

A person is standing in the center of a long, brightly lit server room aisle, looking down at a tablet. The aisle is flanked by tall server racks on both sides, which are filled with various components and glowing lights. The perspective is looking down the length of the aisle, creating a sense of depth. The lighting is a mix of cool blues and bright whites, with some warm yellow lights visible on the right side of the racks.

# A five year journey with ESS

**IBM Storage Scale Days 2024**

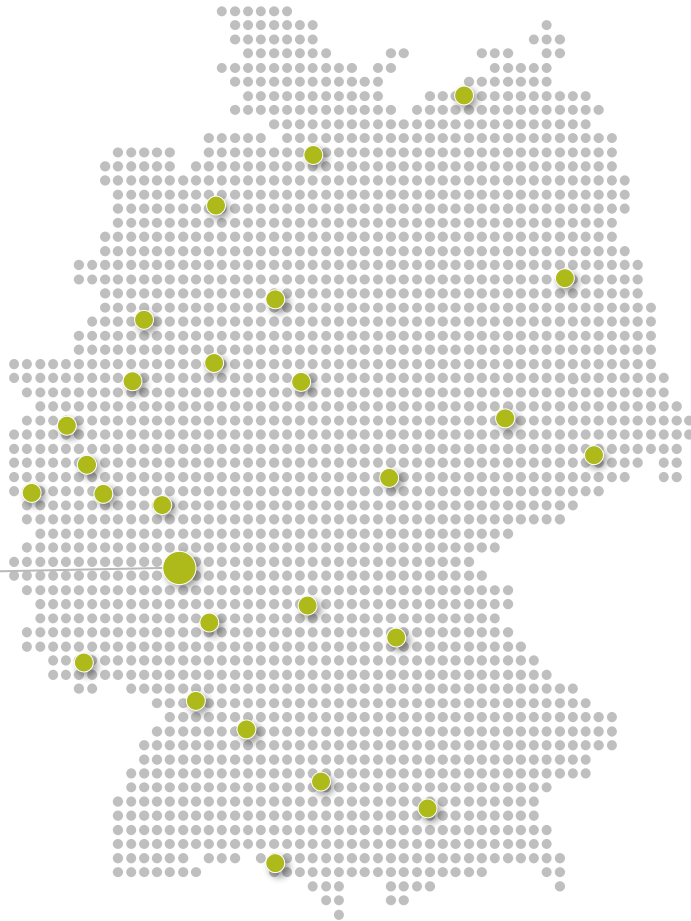
March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

[jochen.zeller@sva.de](mailto:jochen.zeller@sva.de)

## ABOUT US / COMPANY

27  
Locations

Wiesbaden



7 TOP  
Industries



Automotive



Retail



Public



Machine &  
Plant  
Construction



Finance &  
Insurance




Telecommu-  
nication




Healthcare

# ABOUT US / COMPANY

## Employees



## Sales volume in Mio €



## WHAT MAKES US UNIQUE?

/ SVA ...

**BESTER  
IT-DIENSTLEISTER 2023**

COMPUTERWOCHE ChannelPartner

**PLATZ 1** IN DER UMSATZKLASSE  
ÜBER 1 MRD. €

SVA SYSTEM VERTRIEB  
ALEXANDER GMBH

WirtschaftsWoche

**BESTER  
Mittelstands  
Dienstleister  
2023**

SVA System Vertrieb  
Alexander GmbH

Im Vergleich: 19 Unternehmen  
Partner: ServiceValue GmbH  
Ausgabe 11/2023

**Bester MSP**  
MANAGED SERVICE PROVIDER

**Platz 1**

SVA  
System Vertrieb Alexander GmbH

ChannelPartner

COMPUTERWOCHE

› Owner-operated

› High employee satisfaction  
and loyalty


› Long-term technical support  
for solutions delivered by us

› Customer satisfaction and  
recommendation above average

## / HOW IT ALL BEGAN ... 2019



# / BACK TO THE VERY BEGINNING, SPRING 2019 ...

- The customer has ordered a 25 Pbyte GPFS appliance including backup concept – for a site in Poland
- Main use case was Autonomous Driving Validation – streaming large videos to HiL stations
-  2019 → ESS + appliance?
- The ESS was already many things, but far from an appliance
- The installation and upgrade is probably the biggest evolution in ESS.  
Today - after the basic configuration - everything is very simple to do via `essrun` command.



