



Backup Infrastructure at CERN

Aurélien GOUNON
IT-ST-TAB




Agenda

- Service overview
- Infrastructure
- Recent changes
- Impact and benefits
- Perspective and challenges

Service Overview

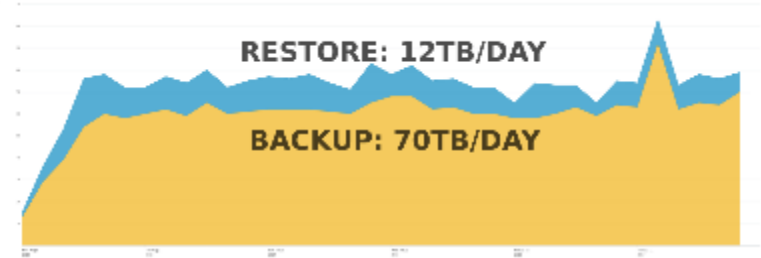
- What we backup
 - 11PB of data
 - 52TB of exchange mailboxes
 - 80 databases
 - > 1100 servers
 - 2'300 million files
- We don't backup physics data and PC

Service Overview




Daily traffic

Total volume
breakdown



Infrastructure

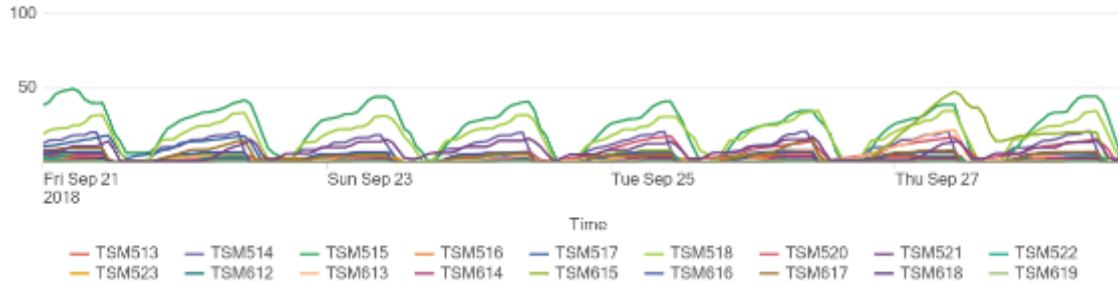
- 2 IBM TS3500 libraries
 2 buildings
- 55 IBM TS1140 drives
- 13 000 IBM 3592 tapes
- 7 Fiber channel switches
- 19 TSM servers
- 2 library managers



Infrastructure

A two step process

Backup pool fill percentage



Recent changes

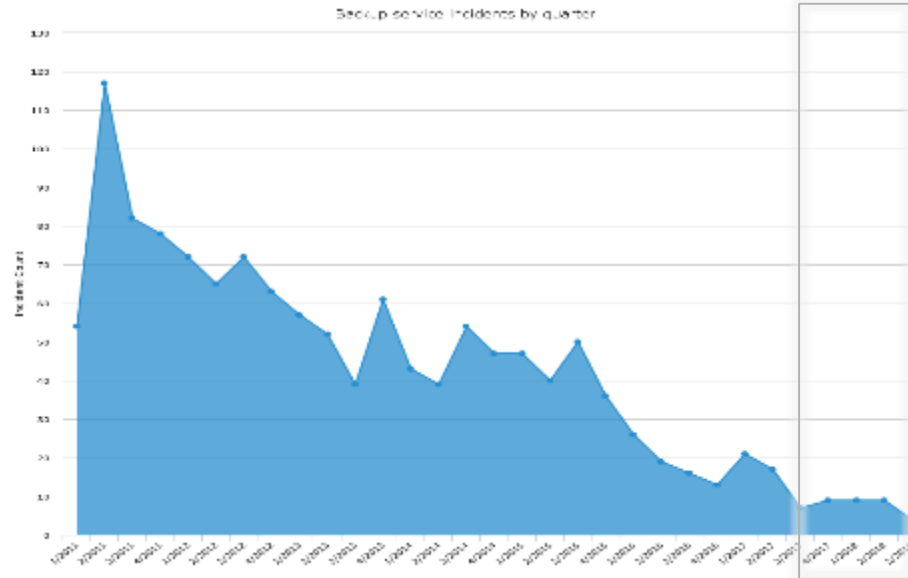
- Using simple elements that scale: software raid and SAS expanders
 - 16 to 32TB of local storage
 - 9 SW RAID 1E volumes
 - Up to 200 TSM FILE volumes
 - 890MB/s backup speed
 - 340MB/s migration speed with 2 processes



Impact and benefits

- A mature and reliable service

Number of tickets opened to the Backup Service



Past 12 months

Perspective and challenges

- Automation
 - Use of Rundeck for automation of recurrent tasks and maintenance
- License limitations
 - Currently licensed for up to 15.9 PB, license renewal in 2020
 - Considering alternatives to TSM for big users

Perspective and challenges

- Future
 - Backup to EOS CTA with natively supported and scaling transfer clients
 - Provide CephFS filers backup

Thank you, questions?





www.cern.ch