SSSD23 - IBM ESS as super fast and super scalable Backup Storage for Commvault

jochen.zeller@sva.de

SVA ISS Team

ABOUT US

/ COMPANY

26
Locations

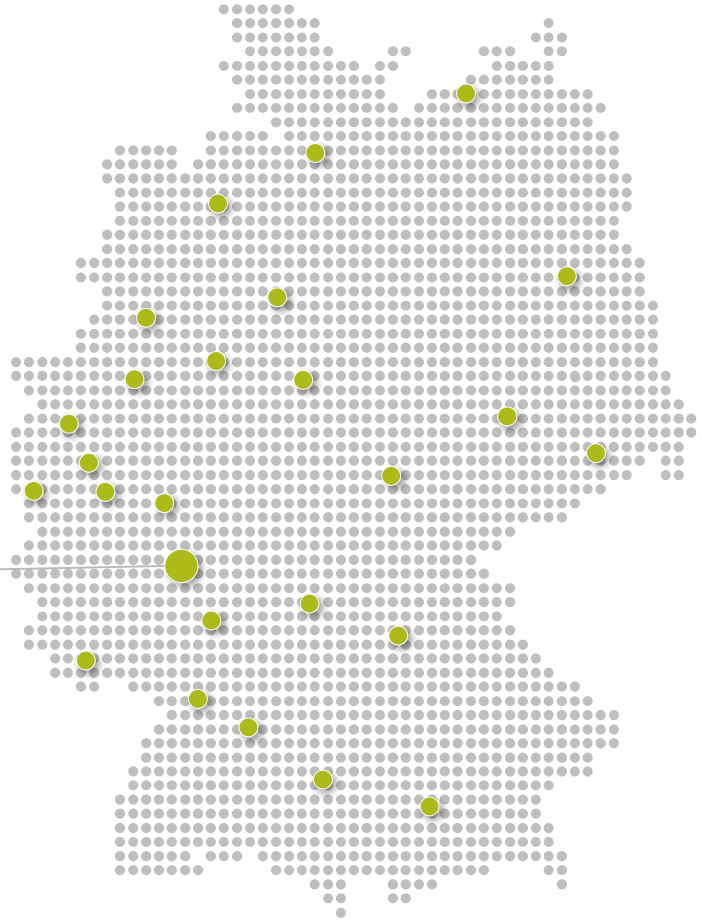

7 TOP
Industries

- Automotive
- Retail
- Public
- Machine & Plant Construction
- Finance & Insurance
- Telecommu-nication
- Healthcare

WHAT MAKES US UNIQUE?

/ SVA ...

- Owner-operated
- High employee satisfaction and loyality
- Long-term technical support for solutions delivered by us
- Customer satisfaction and recommendation above average

/ WHY SPECTRUM SCALE / ESS AND BACKUP?

- Backup systems typically require large data storage
- In the past, tape was preferred as large data storage, HDD / SSD was just a cache
- Today, most backup data is kept on HDD / SSD / NVMe for fast parallel restore
- Large, fast, parallel = ESS

ESS3500 capacity model
up to 816 x 20TB HDD
0,3 to 12 Pbyte usable capacity
Up to 48GB/sec

