

# SOFTDRIVE.NL, CVMFS FOR THE MASSES

DENNIS VAN DOK

Generic Components of the eScience Infrastructure  
Ecosystem — 14th IEEE eScience Conference  
Amsterdam, Monday 2018-10-29

# GRID COMPUTING A.K.A. PAAS

Large scale common science infrastructure for high throughput batch computing.

- the only guaranteed environment is the base OS and some middleware
- no persistent local storage between jobs
- bring your own software

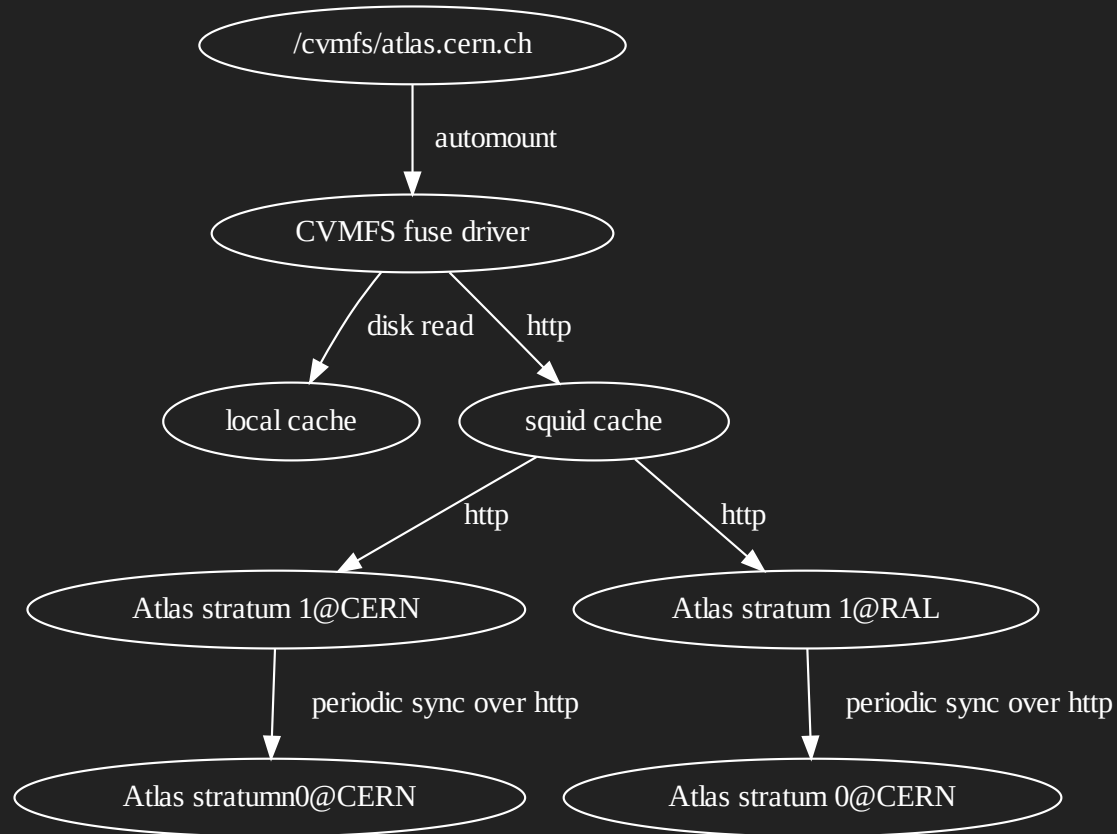
# CHALLENGES OF SOFTWARE DISTRIBUTION

- Bringing software with every job incurs much overhead
- Projects to develop common software distributions have a slow upgrade cycle
- Negotiating a locally writable software area for each site takes time, effort and coordination

# SOFTWARE DISTRIBUTION WITH CVMFS

- CVMFS spun off the the CERN Virtual Machine
- content delivery based on http
- data is distributed as objects referenced by hashes
- read-only, so trivial to replicate massively
- transactionally consistent indices
- garbage collected

# ARCHITECTURE



# DRAWBACKS

CVMFS is great for large organisations. But for small teams it can be a real challenge:

- set up and maintain a repository
- take care of a Stratum-0 server
- negotiate the replication at Stratum-1 sites
- negotiate with sites to include the repository in their CVMFS configuration

I imagined dozens of small e-science groups knocking on my door to get their repositories mounted.

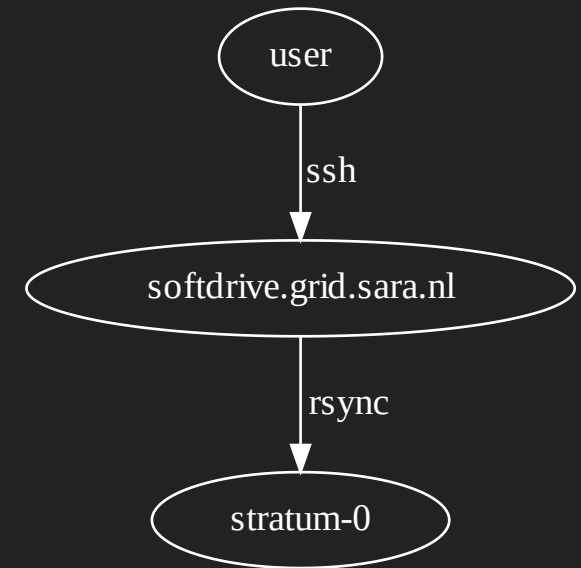
# OUR SOLUTION: SOFTDRIVE.NL

Nikhef and SURFSara have jointly set up [/cvmfs/softdrive.nl](https://cvmfs.softdrive.nl) to offer a single CVMFS repository for all e-science users in the Netherlands.

