

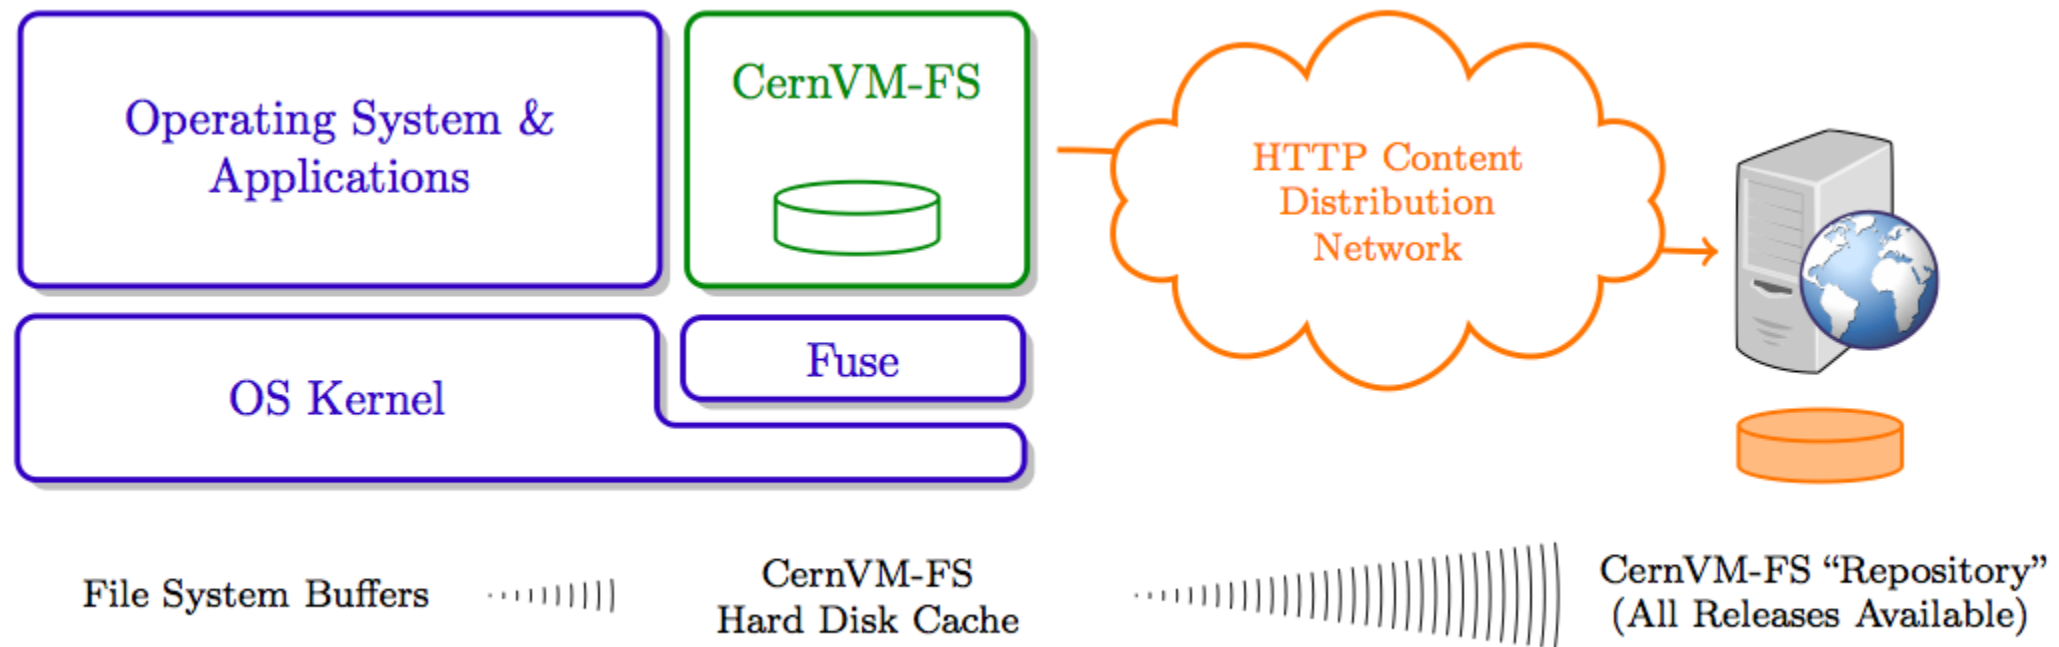


NOvA CVMFS Experience

Gavin S. Davies, Iowa State University
Raphael Schroeter, Harvard University
On behalf of the NOvA Collaboration

FIFE Stakeholders Meeting
August 7th 2013

What is CVMFS?




