




CernVM File System



Gathering CVMFS statistics



Daniel-Florin Dosaru • EP-SFT Summer Student

Supervisors: Radu Popescu, Jakob Blomer

About me

Bachelor's student in Computer Science and Engineering
Politehnica University of Bucharest, Romania



dosarudaniel@gmail.com
[Linkedin](#), [Github](#)

Hobbies:

- Open Source Programming, Operating Systems, Embedded programming, Distributed systems, Computer Architecture and Experimental Physics
- Running, swimming, climbing

Gathering CVMFS statistics

This project involved implementing, in the CVMFS server tools, a system for gathering useful statistics about the publication and the garbage collection processes:

For example:


- number of files added/removed/changed/duplicated
- number of bytes added/removed (uncompressed and compressed)
- number of deleted data objects from the backend storage
- publish, ingest, garbage collector run time

Short description

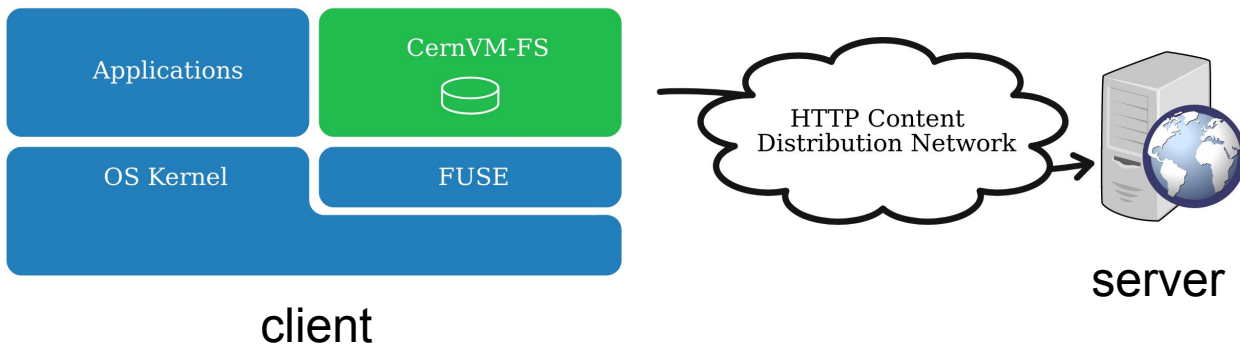


CVMFS = CernVM-File System

- provides a software distribution service for the WLCG
- implemented as a POSIX read-only file system in user space (a **FUSE** module)
- uses outgoing HTTP connections only to avoid firewall issues



WLCG
Worldwide LHC Computing Grid

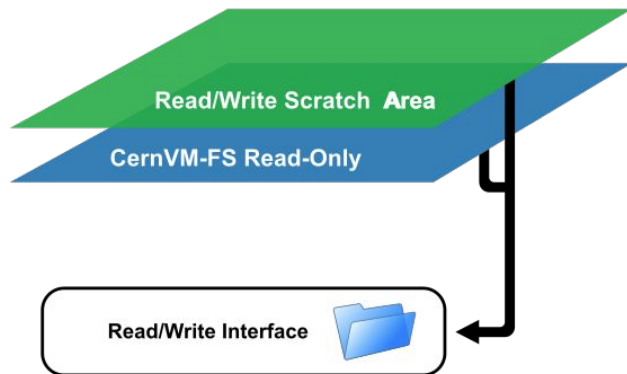


CVMFS publish

CVMFS is not a general purpose file system, but follows a publish/subscribe model.

```
[root@phsftpc2r20:~]# cvmfs_server transaction test.cern.ch  
[root@phsftpc2r20:~]# # add changes to /cvmfs/test.cern.ch  
[root@phsftpc2r20:~]# cvmfs_server publish test.cern.ch
```

Clients have a read-only view of the repository, while changes are published to the repository from a single machine, which we call the release manager.



