

WLCG CVMFS Stratum 1 Monitoring

Dave Dykstra, dwd@fnal.gov

WLCG Operations Coordination
23 January 2025



WLCG Stratum 1s map



WLCG Stratum 1s hosted distributions

Site	Distributions
ASGC	CERN, EGI, HSF, OSG
BNL	CERN, DESY (for osg), EGI (for osg), HSF, KEK, OSG
CERN	CERN, DESY, HSF, LSST, RAL
DESY	DESY, KEK
FNAL	CERN, DESY (for osg), EGI (for osg), HSF, KEK, OSG
IHEP	CERN, DESY, EGI, HSF, KEK, OSG
IN2P3	LSST
KEK	KEK
NIKHEF	EGI, OSG, RAL
RAL	CERN, DESY, EGI, HSF, KEK, OSG, RAL
Swinburne	CERN, DESY, EGI, HSF, KEK, OSG
TRIUMF	EGI
UNL	CERN (for replicas), DESY (for osg), EGI (for osg), OSG

- Not all repos replicated to all Stratum 1s
- EGI & OSG have a formal process
- Only one repo from KEK & DESY widely replicated

cvmfs-servermon

- Most of the monitoring is done via wlcg-squid-monitor.cern.ch
- Many monitors, but the most important one is based on the `cvmfs-servermon` package
 - Provides a single API point to probe if all repositories have had a successful snapshot recently, if a repository is about to expire, if it has not been a long time since a successful gc, if a check has failed, and if the geo api works
 - Defaults vary per check and can be updated later
 - Can probe a stratum 1 remotely but works much better if installed locally
 - Then wlcg-squid-monitor.cern.ch probes that API point on all the stratum 1s and sends emails to cvmfs-stratum-alarm@cern.ch if status changes and also sends the status to the CERN XSLS Elasticsearch plot
 - Generates lots of messages (more about that later)

Failover monitor

- The next most important monitor is frontier-awstats
 - Uploads data of where all requests come from to wlcg-squid-monitor
 - Has its own web page where data can be seen in table form, and the data is also used to look for “failovers”, where many connections are coming not through a site squid
 - Used in combination with squid registration
 - Clients at WLCG sites are configured (via the config repos) to fail over to “backup proxies” at CERN and FNAL if their own squids fail, and that’s where most failovers are seen, but Stratum 1s are also watched directly

CVMFS Failover History

This is a visual representation of direct connections from non-squid machines to central groups of servers. The charts are both **viewers** and **controllers**: they allow for interactive exploration of the failover activity on display. The current filtering selections are noted at each chart's header.

Update


Time:

Sep 04, 2022 05:00 PM - Sep 07, 2022 05:00 PM


Time Zone:

Local

Machine Groups



Hits by site per hour



Distance monitor

- The distance monitor is also based on frontier-awstats, although supplemented with additional data beyond the standard awstats tool
 - IP addresses using every repo are collected
 - Plots the average distance in kilometers that clients are away from Stratum 1s per repo
 - The idea is that the plot can help identify when a repository should be replicated to another Stratum 1 because it is getting a lot of use over long distances

Other monitors

- The CVMFS project monitors select repositories in more detail at <https://cvmfs-monitor-frontend.web.cern.ch/>
 - Not very helpful for alarms, just for overall status
 - Based on manually edited metadata
- Stratum 1s that have frontend squids also get MRTG plots of numbers of requests and data bytes sent each 5 minutes
 - Since not much caching is done, the in/out distinction isn't helpful
- Finally, all the repositories on the RAL and UNL stratum 1s are probed daily to see if their root catalog is greater than 25MB
 - If too large, an email is sent to cvmfs-stratum-alarm@cern.ch
 - Fixable only on the stratum 0

Mailing list noise

- A lot of messages go to cvmfs-stratum-alarm@cern.ch
 - Network problems can send a lot of false alarms
 - A stratum 0 repository outage can raise alarms on a lot of stratum 1s outside of the admins' control
 - In the early days we had monitors comparing stratum 0 revisions to stratum 1 revisions but those resulted in a lot more, worse, false alarms
 - Some admins are much more prompt at fixing problems than others
- We don't currently have anyone whose job it is to interpret the mailing list messages and follow up with stratum 1 or stratum 0 admins
 - I do it sometimes, but don't really consider it to be my job
 - Maybe the people who do that for squids could do it also for CVMFS servers, or other people could do it
 - Then we could suggest that Stratum 1 admins unsubscribe
- The `cvmfs-servermon` API on one stratum 1 can also be easily probed by local monitoring systems

Importance of server status?

- It is very difficult to automatically decide when a stratum 1 should be marked as WARNING or CRITICAL
 - How important is that, especially if there are people to interpret it?
 - If one repository is down for a long time, what should the status be? Doesn't it depend on the importance of the repository? A stratum 0 down for a while might not be a problem for the users, until the whitelist expires, which prevents repository mounts.
 - cvmfs-servermon has a config option for a stratum 1 admin to ignore problems on a repository – does that need a central place to say that it should be ignored on all stratum 1s, for a limited time?

Links

- Stratum One admin guide:
 - <https://twiki.cern.ch/twiki/bin/view/CvmFS/StratumOnes>
- cvmfs-servermon:
 - <https://github.com/cvmfs-contrib/cvmfs-servermon>
- Links to WLCG squid monitors:
 - <https://wlcg-squid-monitor.cern.ch>