/ GPFS + TSM | TODAY IBM STORAGE SCALE + IBM STORAGE PROTECT

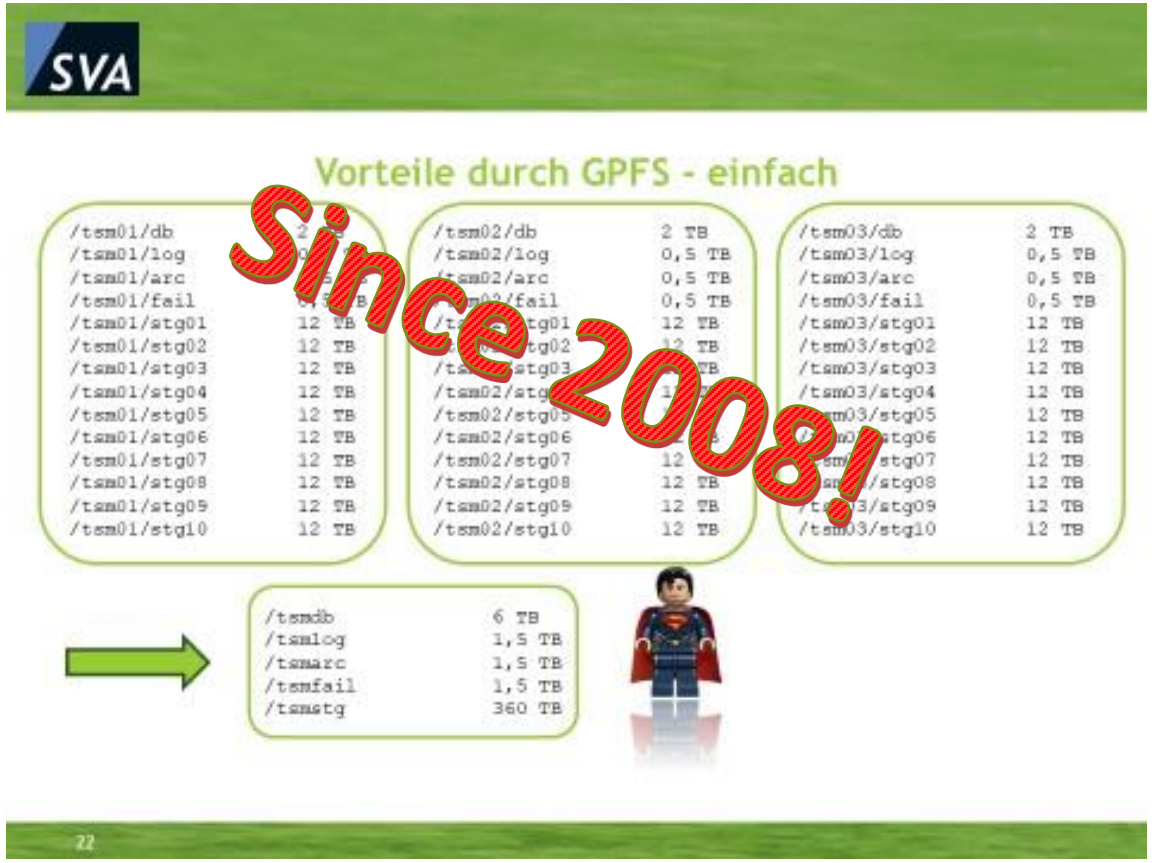

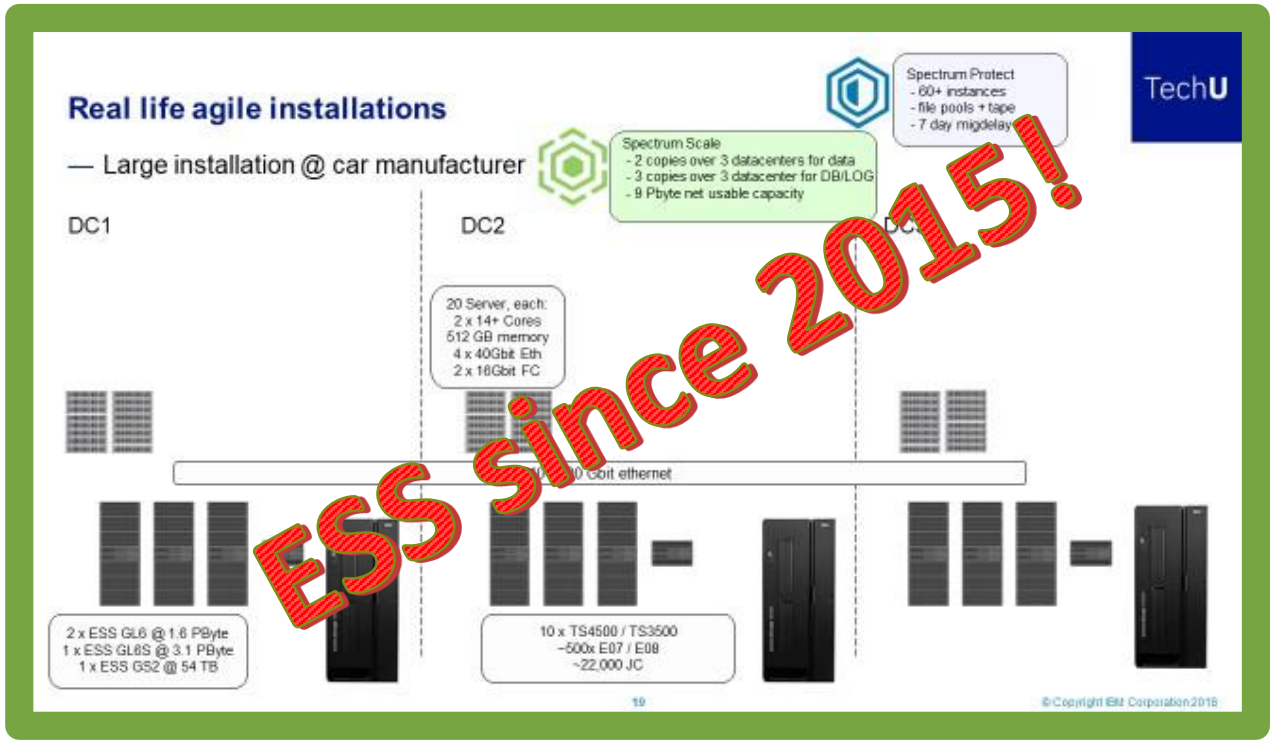

/ SPECTRUM SCALE + VEEAM

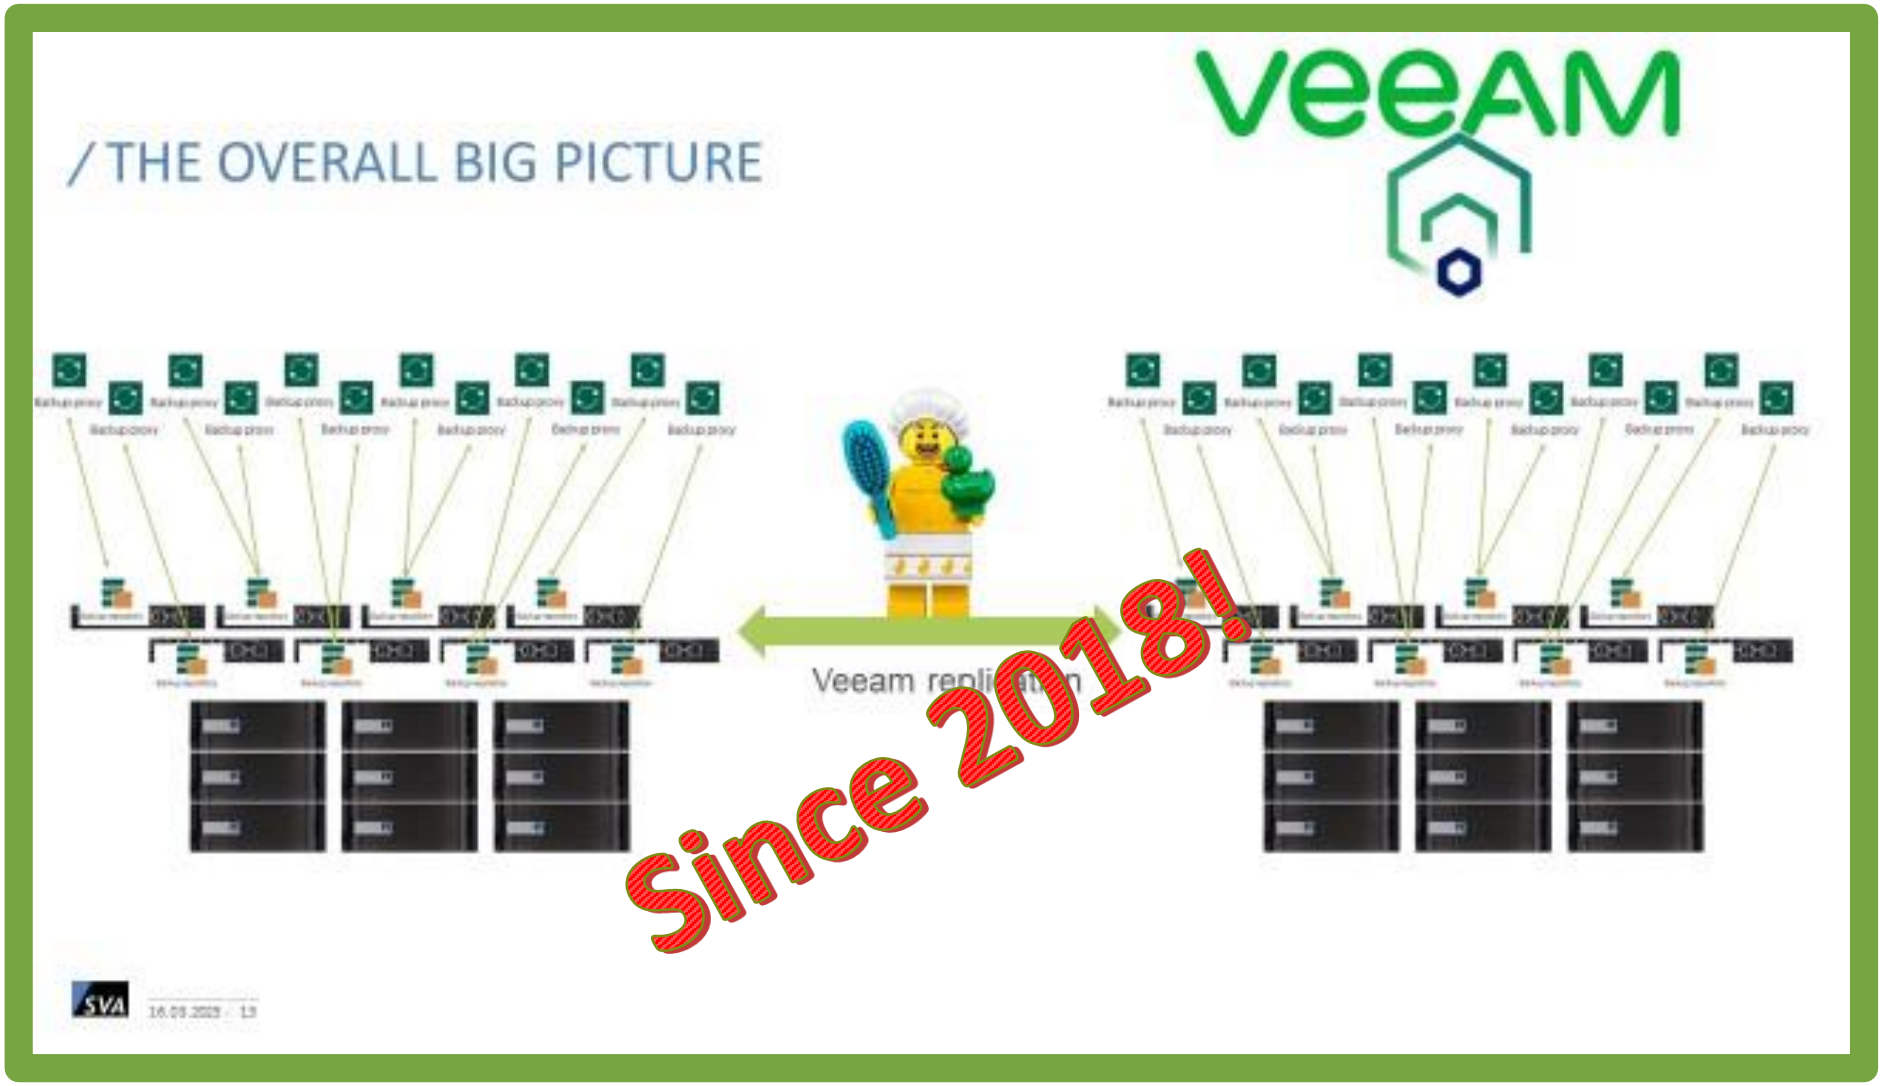

/ NEXT: IBM ESS AND COMMVAULT

- Commvault Data Protection is widely used and growing
- Gartner ranks Commvault among the top 3 leaders
- Commvault scales very well, e.g. by the number of media agents
- Commvault typically stores data on “disk libraries”, all media agents can share the same “disk library”

Sound like perfect for ESS. But:

- Commvault typically stores data duplicated, typical block size 128kB
- Reading deduplicated data is 100% random read @ 128kB
- Restores and tape copies must enrich / rehydrate / duplicate the data again
- Is this small block size optimal for ESS?