- RHEL rpms and kernel not included, manual download
- xcat based installation with gssdeploy script and a time consuming question (42?) and answer game
- update a running cluster means takeover the RG by yourself, mmshutdown and run the update scripts
- IBM Lab Services was a part of the ESS order



# / AND BECAUSE ESS WAS SO COMPLICATED ...

- We did the entire setup in Germany at SVA and delivered a preconfigured „appliance“ to Poland.
- Which means:
  - Unpack 10t of hardware
  - Plug 3.816 disks
  - Cable everything
  - Install EMS + 7 ESSs with gssdeploy and pack them up again, transport the 10t to Poland, do the unpacking and cabling again
- It takes 6 weeks in Germany
- And three days in Poland



# / AND WHEN EVERYTHING WAS FINALLY DONE

- ESS + file systems + CES ready to use, including swift!
- DR ESS with Spectrum Archive EE and rsync automation was done
- Openstack was ready and uses the storage with S3 and GPFS for VMs



→ Management decision, ADAS consolidation to Germany!





## / UNFORTUNATELY, THE FRAGILE SAS FIBER CABLES ARRIVED LIKE THIS




- Some were replaced immediately because they were defective
- Most of them were gradually replaced in countless error situations



- Despite of some downtimes caused by the cables, the data on the ESSes was always safe!

# / THE SETUP



- A Scale cluster was already running in the target data center, with one 6 PB file system
- Build up with “classic” GPFS on Netapp E-Series, Dell PowerVault and one ESS
- We increased the capacity and performance of this file system to 26 PB
- We separated one ESS + tape as a DR system in a remote data center
- The prod cluster is equipped with 10 CES nodes (NFS+SMB), also used as AFM GWs

