

Julien Leduc, Aurelien Gounon - CERN IT

CERN's current Backup and Archive Service hosts 11 PB of data in more than 2.1 billion files. We have over 500 clients which back up or restore an average of 80 TB of data each day. At the current growth rate, we expect to have about 13 PB by the end of 2018. In this contribution we present CERN's Backup and Archive Service based on IBM Spectrum Protect (previously known as Tivoli Storage Manager). We show the architecture and design of the system, the user requirements, the operational issues, the growth challenges and the options for resolving them.

WHAT WE BACKUP

CRITICAL DATA HOSTED IN CERN COMPUTER CENTER

- 52TB OF EXCHANGE MAILBOXES
- 80 DATABASES
- > 1100 SERVERS

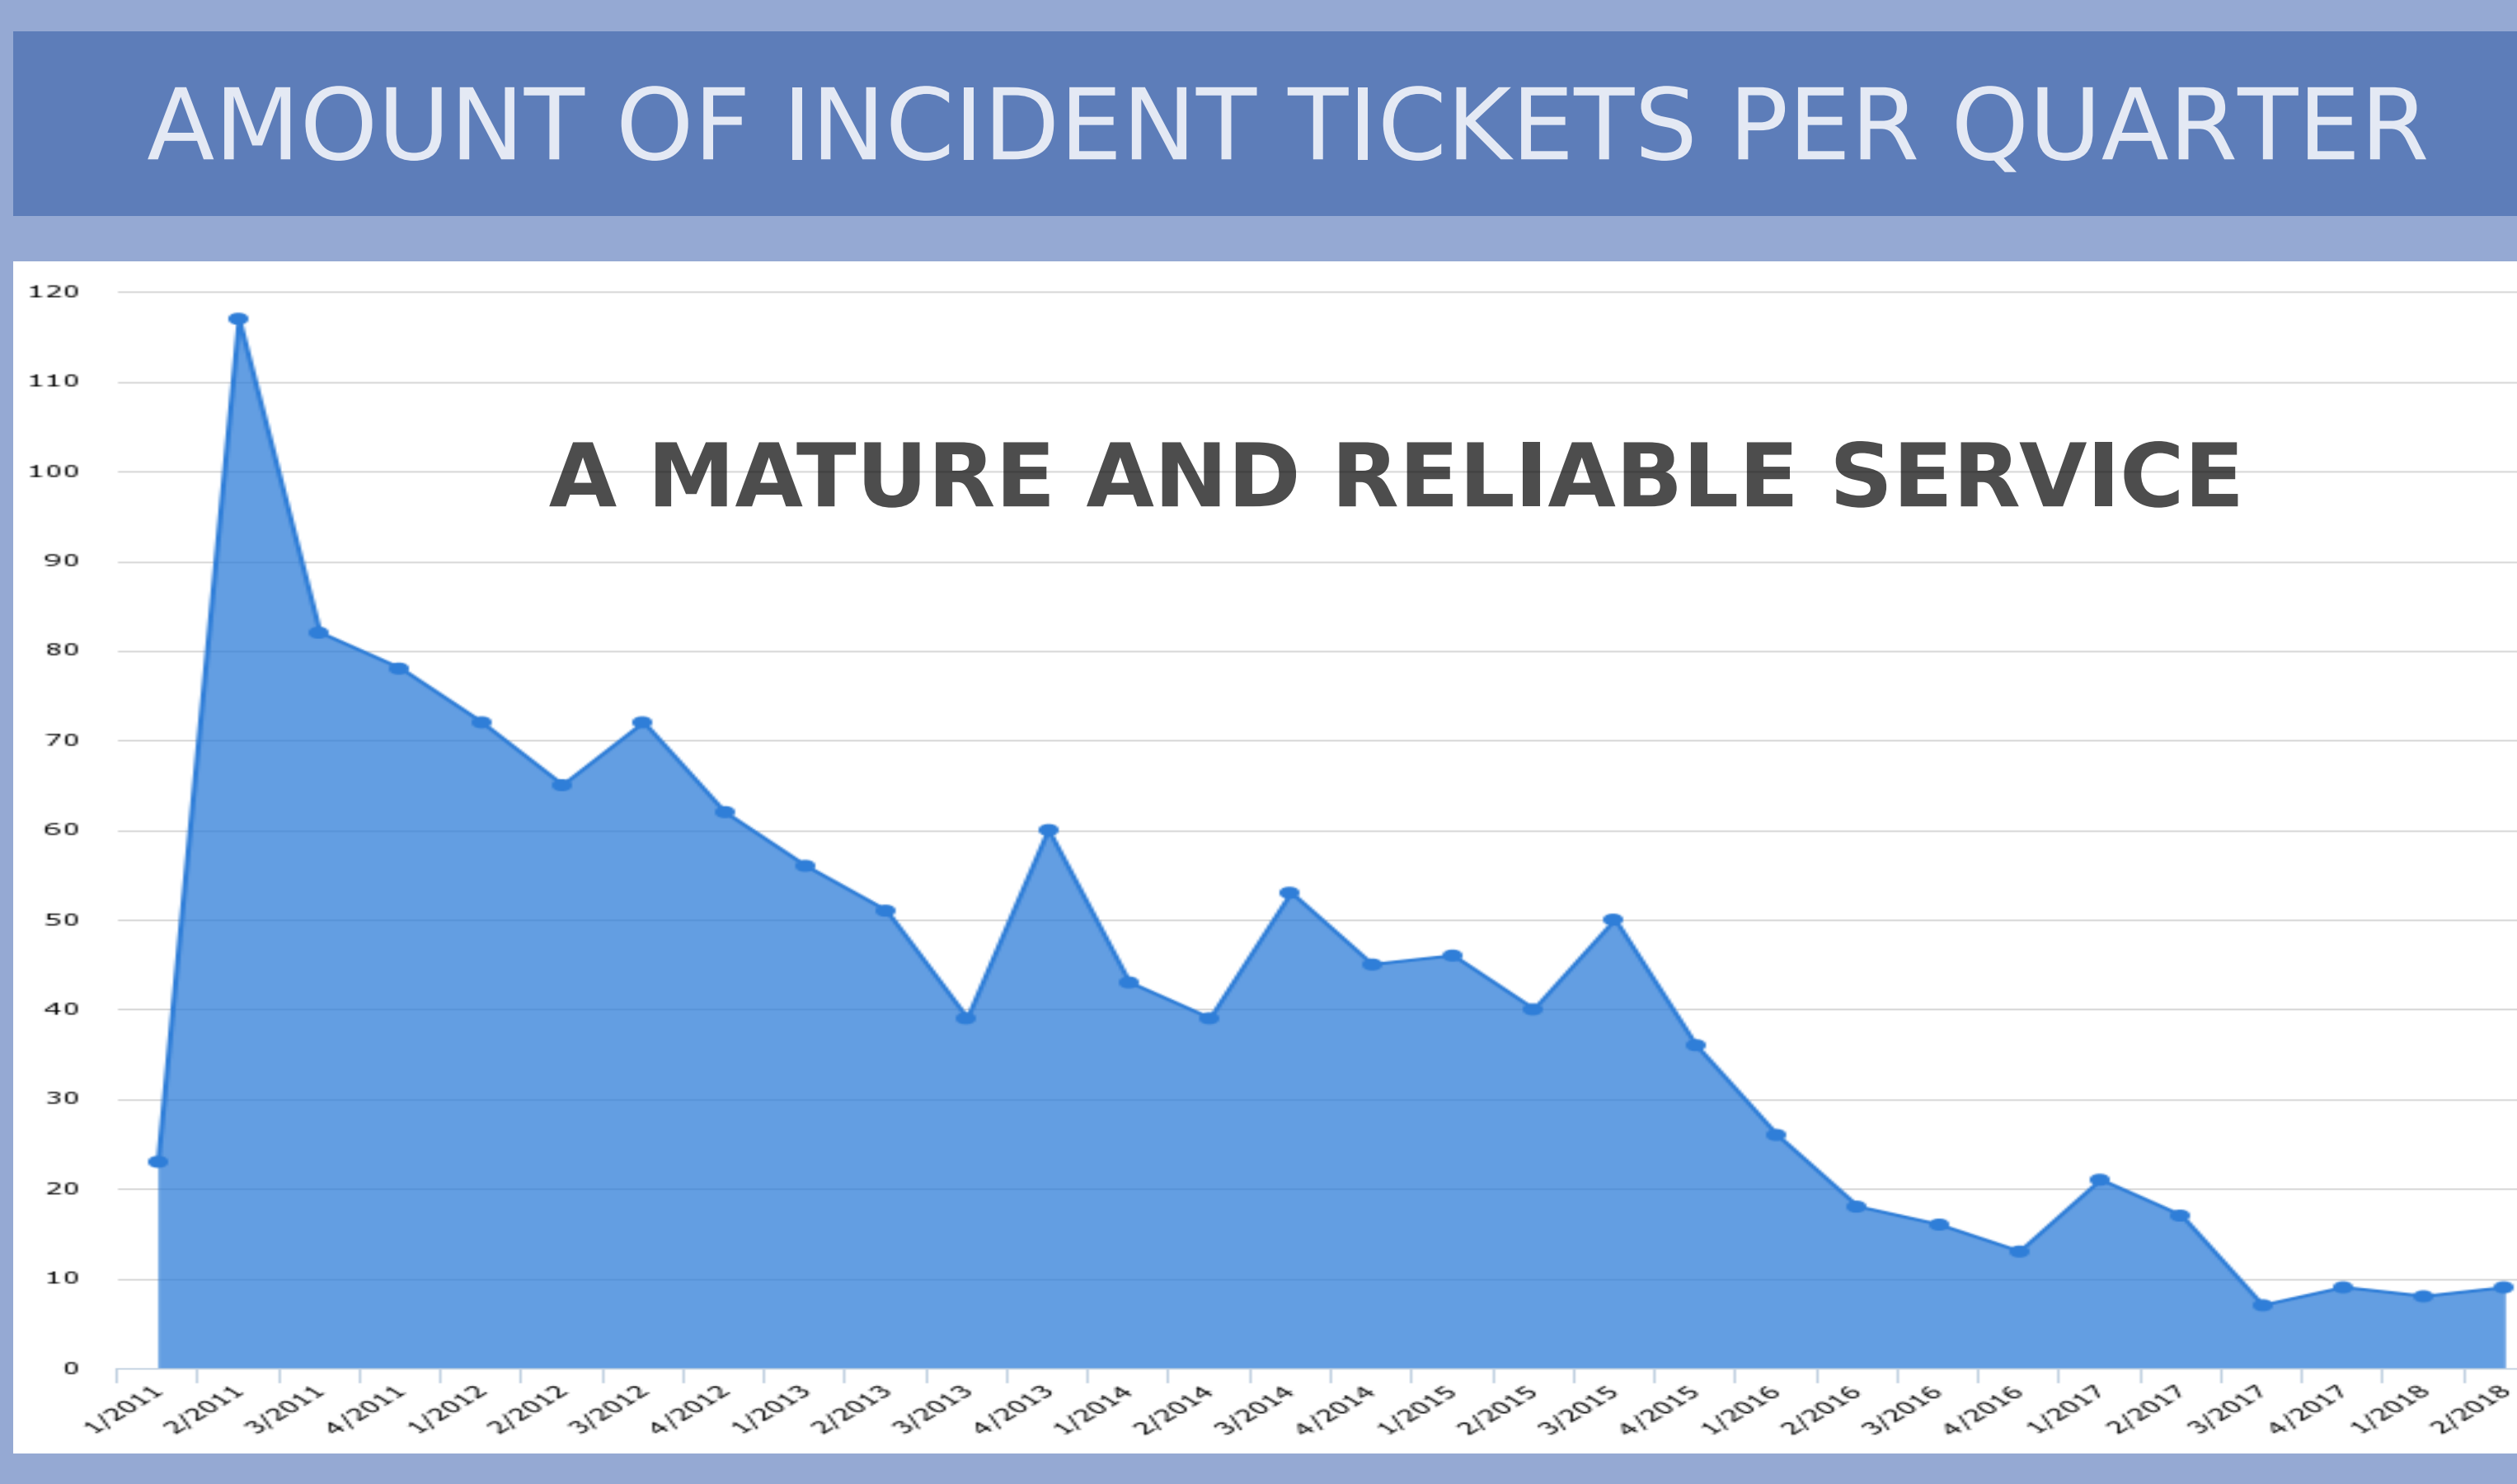
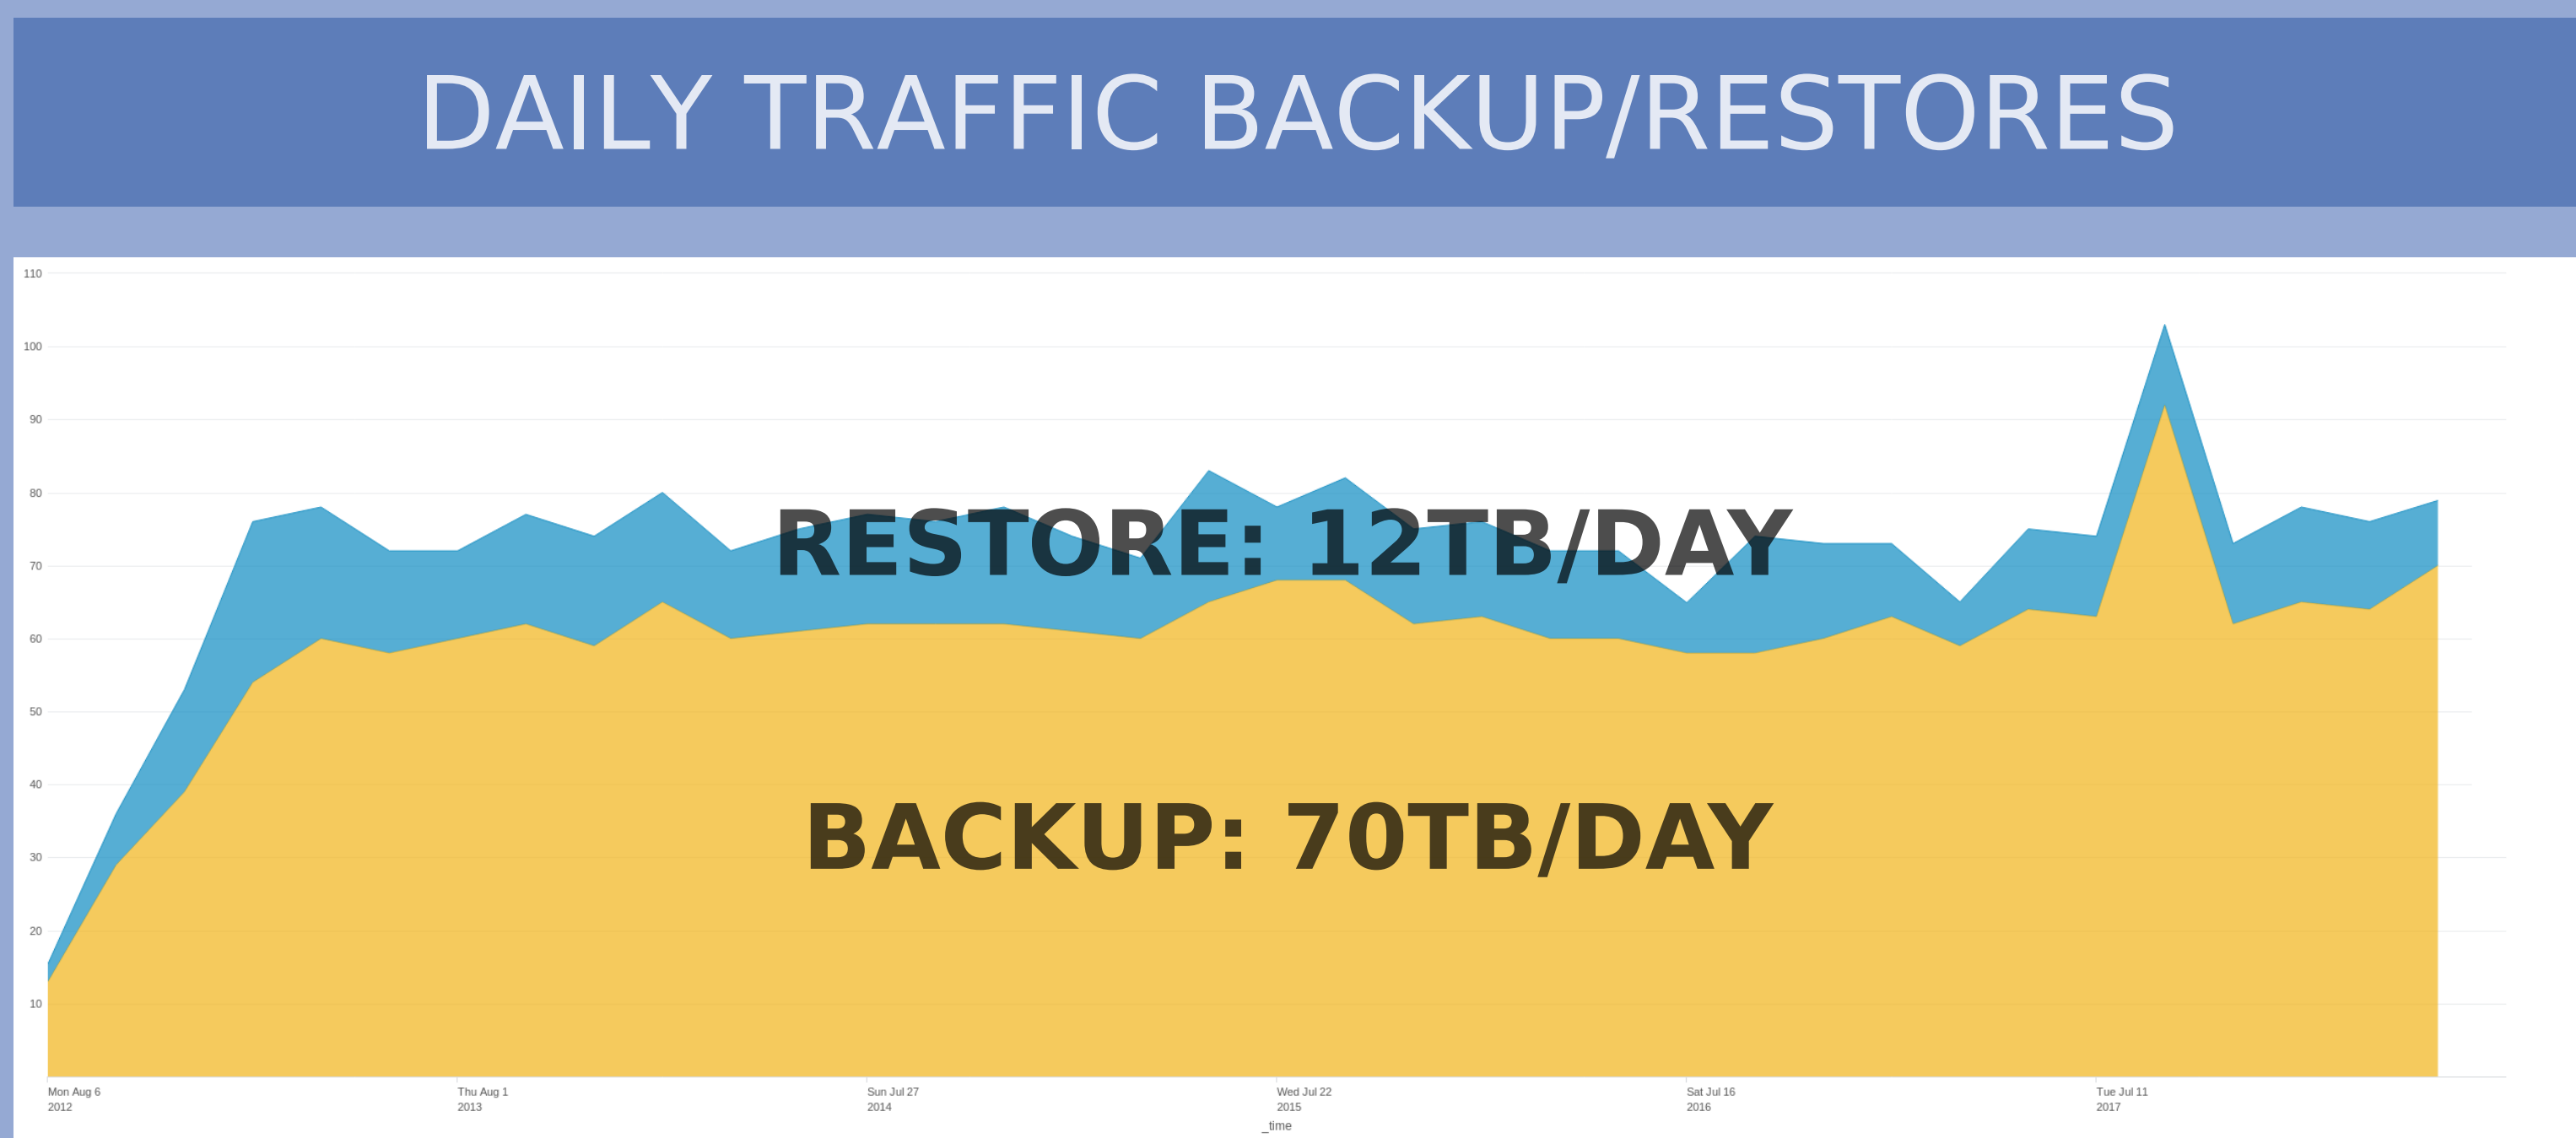
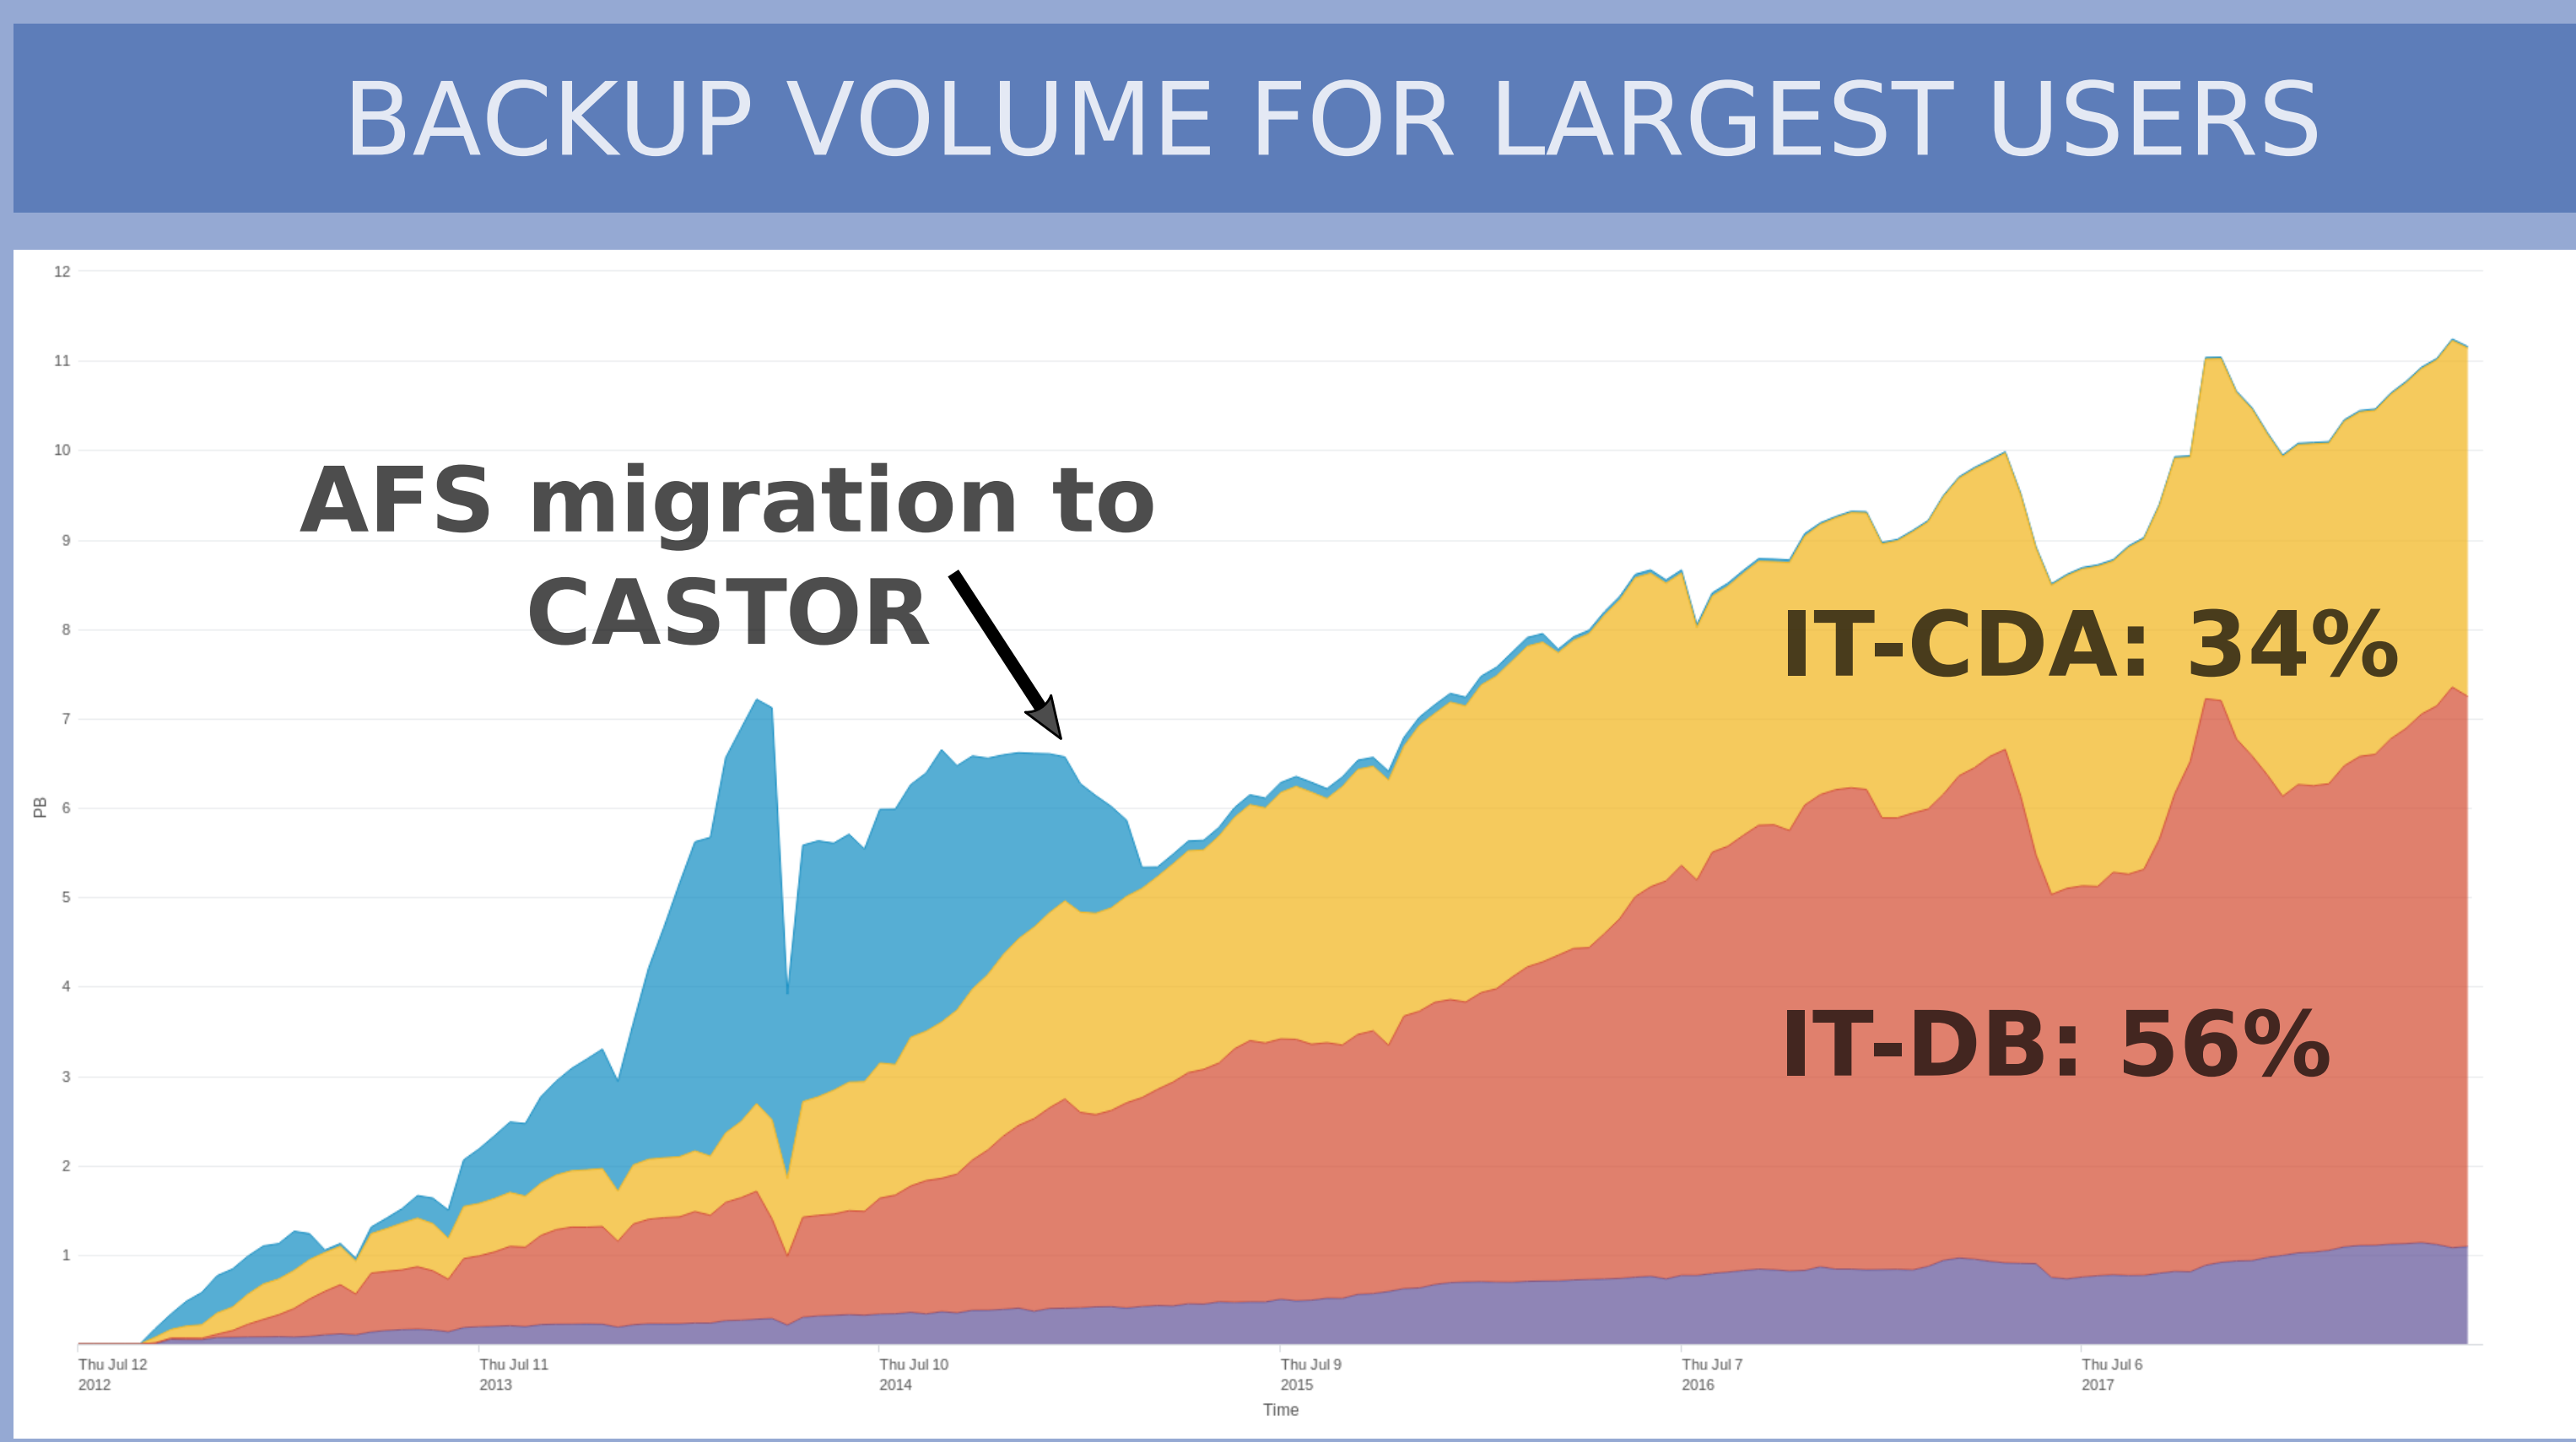
WHAT WE DO NOT BACKUP

- PHYSICS DATA
- DESKTOPS


11
PETABYTES

2'300
MILLION FILES

USAGE AND TRENDS



INFRASTRUCTURE



A SIMPLE COMPUTING ELEMENT THAT SCALES

MONITORING

SERVICE OVERVIEW

ERRORS BY SEVERITY

OUTLOOK

A RELIABLE PERFORMANT SERVICE CAPPED BY LICENSED VOLUME ...STILL GROWING (~20%/YEAR)

ALTERNATIVES?

AFS BACKUPS MOVED TO CASTOR

WHAT ABOUT OTHERS?

DB BACKUP STUDY

Solution	Available	Supported	Cheap	Easy	Fast
RMAN Plugin	✗	✗	?	?	?
Disk	✓	✓	✗	✓	✓
NFS to EOS	✓	✓	✓	✗	~
NFS to Ceph	✓	✓	✓	✗	✓
S3 to EOS	✗	✗	✓	?	?
S3 to Ceph	✓	✗	✓	✓	~
Appliance	✓	✓	✗	✓	?
More TSM	✓	✓	✗	✓	✓