/ COMMVAULT OVERVIEW – DEDUPLICATION

- We were very careful when sizing our first ESS for Commvault
- 2,4 Pbyte ESS5000 SL6 with 550 x 6TB, to have enough HDDs for the 128k IOPS
- With the ESS and 6 media agents, we did some dedup testing with different block sizes

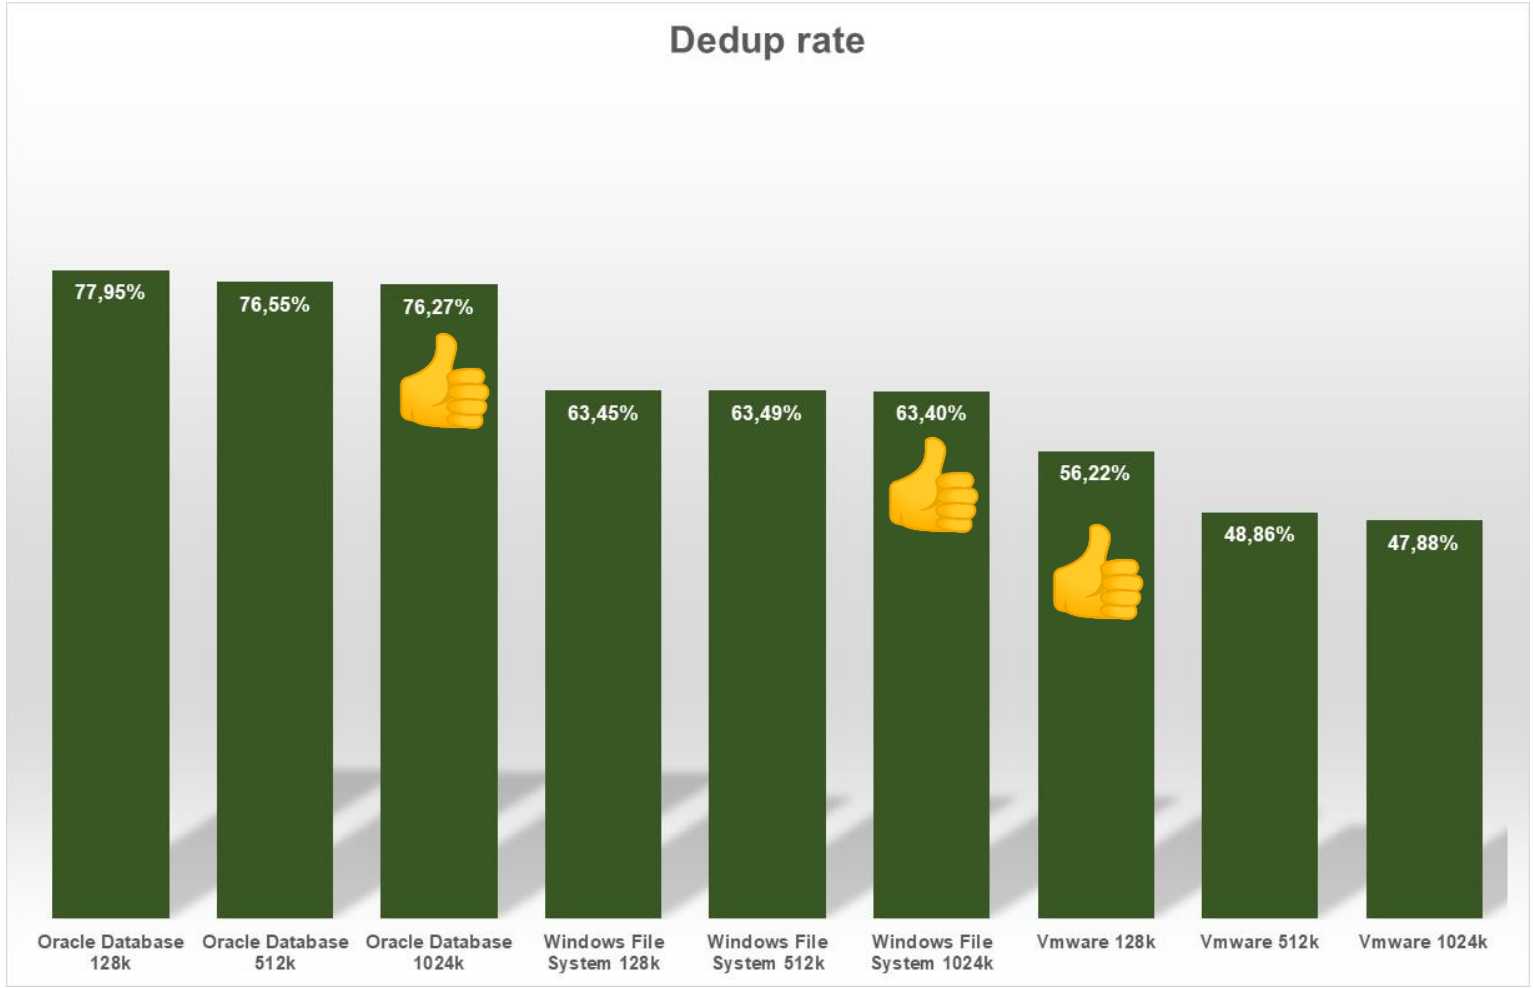

/ COMMVAULT OVERVIEW – DATA STREAM CLASSIC

- Commvault typical data transport