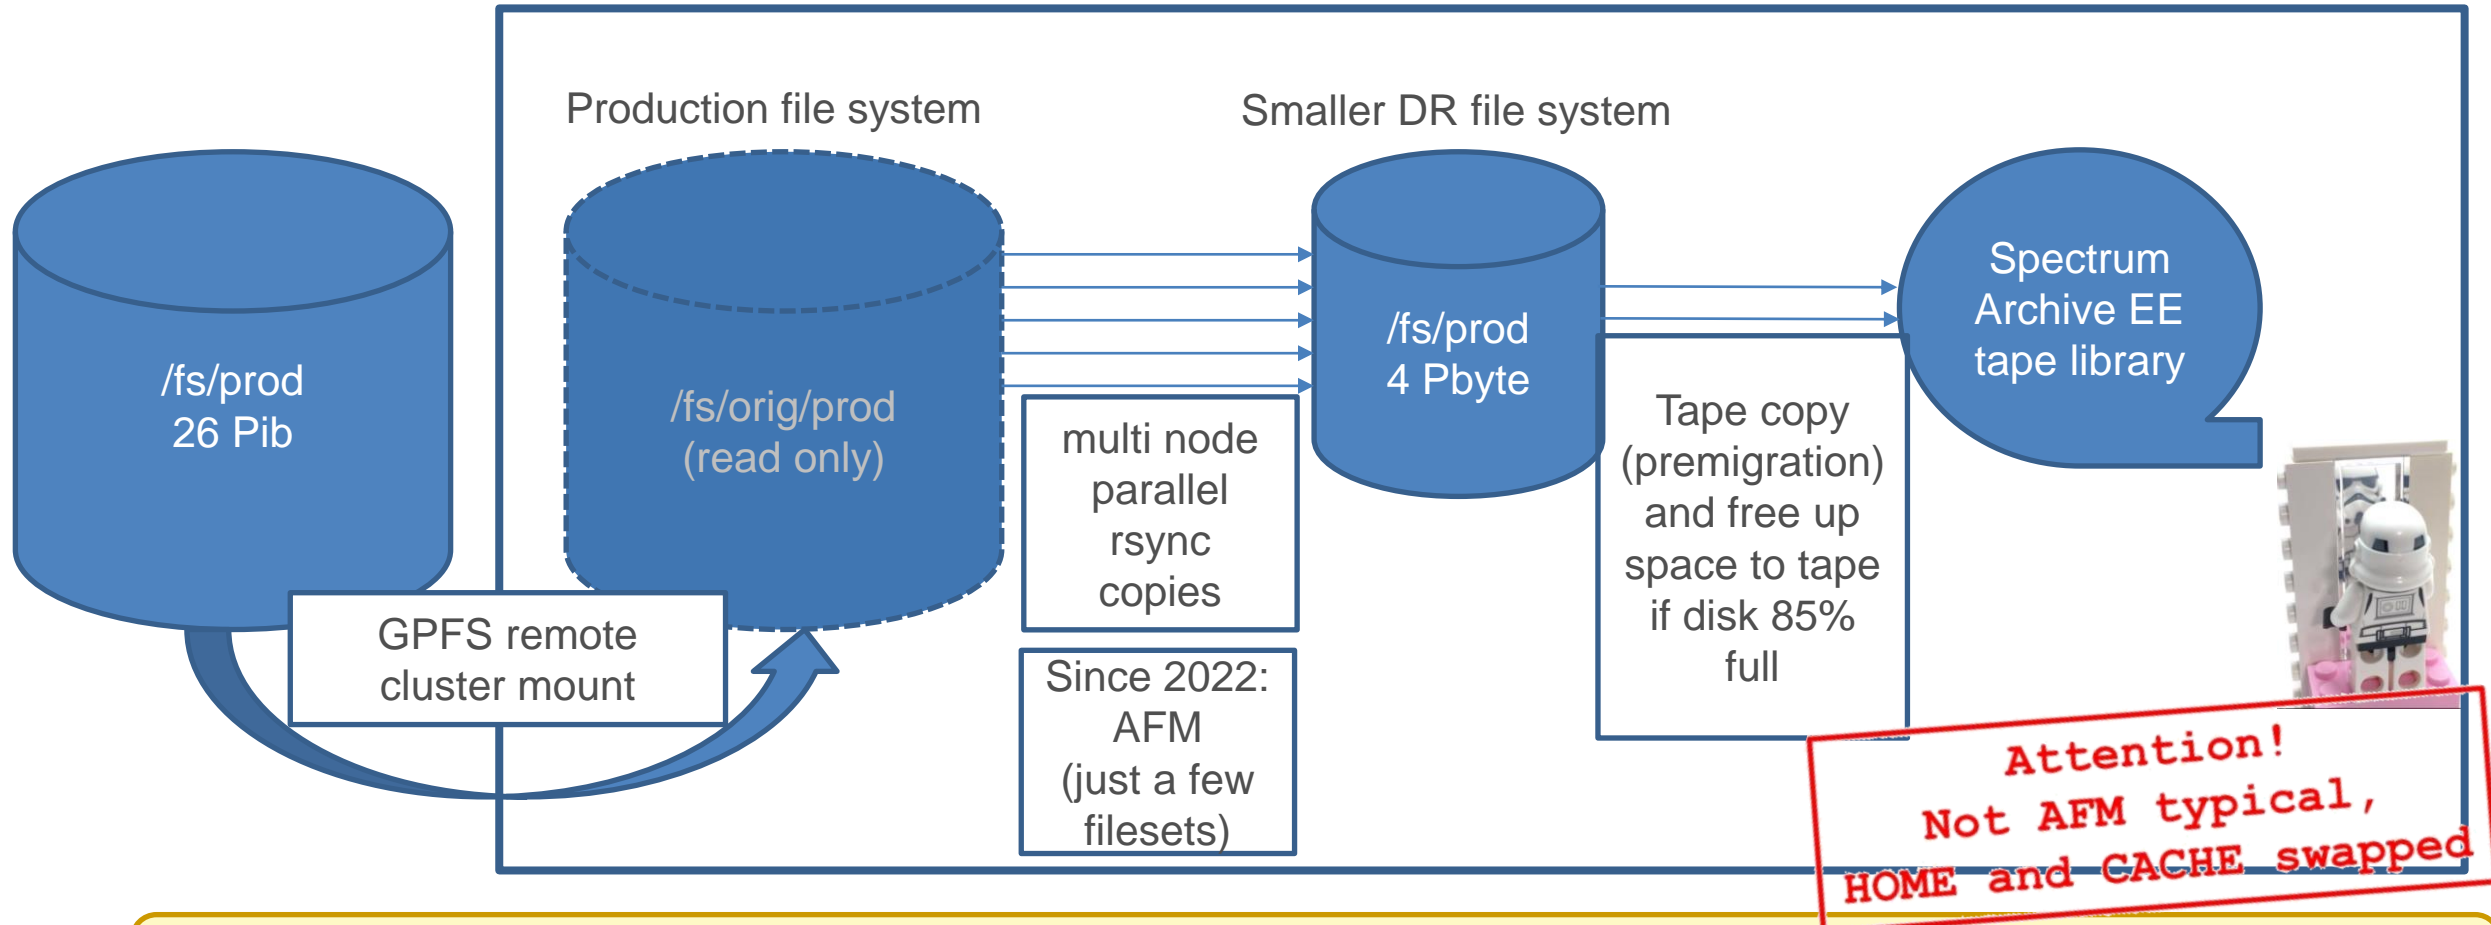


```
mmadddisk prod -F prod-metadata-ESS.cfg --vdisk-set vs_prod_meta01  
--force-nsd-mismatch
```