CVMFS garbage collector

- Internally, CVMFS uses content-addressable **storage** and Merkle trees in order to maintain file data and metadata.
- The garbage collector is figuring out which data objects (compressed file chunks) can be deleted as outdated garbage:

Backend Storage

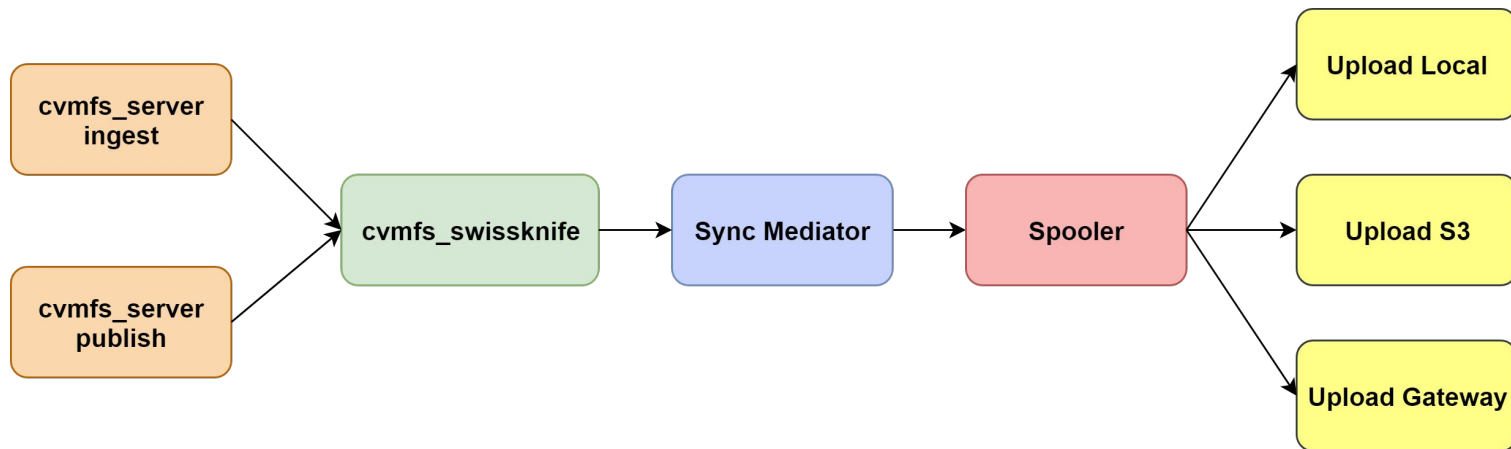
- Local
-  S3
- Gateway

```
[root@phsftpc2r20:~]# cvmfs_server transaction test.cern.ch  
[root@phsftpc2r20:~]# # remove files from /cvmfs/test.cern.ch  
[root@phsftpc2r20:~]# cvmfs_server publish test.cern.ch
```

```
[root@phsftpc2r20:~]# cvmfs_server gc test.cern.ch
```

How to count publish statistics

I modified some of the publish pipeline components:



Mediator: number of added/removed/changed files, uncompressed data size etc.

Upload components: compressed data size, duplicated objects added etc.

Where I store the statistics?

At the end of each publish command the statistics are automatically stored in a **SQLite local file**.

Optionally, the stats can be printed at the console.



Script for submitting publisher statistics to a carbon server:

<https://github.com/cvmfs/cvmfs/pull/2222>

Potential uses

- **Monitoring CVMFS**
 - repository storage
- **Alarm system**
 - discover unusual behaviour
- **Benchmarking**
 - measure server tools execution time



Results



- We installed the tools on the LHCb nightly test release manager
- Gathering stats since 15th of August:

August 2018						
Sun	Mon	Tue	Wed	Thu	Fri	Sat
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

