



EOSHome Backup Status

Roberto Valverde (CERN, IT Storage Services)

3rd Feb 2020 – 4th EOS Workshop

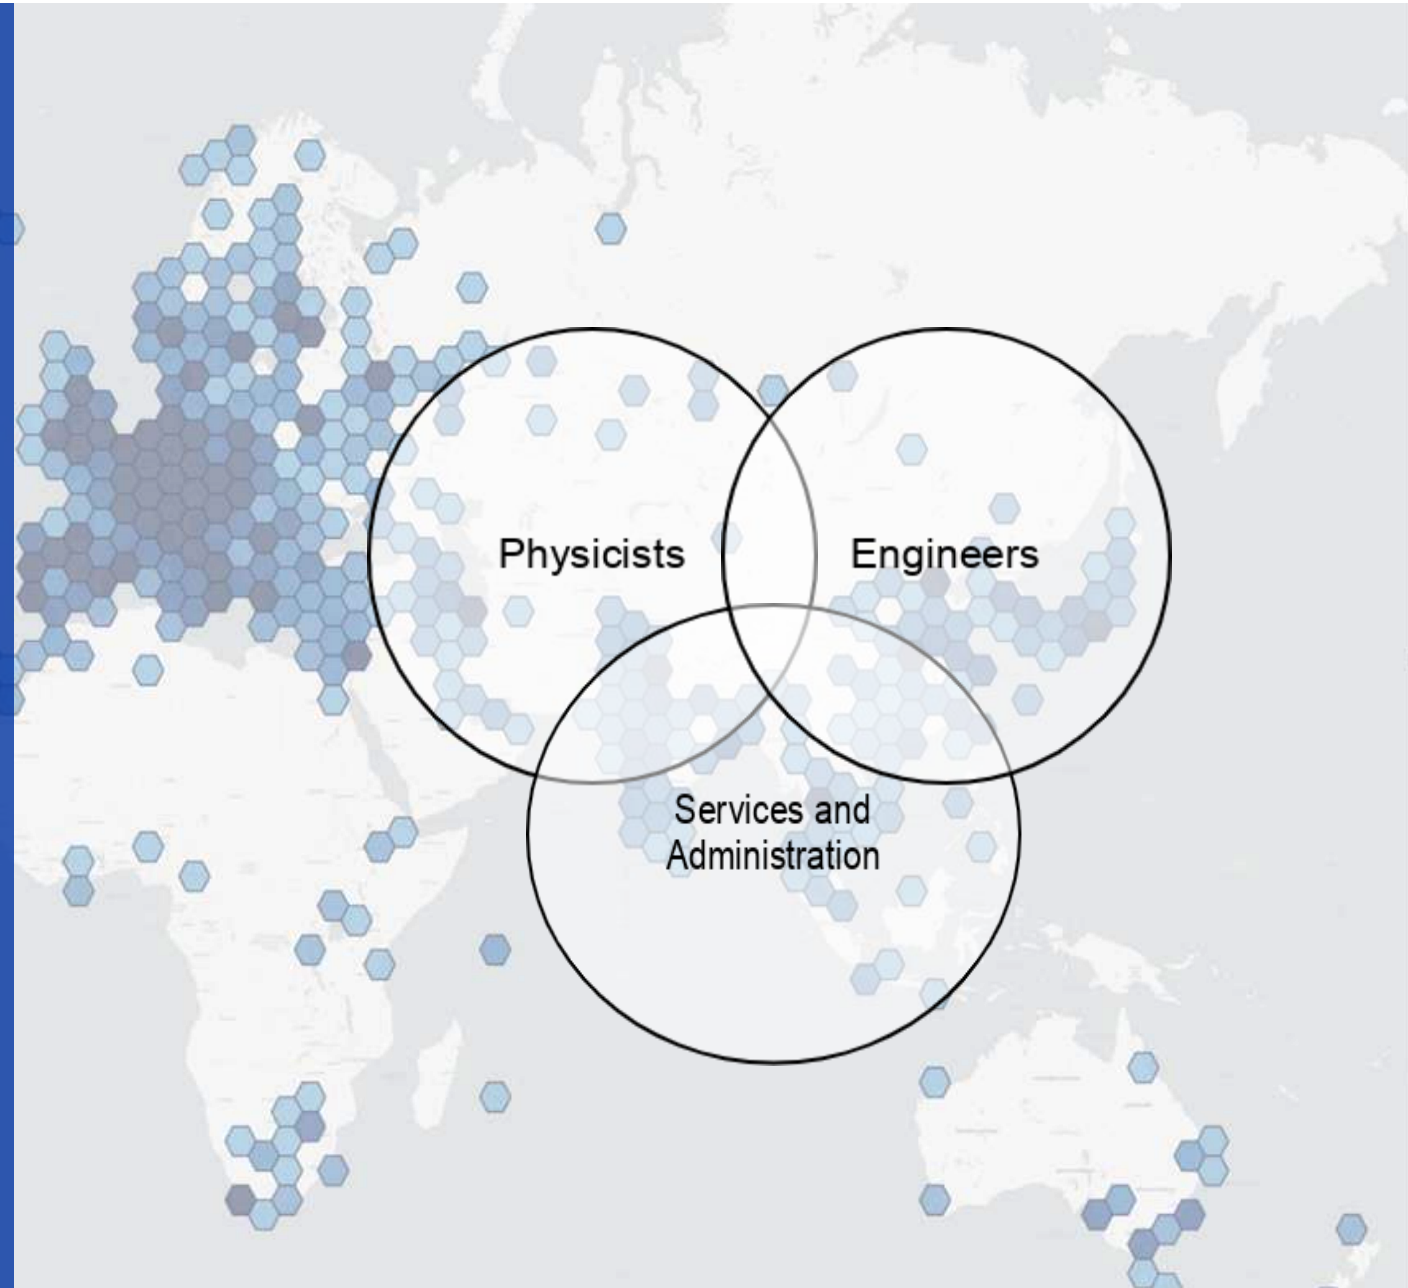


CERNBox



EOS

- 21,000 accounts
- 7PB of storage
- 1.1 Billion files
- 130,000 shares
- 5,000 active users/day



restic - Backups done right!



