applied to every VO content, or on a VO-by-VO basis.
- Allows, besides the login host model, for a grid interface setup.

# Grid Interface

As mentioned, OASIS 2 allows for a grid interface architecture.

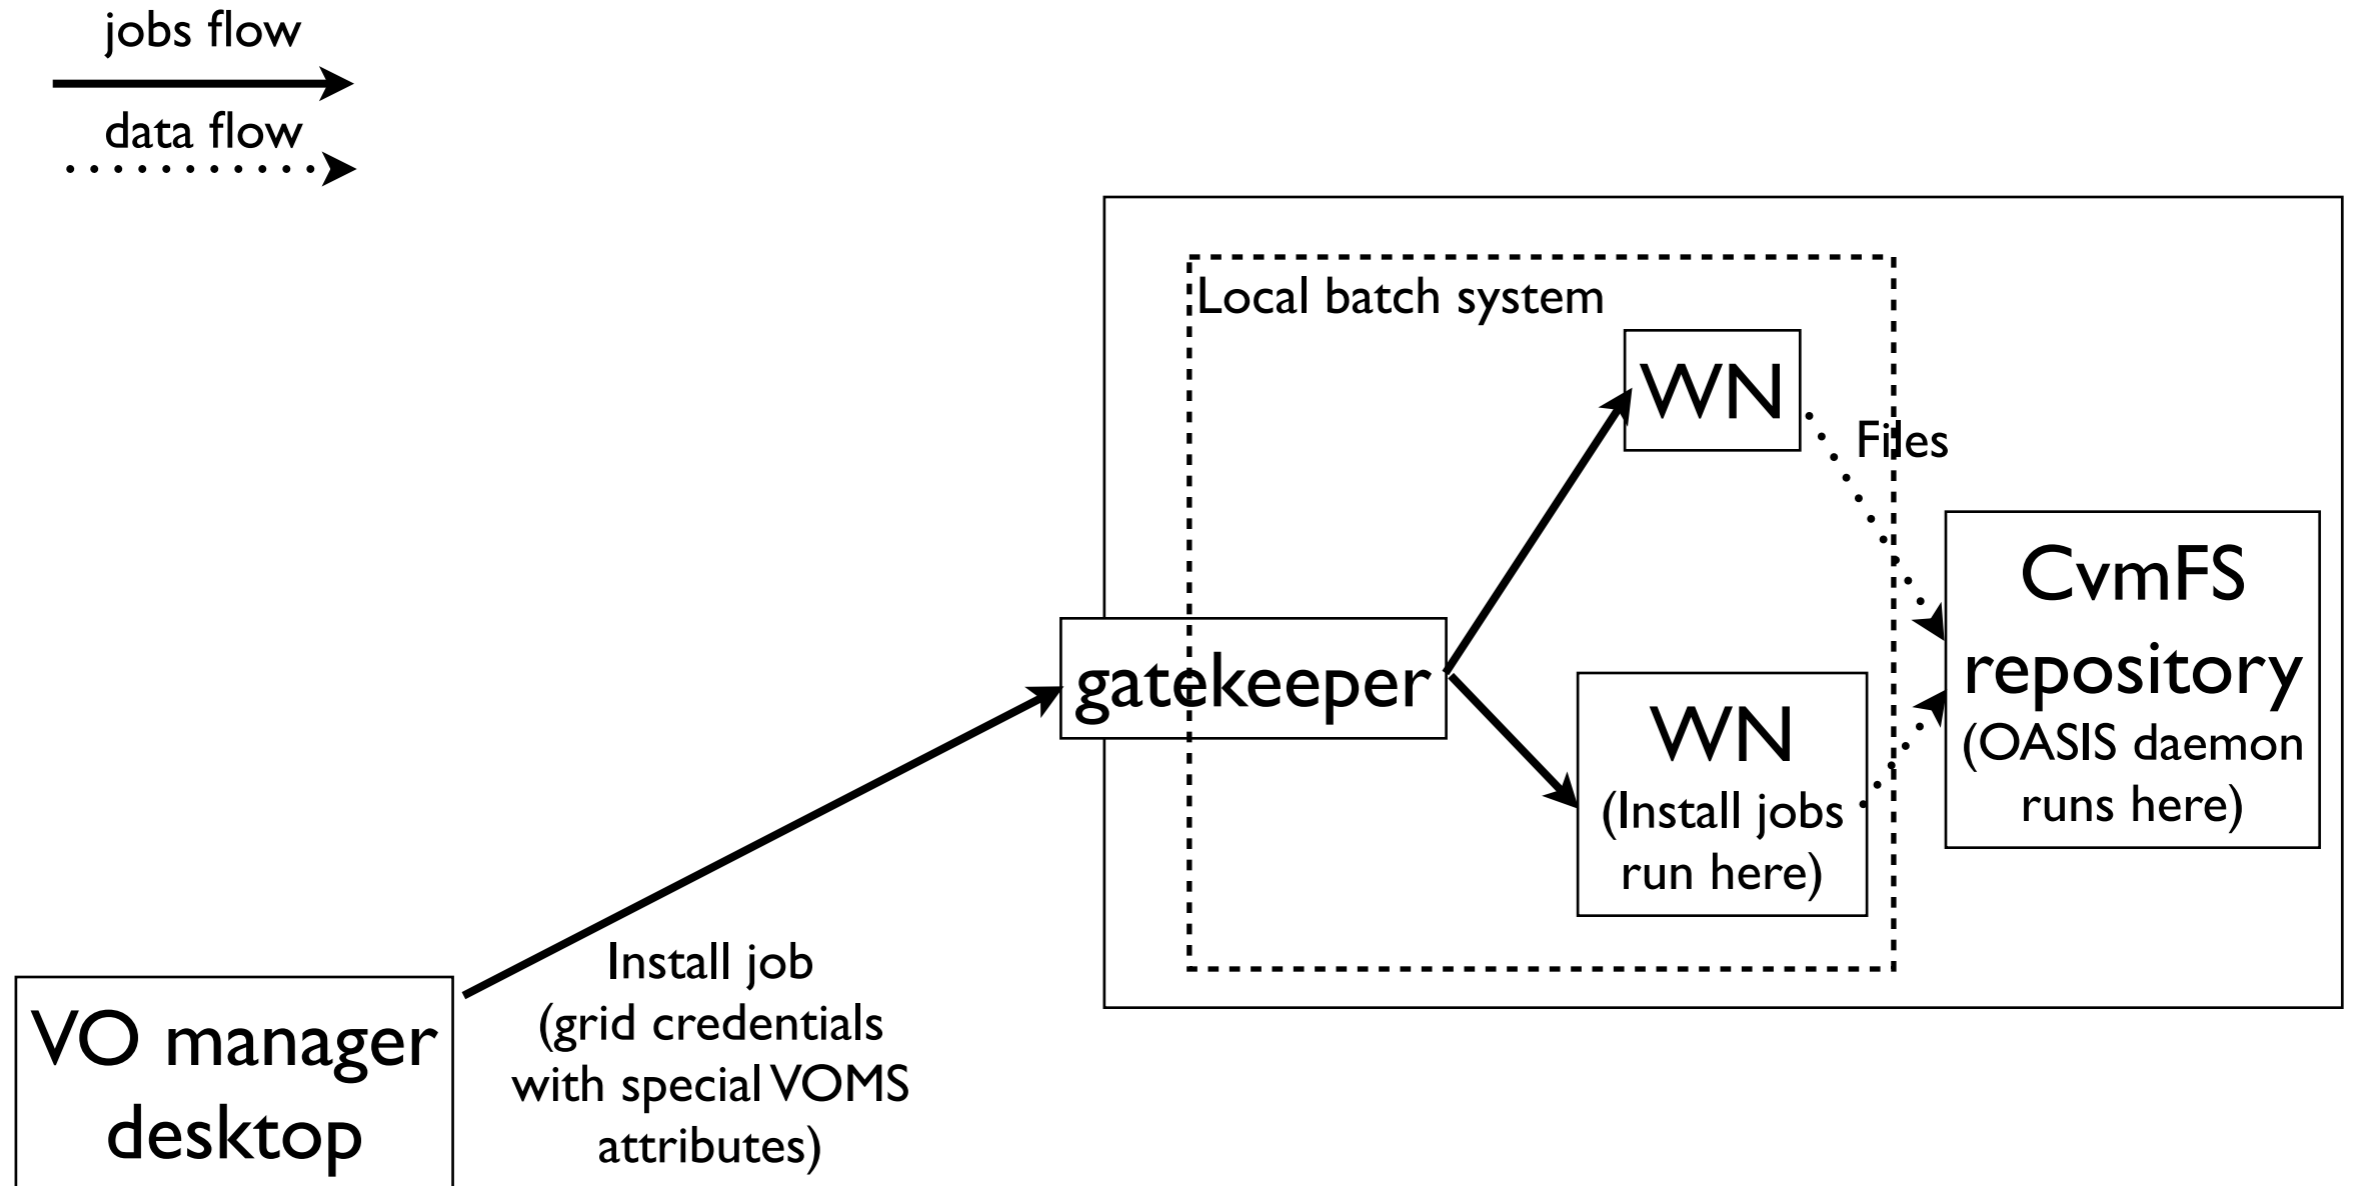
For example, there is one for a Stratum-0 at BNL.