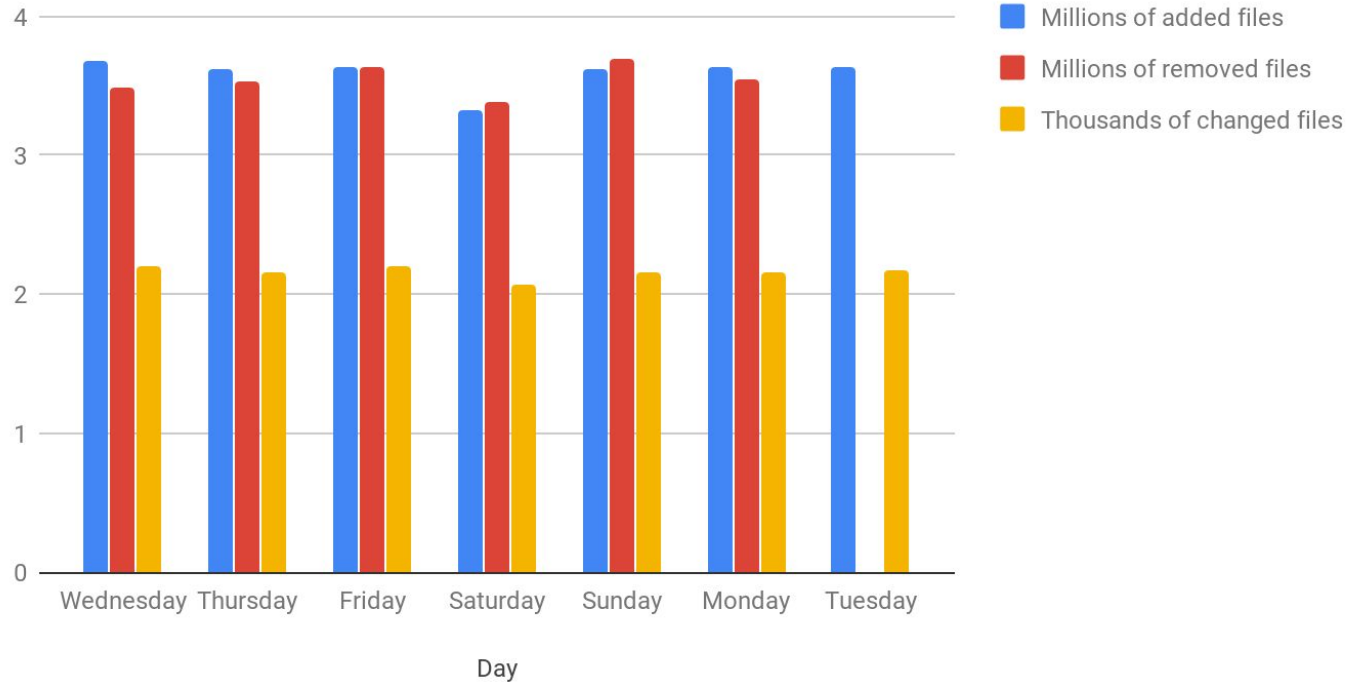
Added

Total added files	25 M
Uncompressed data added	1 549 GB
Compressed data added	501 GB