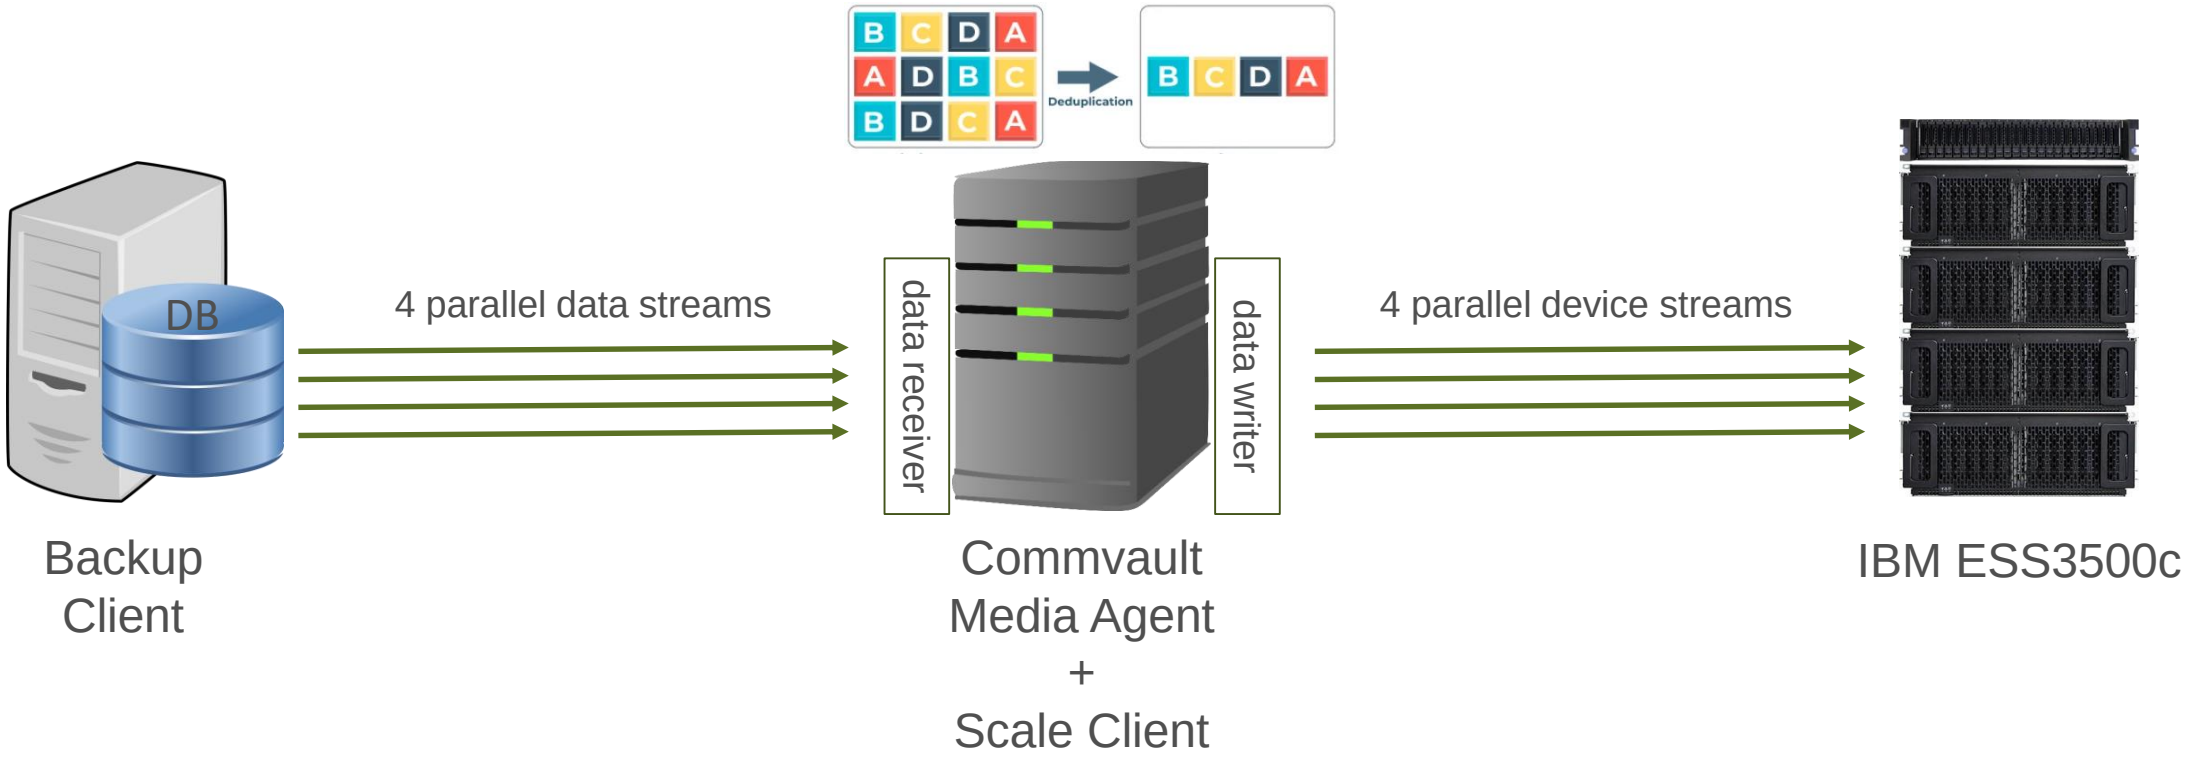

/ COMMVAULT OVERVIEW – DATA MULTIPLEXING

- Commvault multiplexing data transport

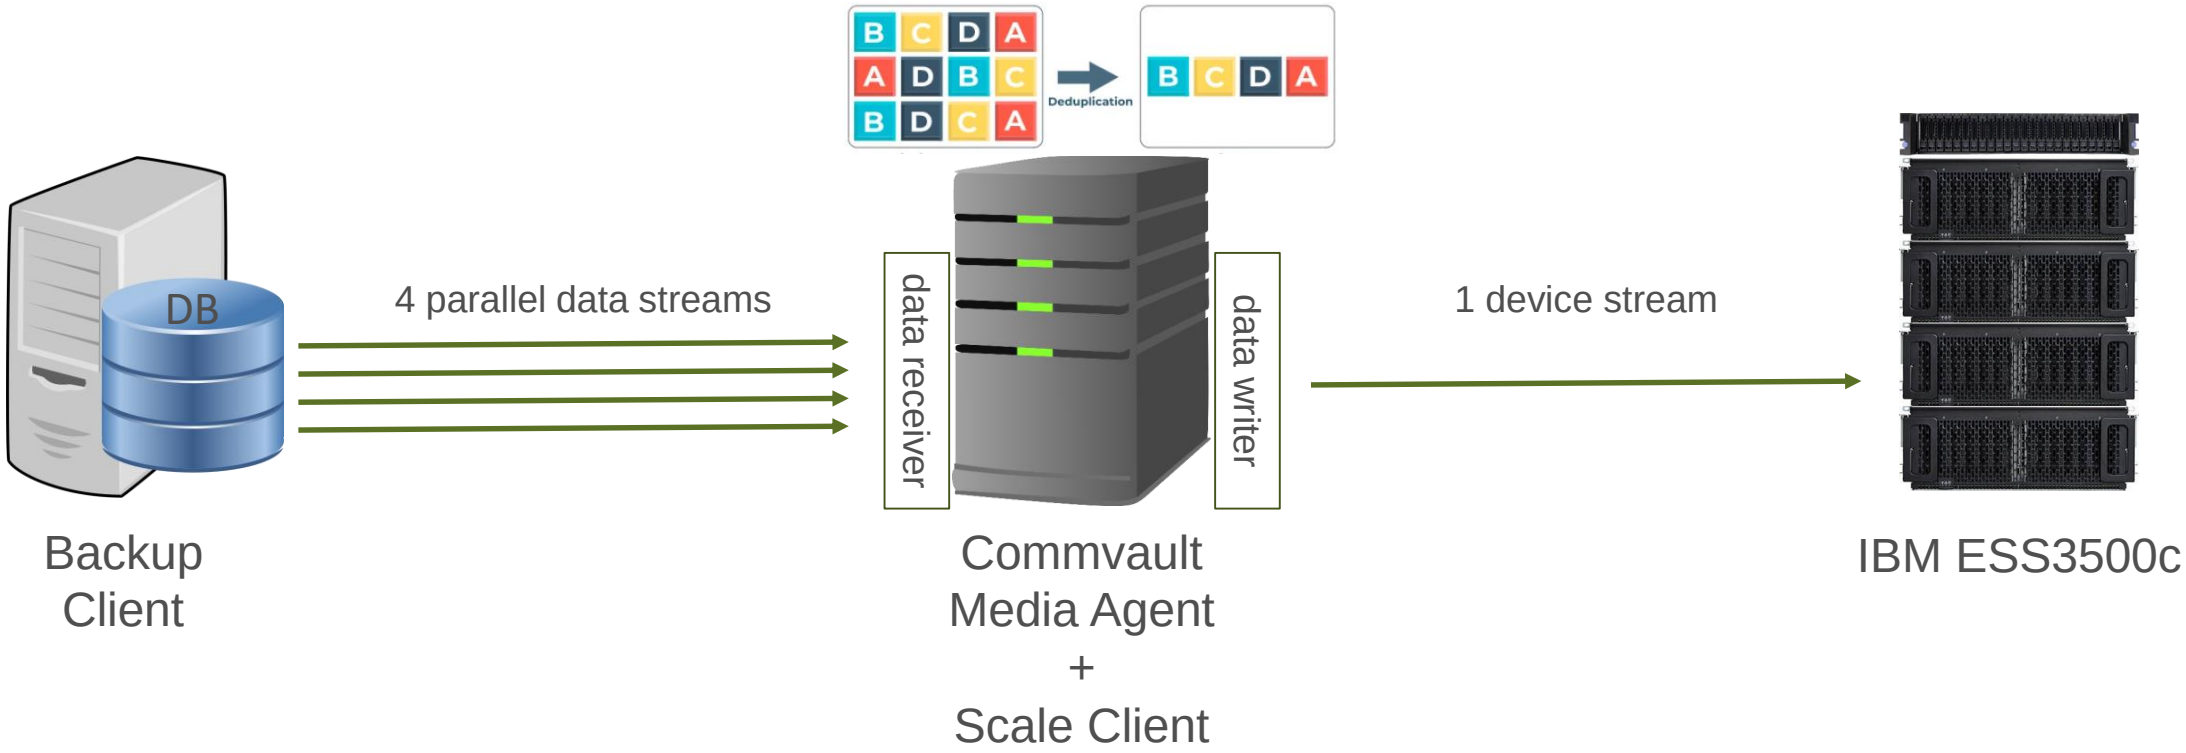

- This can be exciting to use an ESS properly.