VO content managers (carrying specific VOMS attributes) submit a grid job to a dedicated gatekeeper. The job just places new files into a specific directory and that's it. OASIS will notice the presence of these new files and proceed.

VOs have been doing software installation this way for many years. It should look natural to them.

```
$ cd $OSG_APP/myvo/  
  
$ wget http://myinstitution.page.edu/.../newcontent.tgz  
  
$ tar xzf newcontent.tgz  
  
$ rm newcontent.tgz
```

# Grid Interface



# Future work

Future work will be mostly driven by customers requests (explicitly) and from operations experience to improve the user experience and performance. The list of items includes, among other things:

- Improvements in the CLI.
- Writing probes.
- Monitoring (mostly using the CvmFS python bindings).
- Automated update mechanism for the client.
- Puppet mode for automatic configuration.



# For the CvmFS team...

- A simpler mechanism for keys/configuration distribution to make the process of adding repositories to the federation easier.
- More features in the CvmFS API. For example, to allow changes in the revision number from command line.
- To facilitate the deployment of OASIS, it would be nice if CvmFS deployment were easier: the requirement of using AUFS kernel could be a showstopper. CERN needs to commit to update the RPM when needed fast enough.
- Also, deployment of the client w/o root privileges is becoming a need to allow opportunistic platforms. Some alternatives are being discussed.
- Split the publishing into several steps. Or allow plugins code to be run while processing the filesystem and calculating hashes (via callbacks).

# Documentation

**About OASIS:** <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/OasisService>

**Updating content on OaaS:** <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/UpdateOasis>

**Using the grid interface:** <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/OasisUpdateGatekeeperInterface>

**Installing OaaP:** <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/OasisDeployment>

**Joining OaaF:** <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/OasisExternalRepositories>

**Contact:** oasis-dev @ opensciencegrid.org

**Backup slides**

## **Login host specs:**

- 1 CPU
- 1 GB RAM
- 43 GB used


## **Replica host specs:**

- 1 CPU
- 2 GB RAM
- 43 GB used
- Bandwidth:
  - public: 461.16 Mbps out, 1.0 Mbps in
  - private: 1.43 Mbps out, 26.35 Mbps in