- CernVM-FS client provides a virtual file system that loads data only on access
- In this example, all releases of a software package (such as NOvA-ART) are hosted as a CernVM-FS repository on a web server
- CVMFS uses HTTP for data transfer
- No need to install/build NOvA-ART locally on offsite clusters

A Typical CVMFS Setup

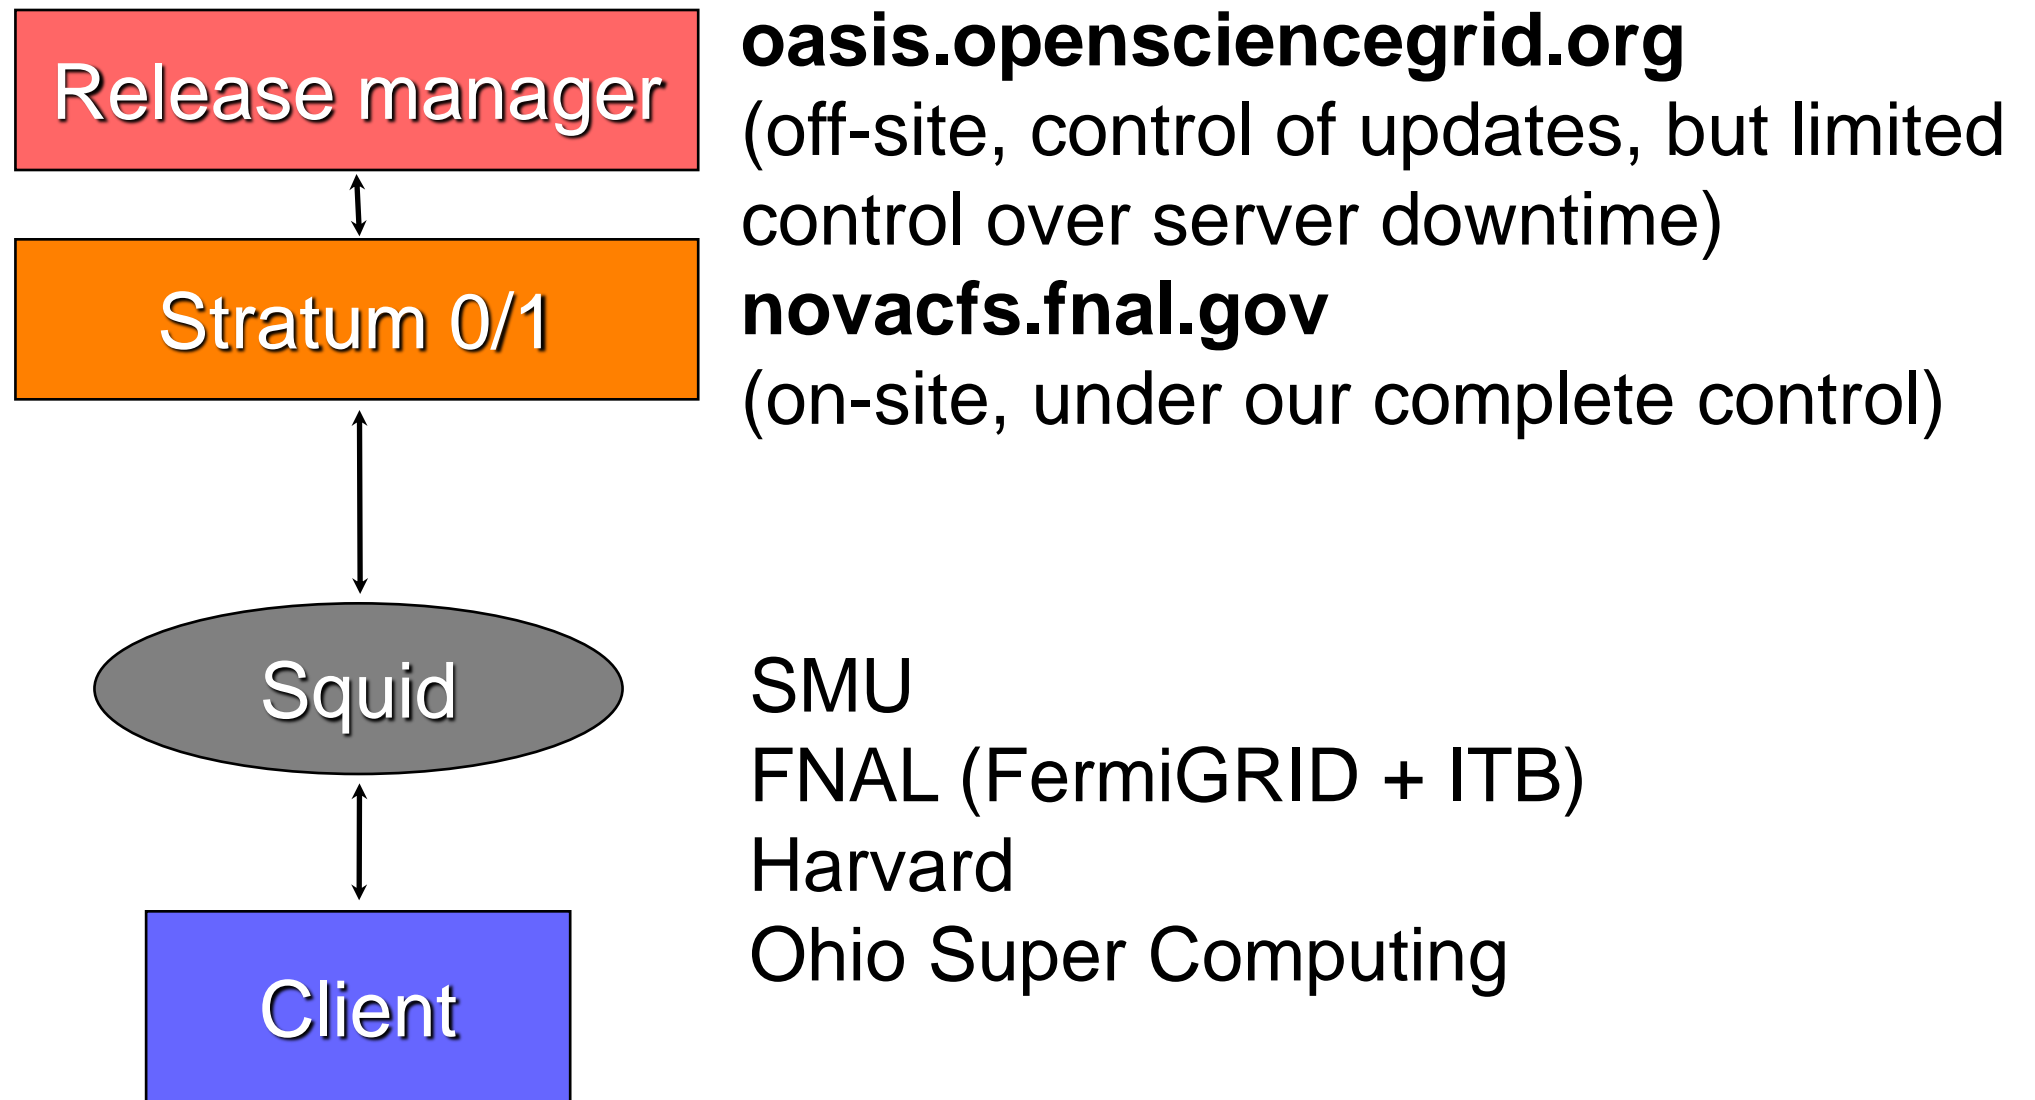
- CVMFS is a file system with a single source of data (repository **Stratum 0**) The stratum 0 is maintained by a dedicated **release manager** machine, where a read-writable copy of the repository is available The CVMFS server tool kit is used to publish the current state of the repository on the release manager machine
- Repositories are replicated to a handful of web servers (**Stratum 1**)
- On the client side, repository data cached by standard web proxies such as **Squid**
- Squid is a caching proxy for the Web supporting HTTP, HTTPS, FTP, ...
- Bandwidth is reduced and response times improved by caching and reusing frequently-requested 'web pages'