# / THE SETUP

```
nohup /opt/rsync.gpfs/bin/rsync -aq --gpfs-attrs -H --specials --devices \
--numeric-ids --delete-during $SOURCE/$ENTRY/$MULTI/ $TARGET/$ENTRY/$MULTI/
```

DR system, remote data center



```
mmlcrfileset prod fset2sync -p afmMode=RO,afmtarget=gpfs://fs/orig/prod/projects/fset2sync --inode-space new
mmlinkfileset prod fset2sync -J /fs/prod/projects/fset2sync
Daily: mmafmctl prod prefetch -j fset2sync --directory /fs/prod/projects/fset2sync
```

# / AND WHAT IS DONE WITH THE STORAGE?

Technical computing

Data exchange      ai based quality inspection (factory)

Autonomous driving development - training

Openshift storage

Machine learning      HPC / HPTC

Sensor Dev

CAE / CAD

Autonomous driving development



... and many more



# / HOW TO FILL UP SO MANY ESS'S? MIGRATE DATA FROM NFS STORAGES!

- Typically: You bought a system in 2015, simply expanded in 2017 and 2019
- Also typically: In 2020, extend maintenance for the first nodes, or replace them with new nodes
- New solution: Replace the nodes by ESS, use AFM to migrate the data! So easy:

```
# mmchfs prod --auto-inode-limit
# mmcrfileset prod isi01share -p afmMode=R0,afmTarget=nfs://isistorage.local:/isi01share
--inode-space new
# mmlinkfileset prod isi01share -J /prod/projects/isi01share
# mmchfileset prod isi01share -p afmEnableAutoEviction=no
# mmafmctl prod prefetch -j isishare01 --directory /prod/projects/isi01share --prefetch-
threads 32
...
```