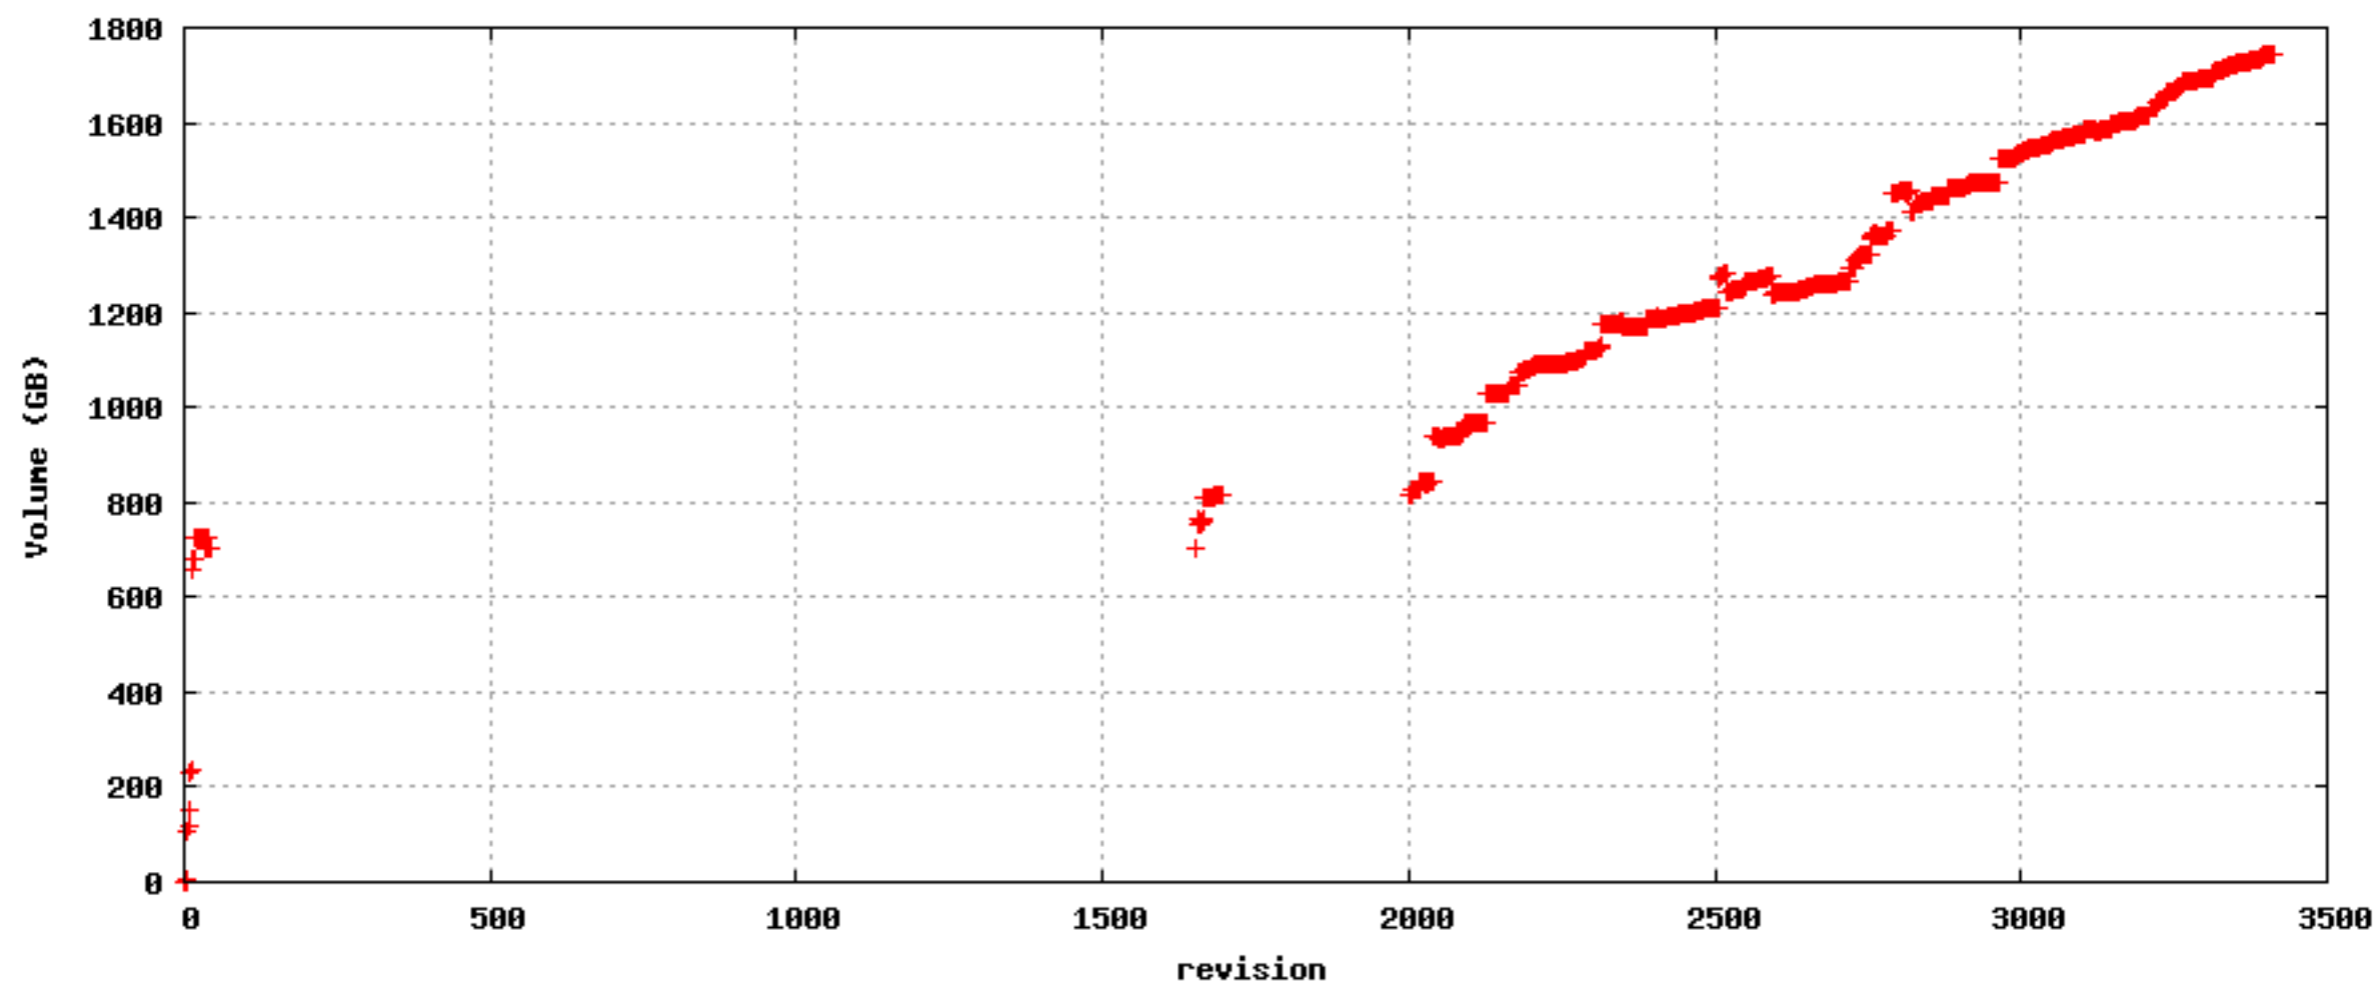
## **Repository host specs:**

- 2 CPUs
- 4 GB RAM
- 6.6 TB used
- Bandwidth:
  - public: 983.14 kbps out, 11.07 kbps in
  - private: 99.73 kbps out, 236.99 kbps in


ORISIS utilization




ORISIS utilization



ORASIS utilization




ORASIS utilization





ORISIS utilization



ORISIS utilization