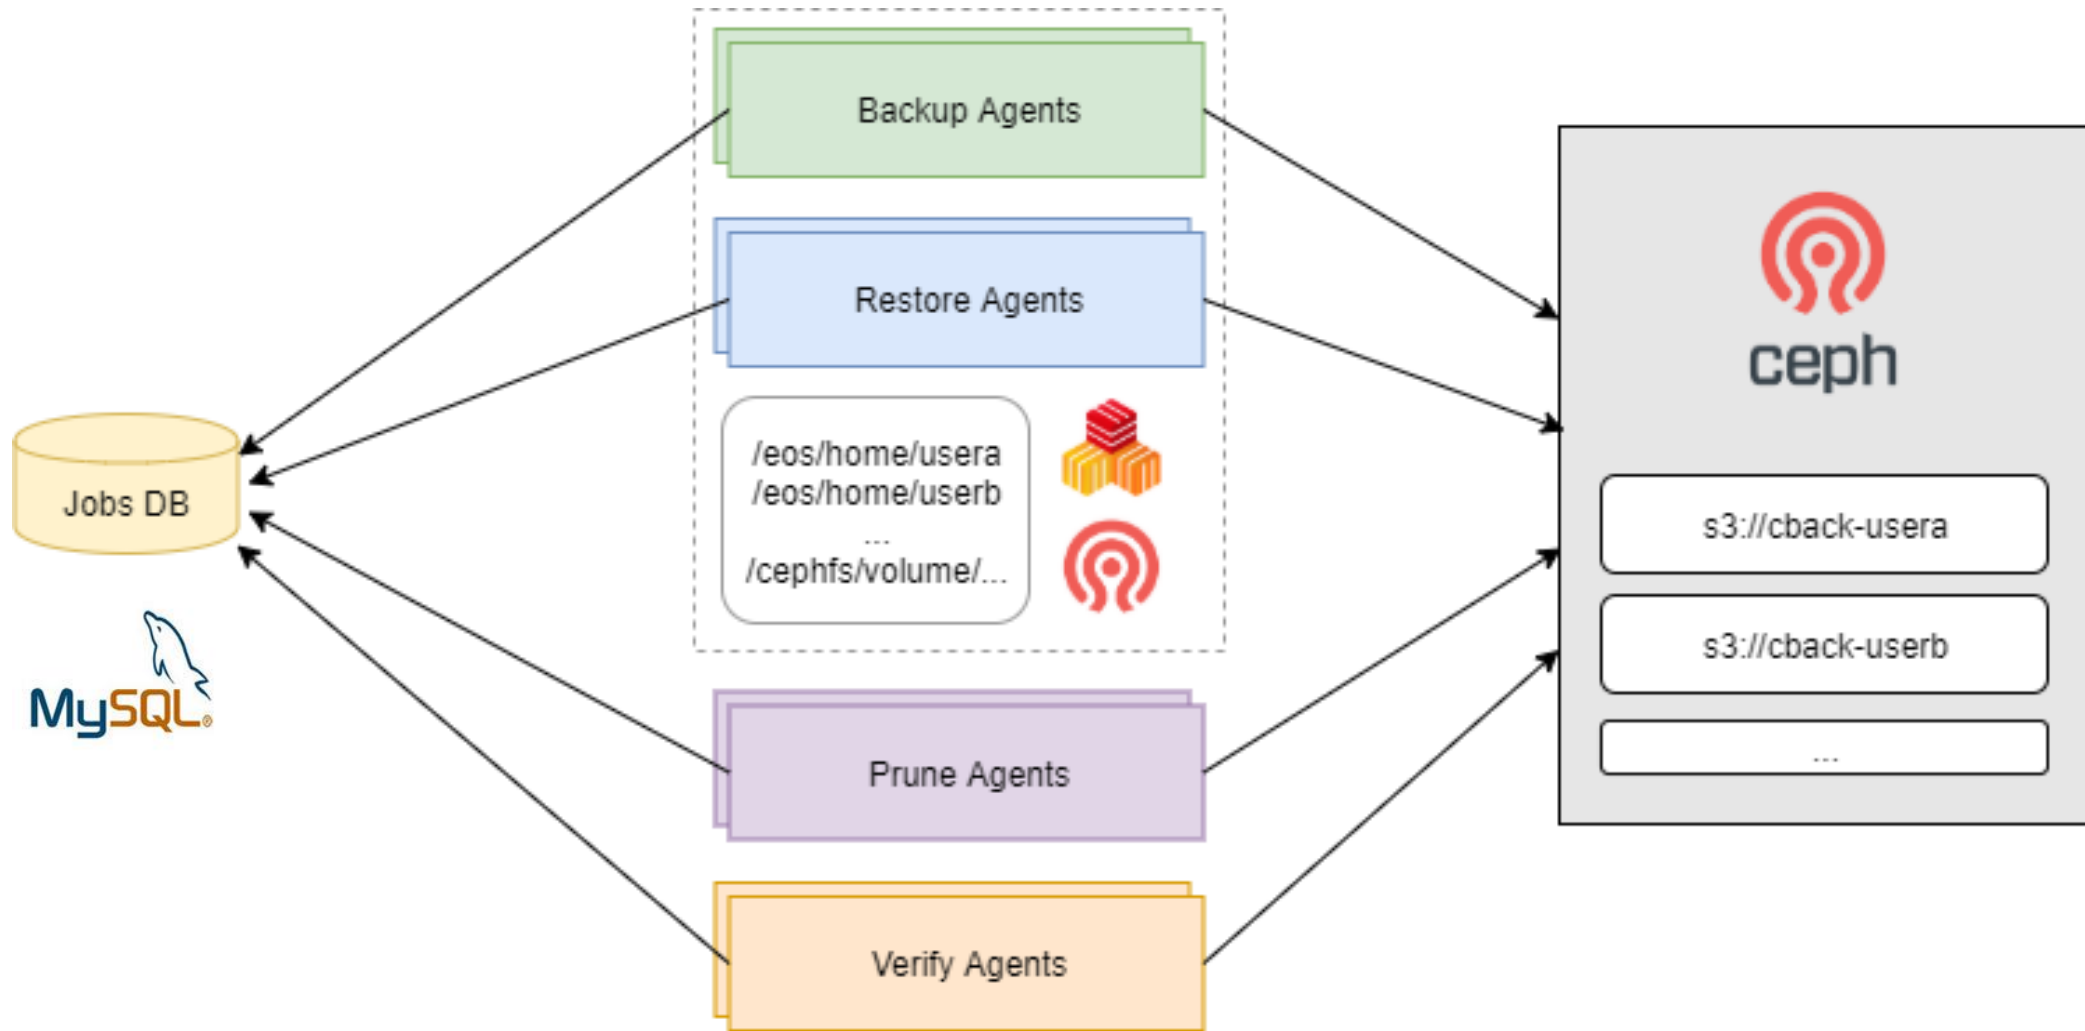
- Easy, fast and secure
- Repository based, incremental snapshots
- Efficient: File & Chunk de-duplication (**CDC**)
- Multiple backends (local, sftp, S3, Swift, Backblaze B2, Azure, Google Cloud, rclone...)
- <https://github.com/restic/restic> ([web](#) | [blog](#) | [forum](#))

```
restic -r <repo> init # Initialize repository
restic -r <repo> backup <path> # Backup a folder
restic -r <repo> snapshots # List snapshots
restic -r <repo> restore <snap_id> -t <target> # Restore
restic -r <repo> mount /mnt/backup # Mount backup repository via fuse
```