# ARCHITECTURE

The system consists of

- a user interface system, where users can log on (with ssh) and upload their software
- a Stratum-0 server which copies the user's files at regular intervals
- Stratum-1 at Nikhef and RAL
- mounted by default on all grid resources in the Netherlands



```
/cvmfs/softdrive.nl/z  
/cvmfs/softdrive.nl/c  
/cvmfs/softdrive.nl/p  
/cvmfs/softdrive.nl/s  
/cvmfs/softdrive.nl/t  
/cvmfs/softdrive.nl/a  
/cvmfs/softdrive.nl/r  
/cvmfs/softdrive.nl/f  
/cvmfs/softdrive.nl/k
```



# RULES

- User requests account at SURFSara
- Standard quota of 2GB (could be extended)
- Manage software on `softdrive.grid.sara.nl`
- Copy software to  
`/cvmfs/softdrive.nl/$USER`
- Run the `publish` command which touches the `softdrive.modified` file

# MECHANISM

- Automated rsync from Stratum-0 server at Nikhef
- Two stage process:
  1. rsync the softdrive.modified files
  2. rsync those directories with updated softdrive.modified files

# QUIRKS

Catalog size exploded when monitoring was put in place. The monitoring triggered an update every five minutes and thereby a completely new, full catalog of all files.

This was ultimately understood and remedied by making subcatalogs per user.

# USER EXPERIENCE

To complement the technical implementation, the total user experience was taken care of by having proper documentation, monitoring and guidance.

# DOCUMENTATION

The user documentation is right there when logging on to the system. The message of the day, printed for login shells, gives a summary of the workings of the system and how to publish data.

More extensive documentation was [written and placed on-line](#).

# MONITORING

End to end monitoring of the system is done by automatically triggering a change to the system every hour and measuring the time it takes for the data to reach a client machine. Alerts are raised if the delay reaches a certain threshold, prompting the technicians to inspect what went wrong.

# SUMMARY

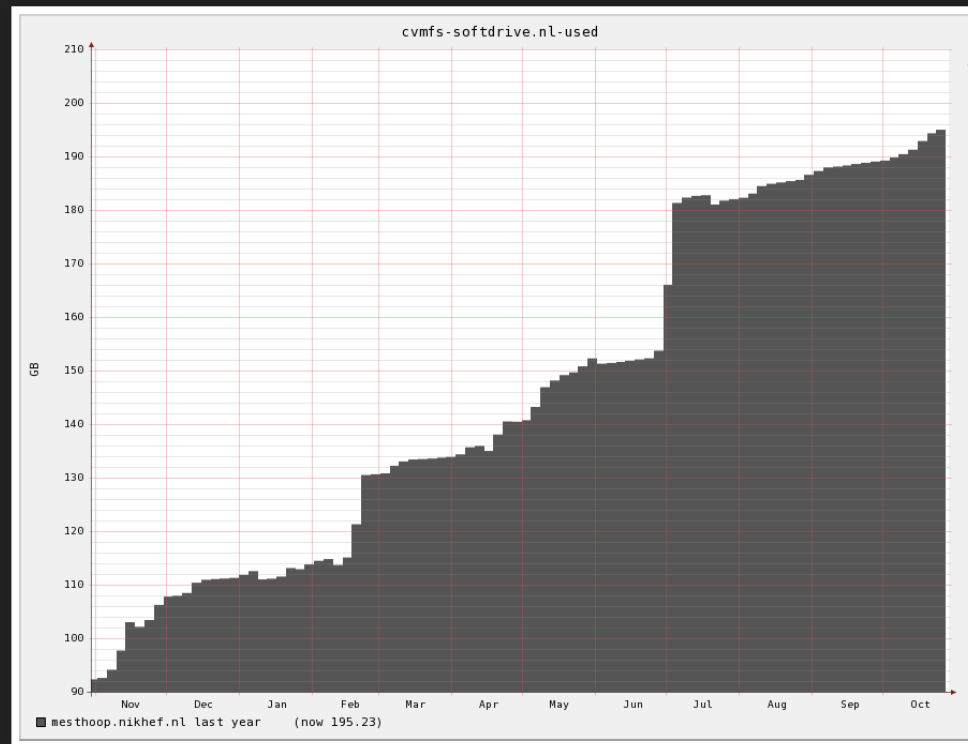
The softdrive model has proven to be successful; it is easy for users to maintain their own software; the software is lightweight and the maintenance burden on the administrators is very light.

There is no plan at this point to add more bells and whistles to the system.

Even as the PaaS infrastructure dwindles in favour of IaaS (infrastructure as a service), the CVMFS system could still be a viable component for delivering software.

# SOME NUMBERS

- 25 active users last 6 months
- 393k files, 178 GB





## INTERESTED?

Some other national grid infrastructures offer something similar to softdrive, but I've not heard of anyone interested in cloning our setup. If you have plans to provide CVMFS to your users, and would perhaps like to use (parts of) the softdrive system, don't hesitate to contact me.

# ACKNOWLEDGEMENTS

- Coen Schrijvers and colleagues at SURFSara for user documentation and monitoring.
- Catalin Condurache (RAL) for the fail-over Stratum-1.
- Ronald Starink for the initial setup of the CVMFS system at Nikhef.
- <http://doc.grid.surfsara.nl/en/latest/Pages/Advanced/>