/ ADVANTAGES OF AN ESS BASED DISK LIB FOR COMMVAULT

- One super fast and super large shared disk lib for all Media Agents (instead of many [small, slow] local)
  - A huge shared disk lib reduces storage management overhead
- Spectrum Scale protocol for data access (instead of NFS for other shared storages)
  - Very fast parallel reads and writes
  - End-to-end checksum
- IBM ESS has a great price performance ration for multi petabyte configuration
- IBM ESS scales linear and horizontal
- Disk errors and rebuilds do not interfere with backup jobs
- An ESS still works well at the limit (98% file system filling level, max IOPS or throughput)
- Oh yeah, did I mention that an ESS is really fast?

/ COMMVAULT + ESS SETUP

hundreds of Commvault
backup clients

DB, Apps, File, VM

Backup &
Restore

20 x TS1160

ESS5000
550 x 6TB HDDs
4 x 100GBE

ESS3500 C
408 x 14TB HDDs
4 x 100GBE

write tape
copy from ESS

write deduplicated
read rehydrated

16 Media Agents, 2-Socket, 2x25Gbit to clients, 2 x 25Gbit to ESS,
2 x 32Gbit FC to tape

/ SPECTRUM SCALE CONFIG PARAMETER USED

- Create a nodeclass for the Media Agents
- Run **`/usr/lpp/mmfs/samples/gss/gssClientConfig.sh`** to set ESS client perf values, e.g.

```
./gssClientConfig.sh -P 16G MediaAgents
```

- Additional configs:

`prefetchAggressivenessRead=1`

Prefetch only if sequential read detected. disable prefetching is also possible (=0), but reduces Commvault throughput

`maxTcpConnsPerNodeConn=4`

multi-connection over TCP. Using more connections may improve the overall performance.

https://www.ibm.com/support/pages/ibm-spectrum-scale-multi-connection-over-tcp-mcot-tuning-may-be-required

/ PERFORMANCE

- Example. Copy VM backup data to tape. 128k random read. 1260MiB/s : 10334 ops/s = 122kiB = 125kB