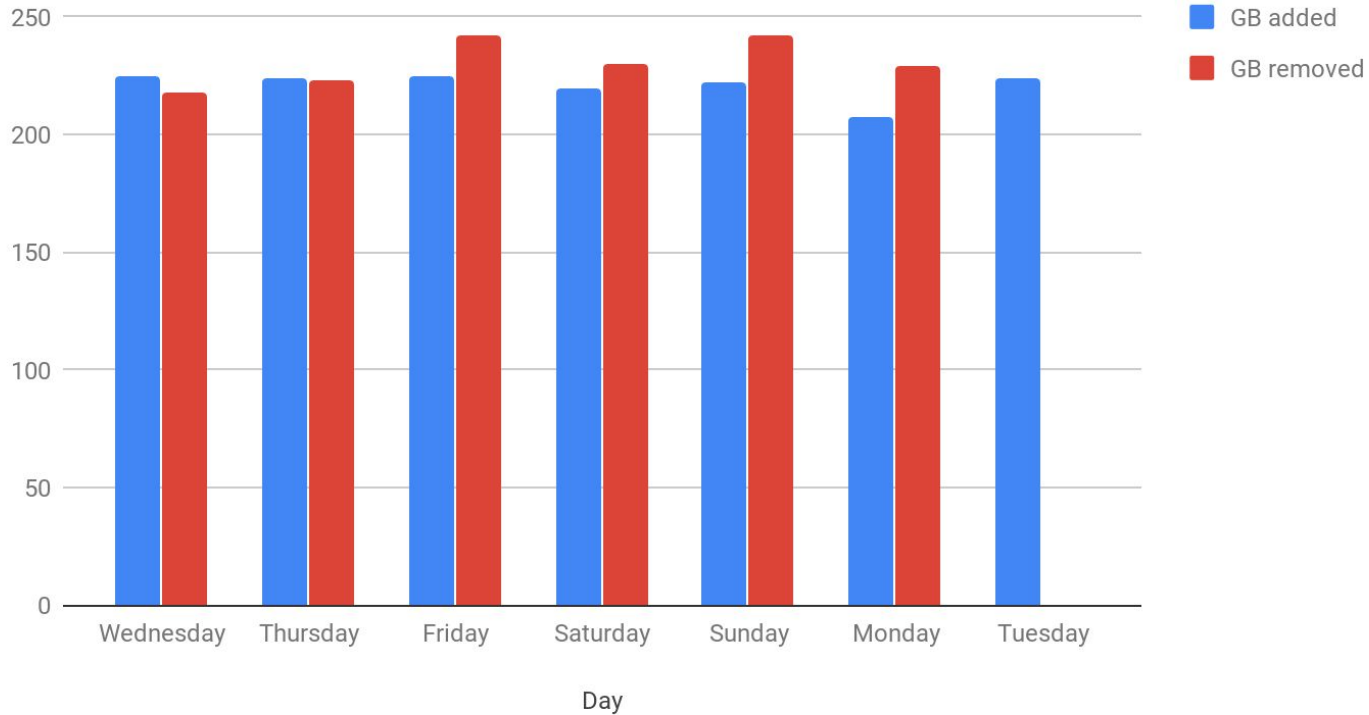
Removed

Total removed files	21 M
Uncompressed data removed	1 386 GB

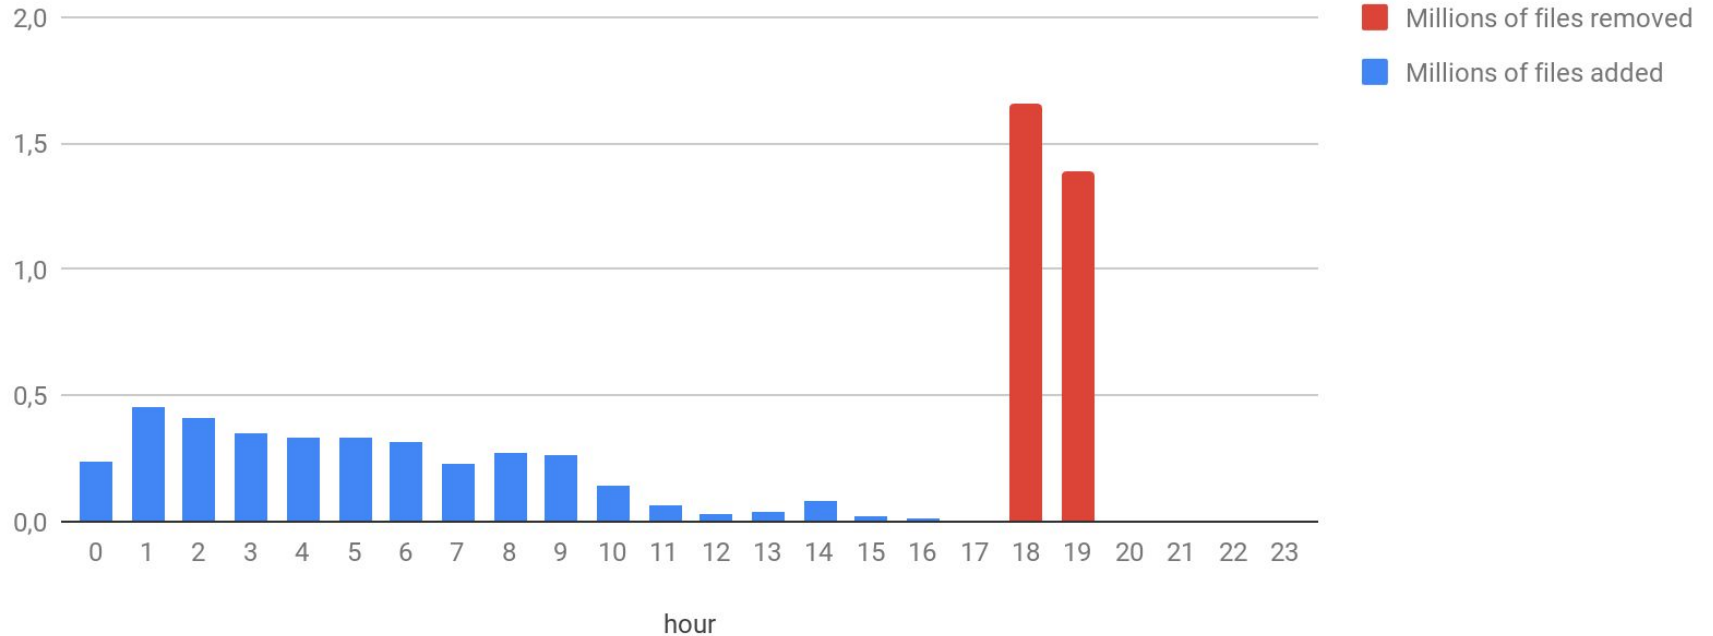
Publish activity over the week




