# CVMFS @Nikhef DENNIS VAN DOK

CernVM virtual workshop 2021, 1–2 February 2021

# TABLE OF CONTENTS

- About Nikhef
- Computing @Nikhef
- CVMFS Servers
- CVMFS Clients
- Softdrive.nl
- Future outlook



# **ABOUT NIKHEF**

- High energy physics lab in the Netherlands
- Sub-atomic physics, accelerator-based and astroparticles
- "the smallest distance scale and the highest attainable energy"



## **PROJECTS**

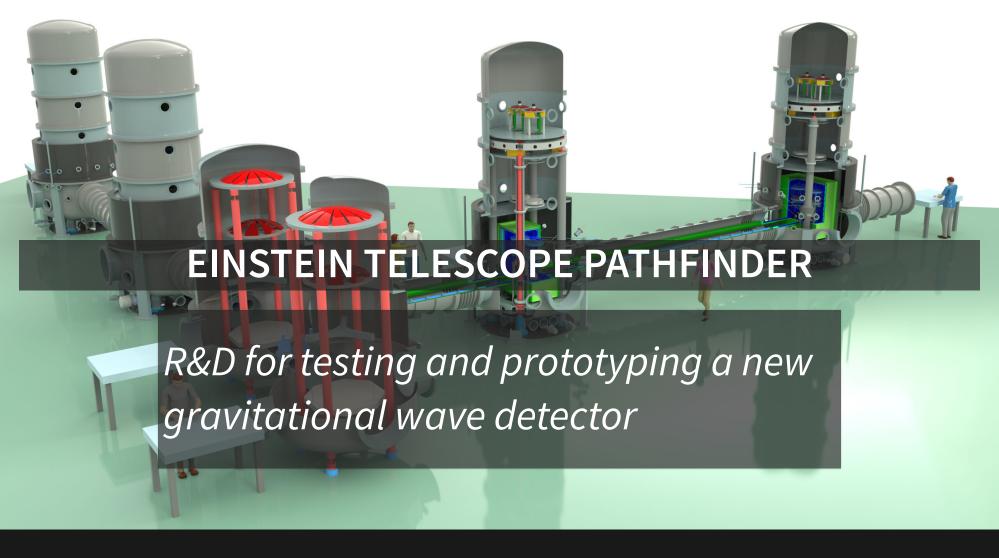
ATLAS, ALICE, LHCb

VIRGO/LIGO

XENON

**KM3NET** 







# COMPUTING @NIKHEF

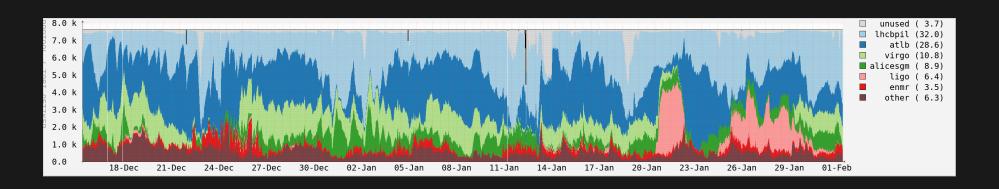
Partnering with SURFsara to form a LHC Tier 1 since 2001.

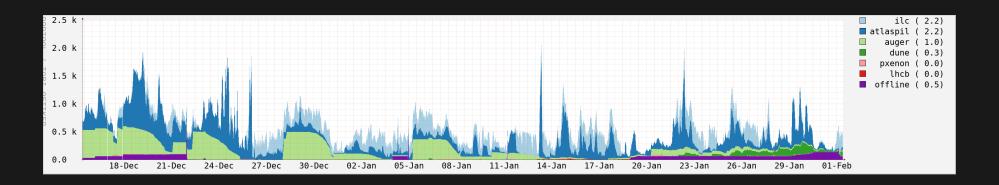
Partnership evolved into a shared national infrastructure for large-scale science computing.

This is well beyond the scope of Nikhef's own mission, but strategically important for a sustainable computing infrastructure.



## RECENT USAGE NUMBERS







# **CVMFS SERVERS**



#### STRATUM 0

We run a stratum 0 server for a long time and it is in need of an upgrade from CentOS 6 to CentOS 7. This is currently on-going.

The only repository of note that we have at the Stratum 0 is softdrive.nl.



#### STRATUM 1

The stratum 1 server has been recently upgraded to CentOS 7. Icinga monitoring is in place, the system has a frontier-squid reverse proxy monitored by CERN.

We carry a bunch of repos from opensciencegrid.org, gridpp.ac.uk and egi.eu.

On purpose we did not include the CERN repos.



The hardware is fully virtual, with a iscsi storage backend on a Fujitsu DX200 flexible tiering storage system. Not the ideal match, but it scales well enough so far.

Current disk use is 6.6TB, coming down from 11TB preupgrade. The difference is probably due to no garbage collection.