Main features

- **restic orchestration tool**
- **S3 backend for backup storage**
- **eos-fusex, cephfs support as source filesystem**
- **Incremental backup, only when is needed**
- **Individual backups, scattered in time**
- **Scalable backup and restore agents**
- **Automatic backup verification**
- **User-triggered restore jobs**
- **Admin CLI**

Architecture



Command line interface



--docopt
tabulate

CERNBack - Backup Status Summary

Total	Pending	Running	Completed	Failed	Coverage
21562	2623	20	18919	0	87.41 %

Not enabled: 2622

Running Backups

Id	User	Path	Status	Worker
20450		/eos/prc	Running	cbox-restic-backup-12.cern.ch:11469
16671		/eos/hon	Running	cbox-restic-backup-08.cern.ch:5075
15192		/eos/prc	Running	cbox-restic-backup-11.cern.ch:22575
16214		/eos/hon	Running	cbox-restic-backup-01.cern.ch:15017
16216		/eos/hon	Running	cbox-restic-backup-07.cern.ch:8436
13777		/eos/hon	Running	cbox-restic-backup-06.cern.ch:3127
16019		/eos/hon	Running	cbox-restic-backup-02.cern.ch:30713
17798		/eos/hon	Running	cbox-restic-backup-10.cern.ch:14824
17353		/eos/hon	Running	cbox-restic-backup-08.cern.ch:24278
13337		/eos/hon	Running	cbox-restic-backup-12.cern.ch:32193
15852		/eos/hon	Running	cbox-restic-backup-11.cern.ch:18410
17026		/eos/hon	Running	cbox-restic-backup-06.cern.ch:5871
16712		/eos/hon	Running	cbox-restic-backup-09.cern.ch:7844
15825		/eos/hon	Running	cbox-restic-backup-03.cern.ch:28414
1718		/eos/prc	Running	cbox-restic-backup-01.cern.ch:11526
17787		/eos/hon	Running	cbox-restic-backup-03.cern.ch:11327
16884		/eos/hon	Running	cbox-restic-backup-05.cern.ch:13475
16533		/eos/hon	Running	cbox-restic-backup-05.cern.ch:1808
17067		/eos/hon	Running	cbox-restic-backup-02.cern.ch:15416
6747		/eos/hon	Running	cbox-restic-backup-04.cern.ch:13157


CERNBack CLI.

Usage:

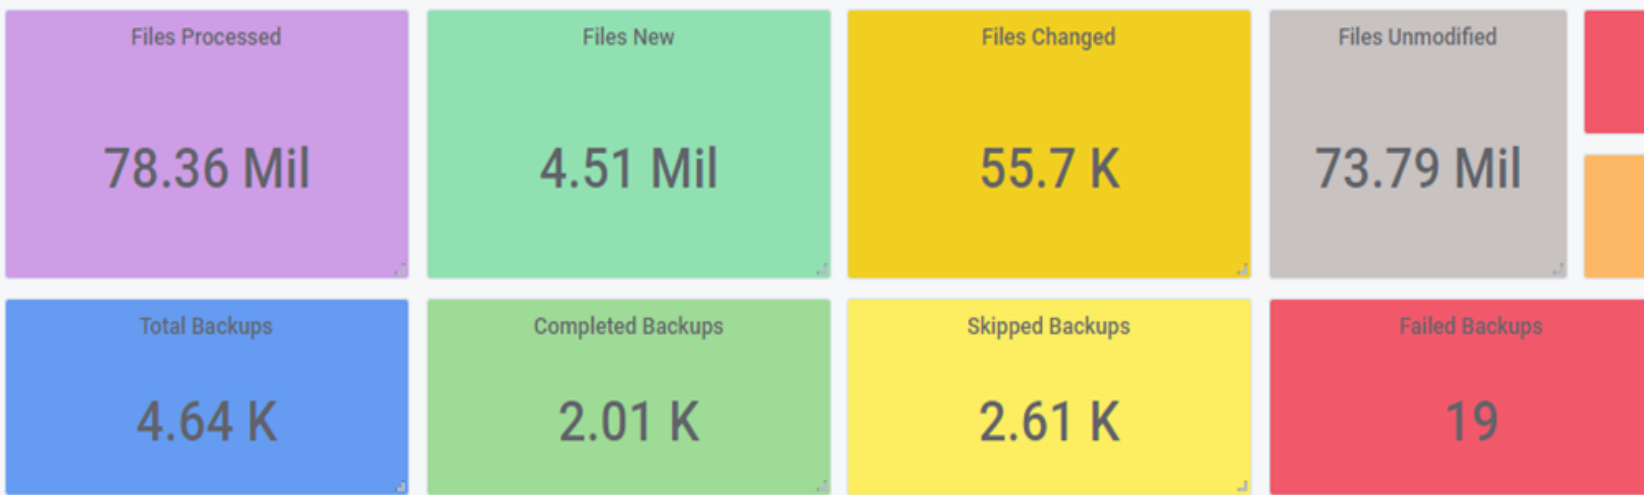
```
cback backup status [<user_name>]
cback backup ls [failed|running|pending|completed|disabled] [--format <form
cback backup add <user_name> <instance> <path> [--bucket-prefix=<prefix>][
cback backup (enable|disable|reset|(delete|rm) [--force]) <backup_id>
cback prune status [<user_name>]
cback prune (enable|disable|reset) <backup_id>
cback prune ls [failed|running|pending|completed|disabled]
cback restore status [<id>]
cback restore ls [failed|running|pending|completed]
cback restore add <backup_id> [--pattern=<pattern>][[--snapshot=<snapshot_i
cback restore (enable|disable|reset|(delete|rm) [--force]) <restore_id>
cback stats [<id>]
cback (-h | --help)
cback --version
```

Options:

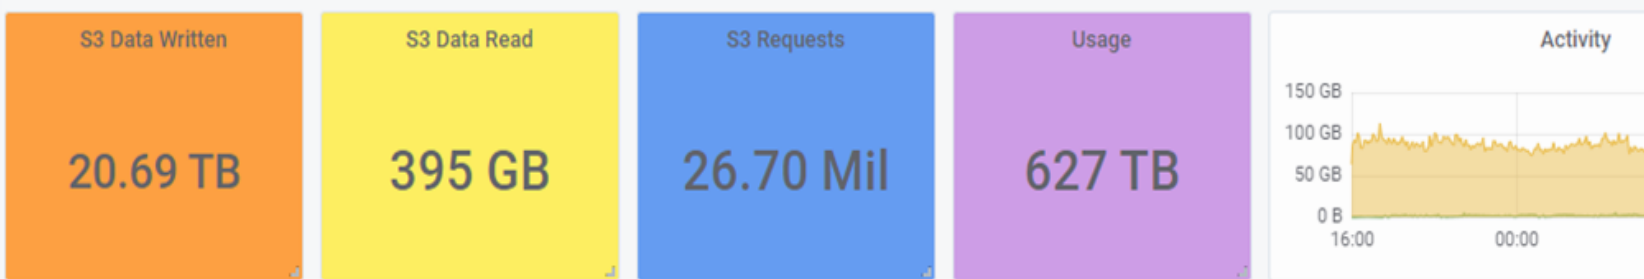
```
-h --help      Show this screen.
--version      Show version.
--enable       Enable the backup after adding it.
--project      Uses project prefix for creating the bucket.
--format       plain|ipal|html|latex|github
```



Backup Activity



S3 Statistics




- ~19k users backed-up daily (88% of total)
 - ~630TB of S3 storage
 - ~80M files processed per day
 - ~20TB traffic / day
- 24h, 12 x 8c, 16G VMs

Next

- Finish to add users and project spaces
- Integration of the restore functionality into CERNBox infrastructure
- Multi-backend/Multi-replica backups
- Using restic for backup on tape?

Thank you!

rvalverd@cern.ch



home.cern