Publish activity over the week - uncompressed data size



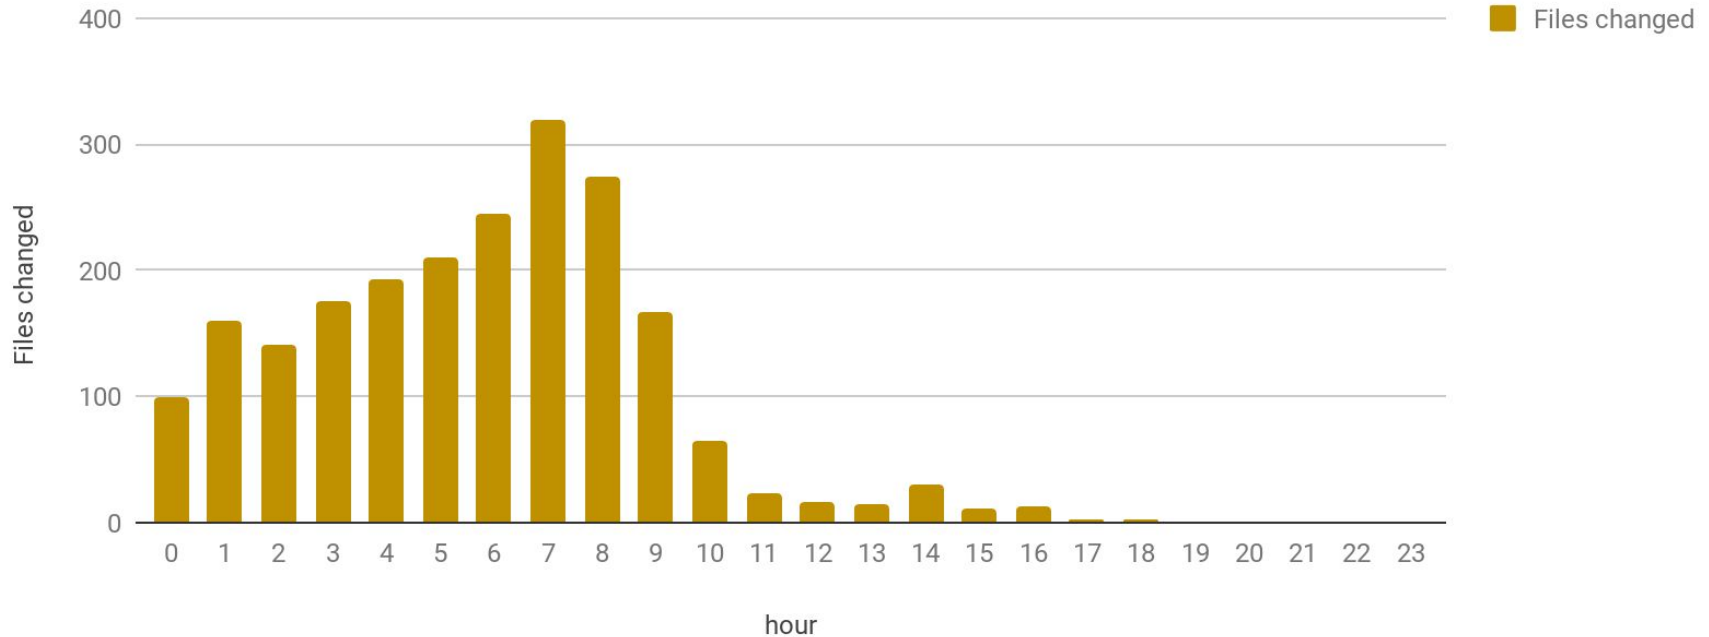
Average publish activity over the week per hours (files added/removed)