- Cookbook : <https://www.ibm.com/docs/en/storage-scale/5.1.9?topic=afm-data-migration-by-using-migration-enhancements>
- ~~Please vote: Enhance AFM to transfer NFSv4 ACLs from NFS servers like Isilon and store the ACLs correct as GPFS ACLs~~

~~<https://ibm-sys-storage.ideas.ibm.com/ideas/GPFS-I-1050>~~



# / OFF TO THE CLOUD WITH IT




- Scale is able to sync filesets with cloud storage (AWS, Azure, Google)
- Use cases in this installation:
  - sync data between Scale fileset and cloud storage to work onprem and in cloud with the same data
  - Upload/archive selected files to cloud (AFM MU – manual update - mode)
- AFM to cloud is simple if the AFM gateway nodes have direct https access to the cloud storage:

Bucket: <https://gpfs-upload.s3.eu-north-1.amazonaws.com>

Region: eu-north-1

```
# cat aws.key
akey:JVAKESAAFVDSQ62SLMDRA
skey:mGbj4hGzNWK79f916fddsnGHJUmts30Fda2Fr4Bn
# mmafmcoskeys gpfs-upload:eu-north-1@s3.eu-north-1.amazonaws.com set --keyfile aws.key
# mmafmcosconfig prod fset2AWS --endpoint https://eu-north-1@s3.eu-north-1.amazonaws.com --dir fset2AWS
--prefix fromGPFS --mode iw --object-fs --debug --cleanup --bucket gpfs-upload --acls --fast-readdir
```

# / OFF TO THE CLOUD WITH IT



- Scale is able to sync

- Use cases in

- sync d

- I

- All

- B

- Reg

Question: Who are you – tape or cloud?

Question: Are you already using AFM?

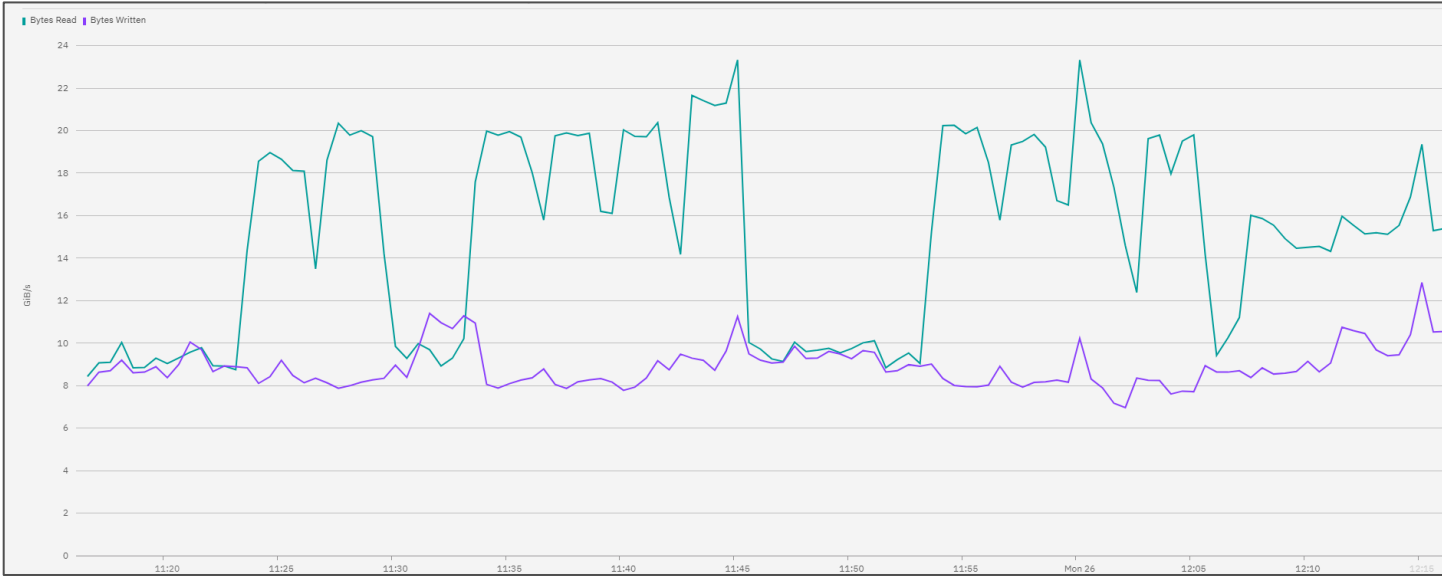
Question: Are you already using AFM with object storage?