Server Structure



- The Stratum 0 repository server is able to serve clients directly
- A large number of clients is better served by a set of Stratum 1 replica servers, that may be distributed in different locations
- Multiple Stratum 1 servers improve the reliability, reduce the load, and protect the Stratum 0 master copy of the repository from direct accesses
- Stratum 0 server, Stratum 1 servers and the site-local proxy servers can be seen as **content distribution network**

NOvA CVMFS



- Currently two different CVMFS servers,
 - One in-house at FNAL, other Open Science Grid (OSG) each hosting both stratum 0/1 servers
- Four clients so far: FNAL (access to all servers mounted on FermiGRID & ITB, NOvA VMs also can access them), SMU, Harvard & Ohio have access to OSG server

Novasoft and CVMFS

- NOvA-ART, Novasoft deployed on both CVMFS servers to test latest production release (SVN)
- Our software guarantees to work in any environment
 - Just a matter of packaging the right flavours (SLF5/6 etc)
 - Work in progress for Mac/Ubuntu flavour builds.
- What do we deploy on the CVMFS servers?
 - ART framework externals: ups/upd, art_externals, art_suite, nu_extras
 - ROOT, GEANT4, GENIE, PYTHIA etc.
 - ART group provide the SLF5/6 distribution tarballs
 - novasoft which includes nutools packages
 - IF common packages: SimulationBase, EventGeneratorBase etc
 - `svn+ssh://p-nusoftart@cdcvs.fnal.gov/cvs/projects/nusoftsvn -s nutools`
 - Genie FLUX files
 - Everything is simply distributed via tarballs

Novasoft and CVMFS

- Simple setup Novasoft script on the client side:

```
function setup_novaoffline_cvmfs_fnal {  
    echo -en "\033[00;33mNova Offline Environment Enabled \033[00;32m(CVMFS  
Base)\033[00m\n"  
    export CVMFS_DISTRO_BASE=/cvmfs/novacfs.fnal.gov  
    source $CVMFS_DISTRO_BASE/novasvn/srt/srt.sh  
    export EXTERNALS=$CVMFS_DISTRO_BASE/externals  
    source $CVMFS_DISTRO_BASE/novasvn/setup/setup_novasoft.sh "$@"  
}
```

- `setup_novaoffline_cvmfs_fnal -r S12-11-16` will setup novasoft off the CVMFS server and give access to all the commands needed to run the nova code
- So far, only tested MC production
- MC samples successfully produced on all clients (single, cosmics and GENIE, all done interactively)
- **Part One:** Handled by jobsub (in-house IF job submission) glideins (user transparency, just command line options)
- **Part Two:** worker node receives job, sets up code, runs job
- **Part Three:** Handled by ifdhc (recognises on-site/off-site) to copy files back to FNAL
- Interactively can run everywhere – CVMFS part works perfectly!!
- Full chain (jobsub+cvmfs+ifdhc) works perfectly for FermiGRID and ITB and SMU

CVMFS and cache

- On every offsite cluster should have two or more Squid forward proxy servers that CVMFS can use (currently direct proxy for testing purposes)
- The forward proxies reduce the latency for the local worker nodes
- One Squid server on commodity hardware scales well for at least a couple of hundred worker nodes
- Typical cache required 4 - 8 GB of memory cache and 50 -100 GB of hard disk cache (NOvA request: 3GB memory, 50GB disk)
- Example: Currently to start a NOvA production job we load ~0.5 GB flux file, but then for multiple jobs that same file is available in the cache so is only loaded once for the first job
- Each experiment needs to be sure it has a streamlined method for transferring files so as not to saturate the cache

ART-MAC

- g-2 has a compiled version of ART for Mac OS X running off CVMFS:
- Perfect example to show the power of CVMFS
- art_suite-1.02.04
- compiled for OS X 10.7 (Lion), extra gymnastic to run on 10.8 (Mountain Lion)
- NOVA-ART currently built/compiled:
 - art_externals-0.06.04, art_suite-1.06.00, nu_extras-0.06.04
- OS X 10.6 (Snow Leopard), 10.7 and 10.8
- promising so far, work in progress
 - Likewise for Ubuntu
- On the user side: easy setup, only CVMFS and XCode required

Outlook/Status

- Glidein: submitting production jobs offsite is completely transparent to users – not globally tested but by select individuals
- Issue at Harvard is SL5 vs. SL6 (CentOS) – updated worker nodes to SL6, gatekeeper is SL5 so jobsub failed
 - Update gatekeeper, test jobsub
- Plan for NOvA is to have two build machines (we have one SLF5 – novabuild.fnal.gov, one in future on SLF6)
 - Already build software that lives on bluearc on novabuild01.fnal.gov
 - Changes (eventually) pushed directly to oasis and novacfs servers, no human interaction
- Working towards NOvA production (September?) to be run off CVMFS at FNAL and offsite
- ART-Mac, ART-Ubuntu, etc... majority of NOvA collaborators will be able to run novasoft on their local machines/laptops
- Truly global distribution of software to run on ones laptop
 - Simply change setup script to point to the right place
- Avoids bluearc, CVMFS works great!

End

- Big thanks to Raphael Schroeter (Harvard) who helped put the slides together and has been pushing the NOvA-CVMFS effort forward!

Documentation

- https://cdcvs.fnal.gov/redmine/projects/novaart/wiki/NOvA_Software_in_CVMFS
- https://cdcvs.fnal.gov/redmine/projects/novaart/wiki/Setting_up_NOvA_Software_on_your_Laptop