Average publish activity over the week per hours (data size)



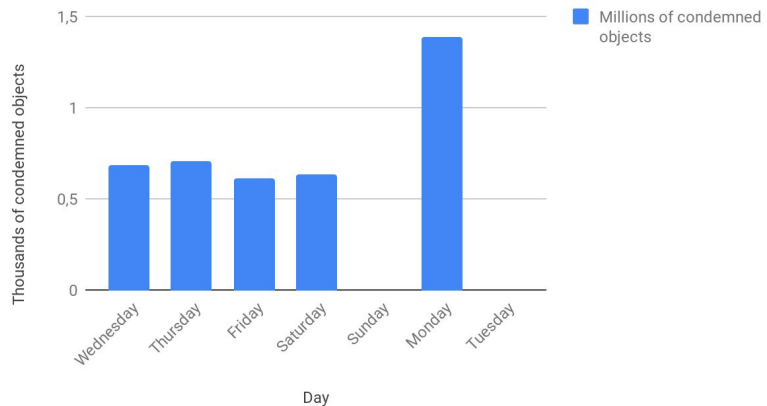
Average publish activity over the week per hours (files changed)



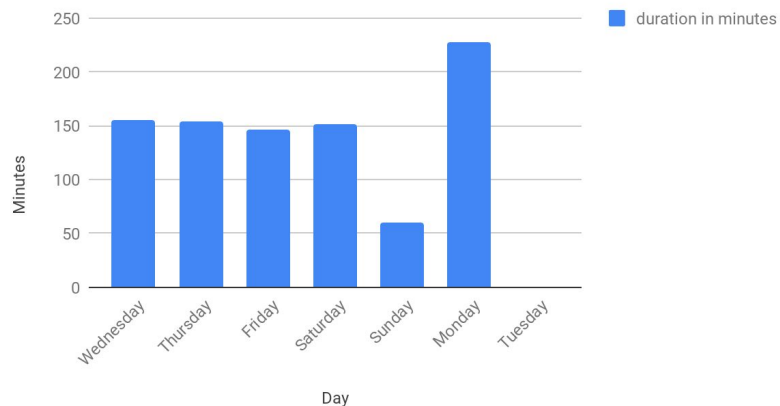
Garbage collector activity:

Day	Start time	Duration	Preserved catalogs	Condemned catalogs	Condemned objects
Wednesday	19:05	2:35	6,666	2,534	681,789
Thursday	19:05	2:34	6,753	2,087	710,525
Friday	19:05	2:26	7,129	1,836	616,224
Saturday	19:06	2:32	6,606	2,606	636,603
Sunday	19:05	1:00	5,345	0	0
Monday	19:07	3:48	6,572	4,386	1,388,537
Tuesday	-	-	-	-	-
Wednesday	19:21	3:39	6,429	4,261	1,397,750
Thursday	19:05	2:27	6,667	2,170	603,889
Friday	19:05	1:46	6,839	2,166	59,509
Saturday	19:04	2:28	7,100	1,926	579,766

Garbage collector activity



Garbage collector run time



More numbers...

Average number of transactions per day	1,070
Average number of added files per transaction	3,286
Maximum number of files in one transaction	51,306 (746 MB)
Maximum data size added in one transaction	~2 GB (for 2,133 files added)
Maximum runtime for a publish command	4 minutes

[All statistics data](#)

Future

- It will be available in the **CernVM-FS 2.6.0** release!
- More metrics can be added.
- TODO: Create an analysis tool for interpreting the release manager statistics



[My contribution on github](#)

[My Summer Student Report](#)

Thank you!

dosarudaniel@gmail.com

daniel-florin.dosaru@cern.ch