Question: Are you already using AFM with cloud storage?

```
# cat  
akey:J  
skey:mGb  
# mmafmcoskeys SP  
# mmafmcosconfig pr  
--prefix fromGPFS --mod  
key  
onaws.com --dir fset2AWS  
et gpfs-upload --acls --fast-readdir
```

# / 5 YEARS AFTER THE TRUCK ARRIVED THE FIRST TIME

- Another truck arrived, but smaller ...
- New ESS environment is optimized for energy savings
- Sizing is based on current capacity and performance values – not on project forecasts

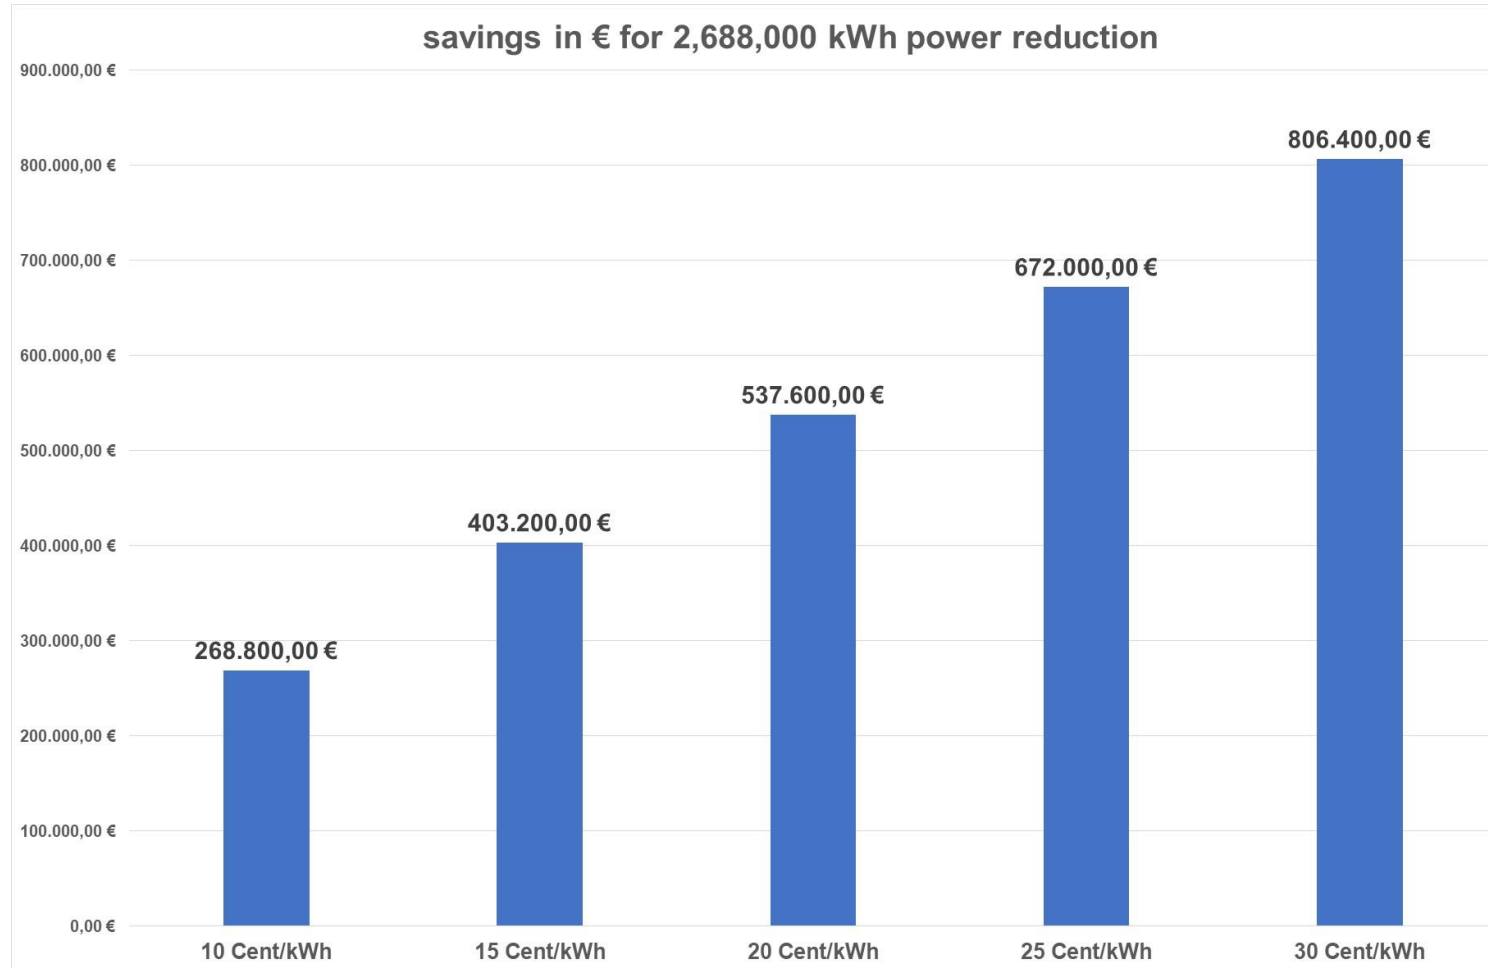


	6x ESSGL6x + 1 x ESSGS4	2 x ESS3500 H6
Capacity	26 PiB	16 PiB
NL-SAS max perf	~330GB/s	~60GB/s
Power cons. 5 years	3,648,000 kWh	960,000 kWh





# / THE BIG PLUS – THE POWER SAVINGS



Power savings for reduced cooling effort not included!



# / SUMMARY

- I could have told you much more about that installation ...  
about NFSv4, AD, strange replication behaviour, illplaced files, EE, ptrash, ... but the time ...
- Despite all the toeing and froing, this installation is still a success story
- I'm really happy about the evolution of the ESS automation scripts, it's close to an appliance
- I like the ESS3500 and looking forward to the SSS6000 hybrid
- I ❤️ AFM
- AFM to cloud is exciting for selected use cases
  - also to get workloads from the cloud back on premises
- GPFS is just magic!



# / THE END!

## / SUMMARY

- I could have told you much more about that installation ...  
about NFSv4, AD, strange replication behaviour, illplaced files, EE, ptrash, ... but the time ...  
Despite all the toeing and froing, this installation is still a success story  
I'm really happy about the evolution of the ESS automation scripts, it's close to an appliance  
like the ESS3500 and looking forward to the SSS6000 hybrid
- I ❤️ AFM
- ... to cloud is exciting selected use cases
- also ... in the cloud back on premises
- GPFS is just ...



05.03.2024 / 18



Many thanks for  
your attention!



## JOCHEN ZELLER

IT Architect

Technical Leader IBM Storage Scale

---

Phone.: +49 151 180 256 77

Mail: [jochen.zeller@sva.de](mailto:jochen.zeller@sva.de)