# LIST OF REPOS

auger.egi.eu	gm2.opensciencegrid.org	pheno.egi.eu
biomed.egi.eu	hyperk.egi.eu	phys-ibergrid.egi.eu
cdf.opensciencegrid.org	icarus.opensciencegrid.org	pravda.egi.eu
cdms.opensciencegrid.org	icecube.opensciencegrid.org	researchinschools.egi.eu
cernatschool.egi.eu	km3net.egi.eu	sbnd.opensciencegrid.org
chipster.egi.eu	lariat.opensciencegrid.org	scotgrid.gridpp.ac.uk
comet.egi.eu	larsoft.opensciencegrid.org	seaquest.opensciencegrid.org
config-egi.egi.eu	ligo.egi.eu	singularity.opensciencegrid.org
connect.opensciencegrid.org	lkeb.softdrive.nl	snoplus.egi.eu
coupp.opensciencegrid.org	londongrid.gridpp.ac.uk	softdrive.nl
darkside.opensciencegrid.org	lsst.opensciencegrid.org	solidexperiment.egi.eu
des.opensciencegrid.org	lucid.egi.eu	southgrid.gridpp.ac.uk
dirac.egi.eu	mice.egi.eu	supernemo.egi.eu
dune.opensciencegrid.org	minerva.opensciencegrid.org	t2k.egi.eu
extras-fp7.egi.eu	minos.opensciencegrid.org	uboone.opensciencegrid.org
facilities.gridpp.ac.uk	mu2e.opensciencegrid.org	uboone.osgstorage.org

13



# **CVMFS CLIENTS**

- 1. Grid cluster (~7500 cores)
- 2. Local batch system (~800 cores)
- 3. Condor cluster (new and growing)



# EXPERIENCE WITH USING CVMFS AT SCALE

Icinga sensor tracks every repository on every node.

Only occasional (mostly transient) warnings reported.

SERVICE STATUS: 1 I/O errors detected; repository revision 4958

Usually a manual cvmfs\_config reload gets things unwedged.



#### **UPGRADE ISSUE**

On very rare occasions a system won't be able to recover from a CVMFS error and will need a reboot. These have been fewer with later versions of CVMFS.

Just recently the upgrade from 2.7.2 to 2.7.5 caused a handful of our worker nodes to hang; they needed to be drained and rebooted.



#### OSGSTORAGE AND STASHCASH

The ligo.osgstorage.org repository is different as data is not public.

Requires use of x509 helper library and authenticated user.

Actual data coming from stashcache.



# SOFTDRIVE.NL

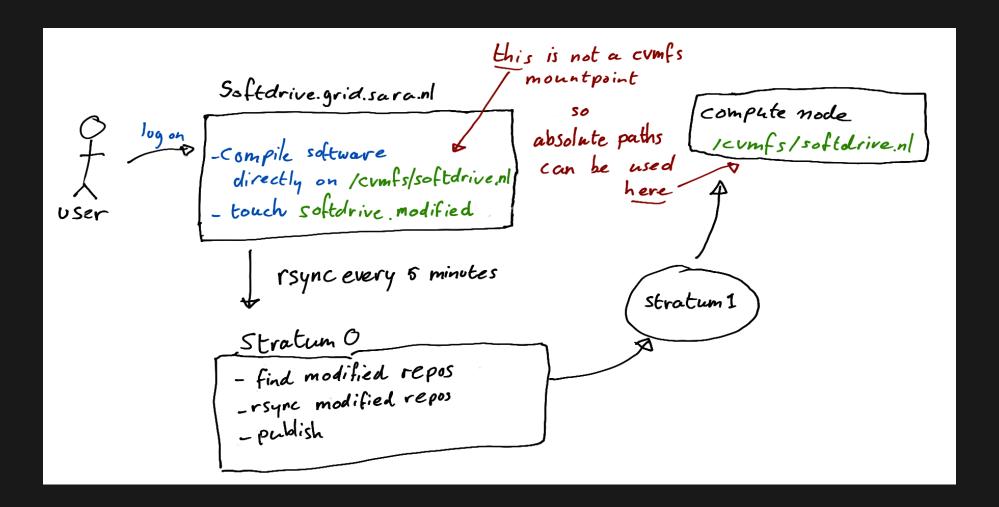
CVMFS 'for the rest of us'; users manage their own directories on this repository.

Presented at the meeting at RAL in 2016.

Co-developed with SURFsara.



#### SOFTDRIVE ARCHITECTURE





#### IMPLEMENTATION DETAILS

Nightly garbage collection (currently broken)

Nested catalogs for bigger users (and the monitoring user which triggered a revision every 5 minutes). This overcame the biggest performance hurdle.

Recently used to include unpacked singularity images (as a proof-of-concept).

# SPIDER @SURFSARA





#### **SPIDER ARCHITECTURE**

The Spider system at SURFsara is a high-throughput data processing platform, similar to Grid but with a broader scope and service offering.

Projects can port their software easily between Grid and Spider using softdrive.nl.



#### **USERS AND USE CASES**

SURFsara advisors have been helping applications to get their software onto different platforms with softdrive.nl.

- minimise effort of (re)installing nad debugging when porting to a new platform
- Centralise management of software used at different sites
- Spreading of singularity images
  - monolithical
  - unpacked



#### Notable projects:

- Tropomi (tropospheric monitoring) uses it to distribute binaries to different platforms.
- LOFAR-SKSP (Solar physics and space weather)
- ProjectMinE (ALS research).



# **FUTURE OUTLOOK**

Upgrading the Stratum-0 immediately.

Explore the possibilities for unpacked container images, template transactions, and the ephemeral writable shell.

Keeping things going with minimal effort.