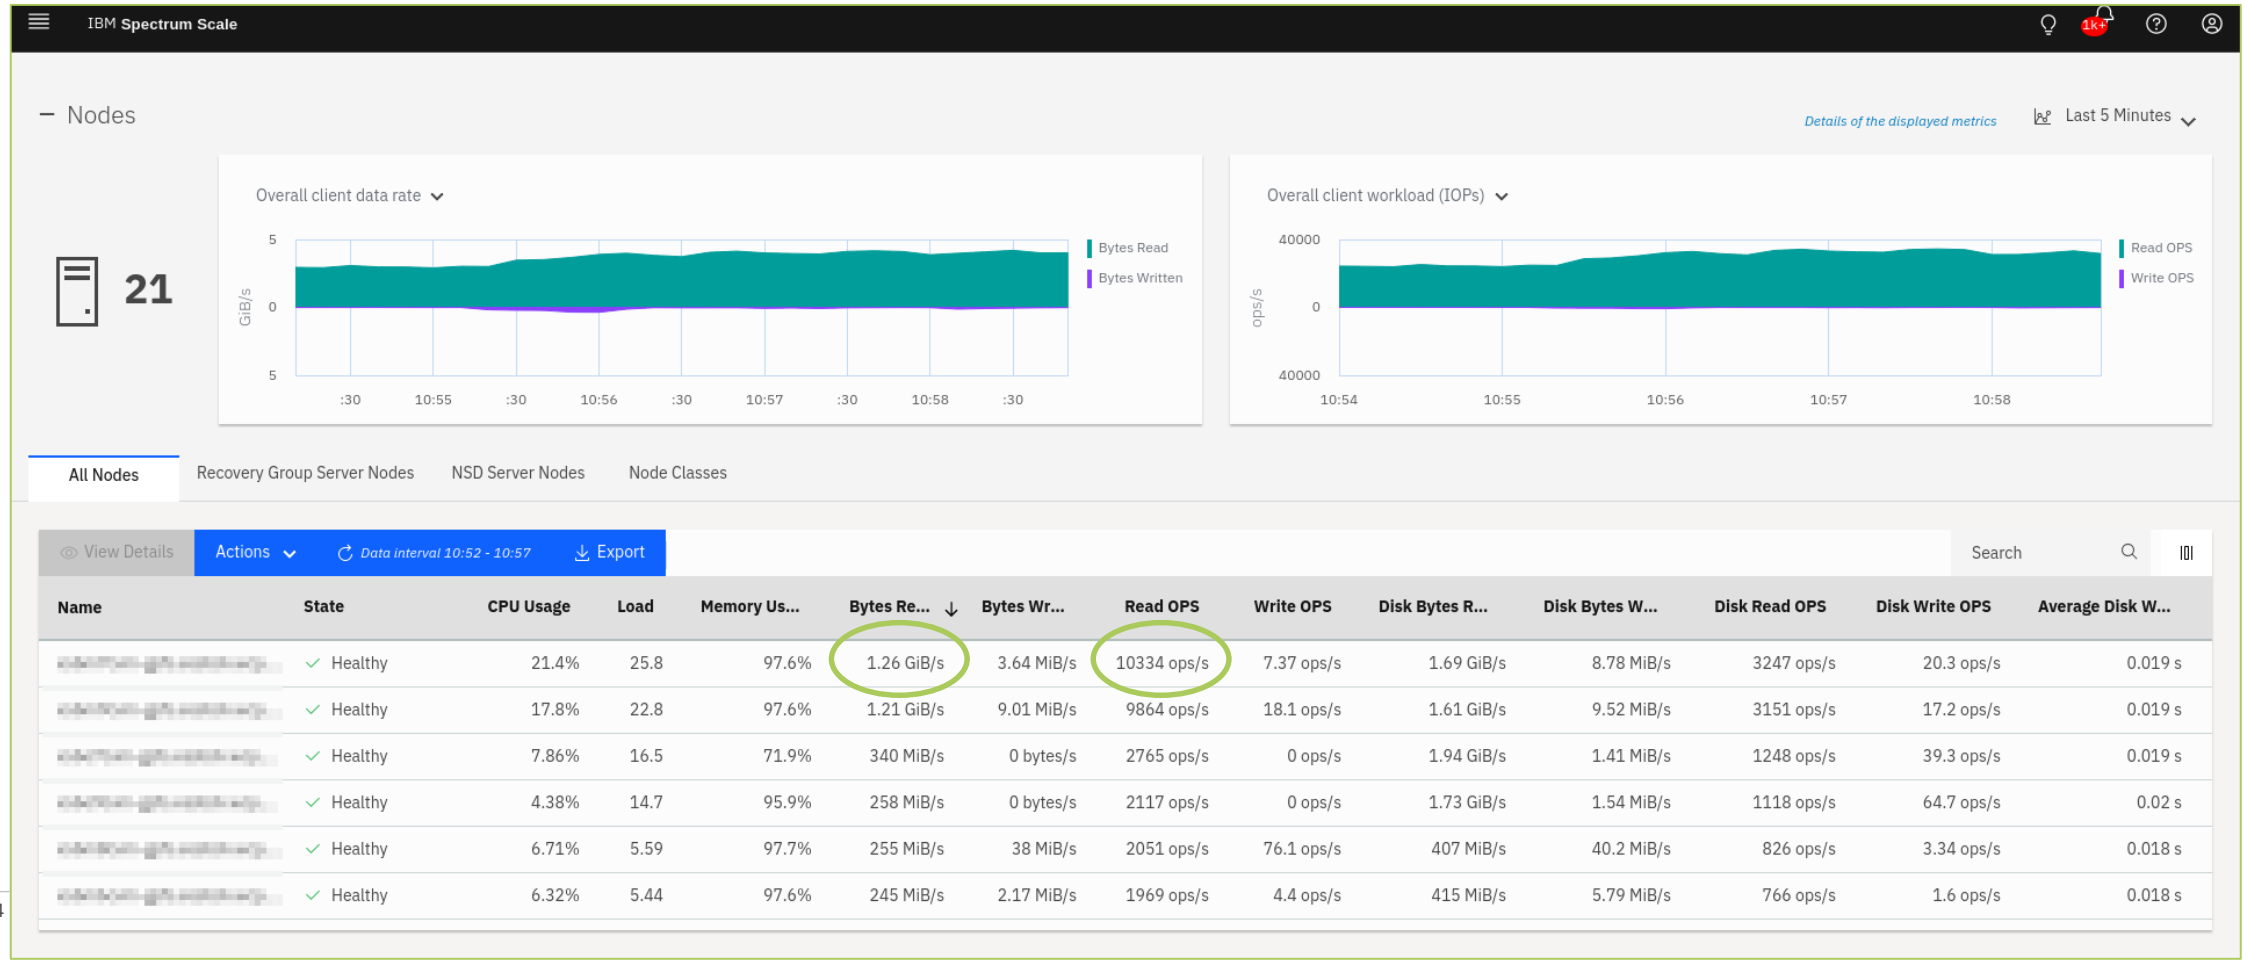

Name	State	CPU Usage	Load	Memory Us...	Bytes Re...	Bytes Wr...	Read OPS	Write OPS	Disk Bytes R...	Disk Bytes W...	Disk Read OPS	Disk Write OPS	Average Disk W...
[redacted]	Healthy	21.4%	25.8	97.6%	1.26 GiB/s	3.64 MiB/s	10334 ops/s	7.37 ops/s	1.69 GiB/s	8.78 MiB/s	3247 ops/s	20.3 ops/s	0.019 s
[redacted]	Healthy	17.8%	22.8	97.6%	1.21 GiB/s	9.01 MiB/s	9864 ops/s	18.1 ops/s	1.61 GiB/s	9.52 MiB/s	3151 ops/s	17.2 ops/s	0.019 s
[redacted]	Healthy	7.86%	16.5	71.9%	340 MiB/s	0 bytes/s	2765 ops/s	0 ops/s	1.94 GiB/s	1.41 MiB/s	1248 ops/s	39.3 ops/s	0.019 s
[redacted]	Healthy	4.38%	14.7	95.9%	258 MiB/s	0 bytes/s	2117 ops/s	0 ops/s	1.73 GiB/s	1.54 MiB/s	1118 ops/s	64.7 ops/s	0.02 s
[redacted]	Healthy	6.71%	5.59	97.7%	255 MiB/s	38 MiB/s	2051 ops/s	76.1 ops/s	407 MiB/s	40.2 MiB/s	826 ops/s	3.34 ops/s	0.018 s
[redacted]	Healthy	6.32%	5.44	97.6%	245 MiB/s	2.17 MiB/s	1969 ops/s	4.4 ops/s	415 MiB/s	5.79 MiB/s	766 ops/s	1.6 ops/s	0.018 s

/ PERFORMANCE

- Example. Copy DB backup data to tape. 1M random read. Parallel backup and other copy jobs. Last 7 days.

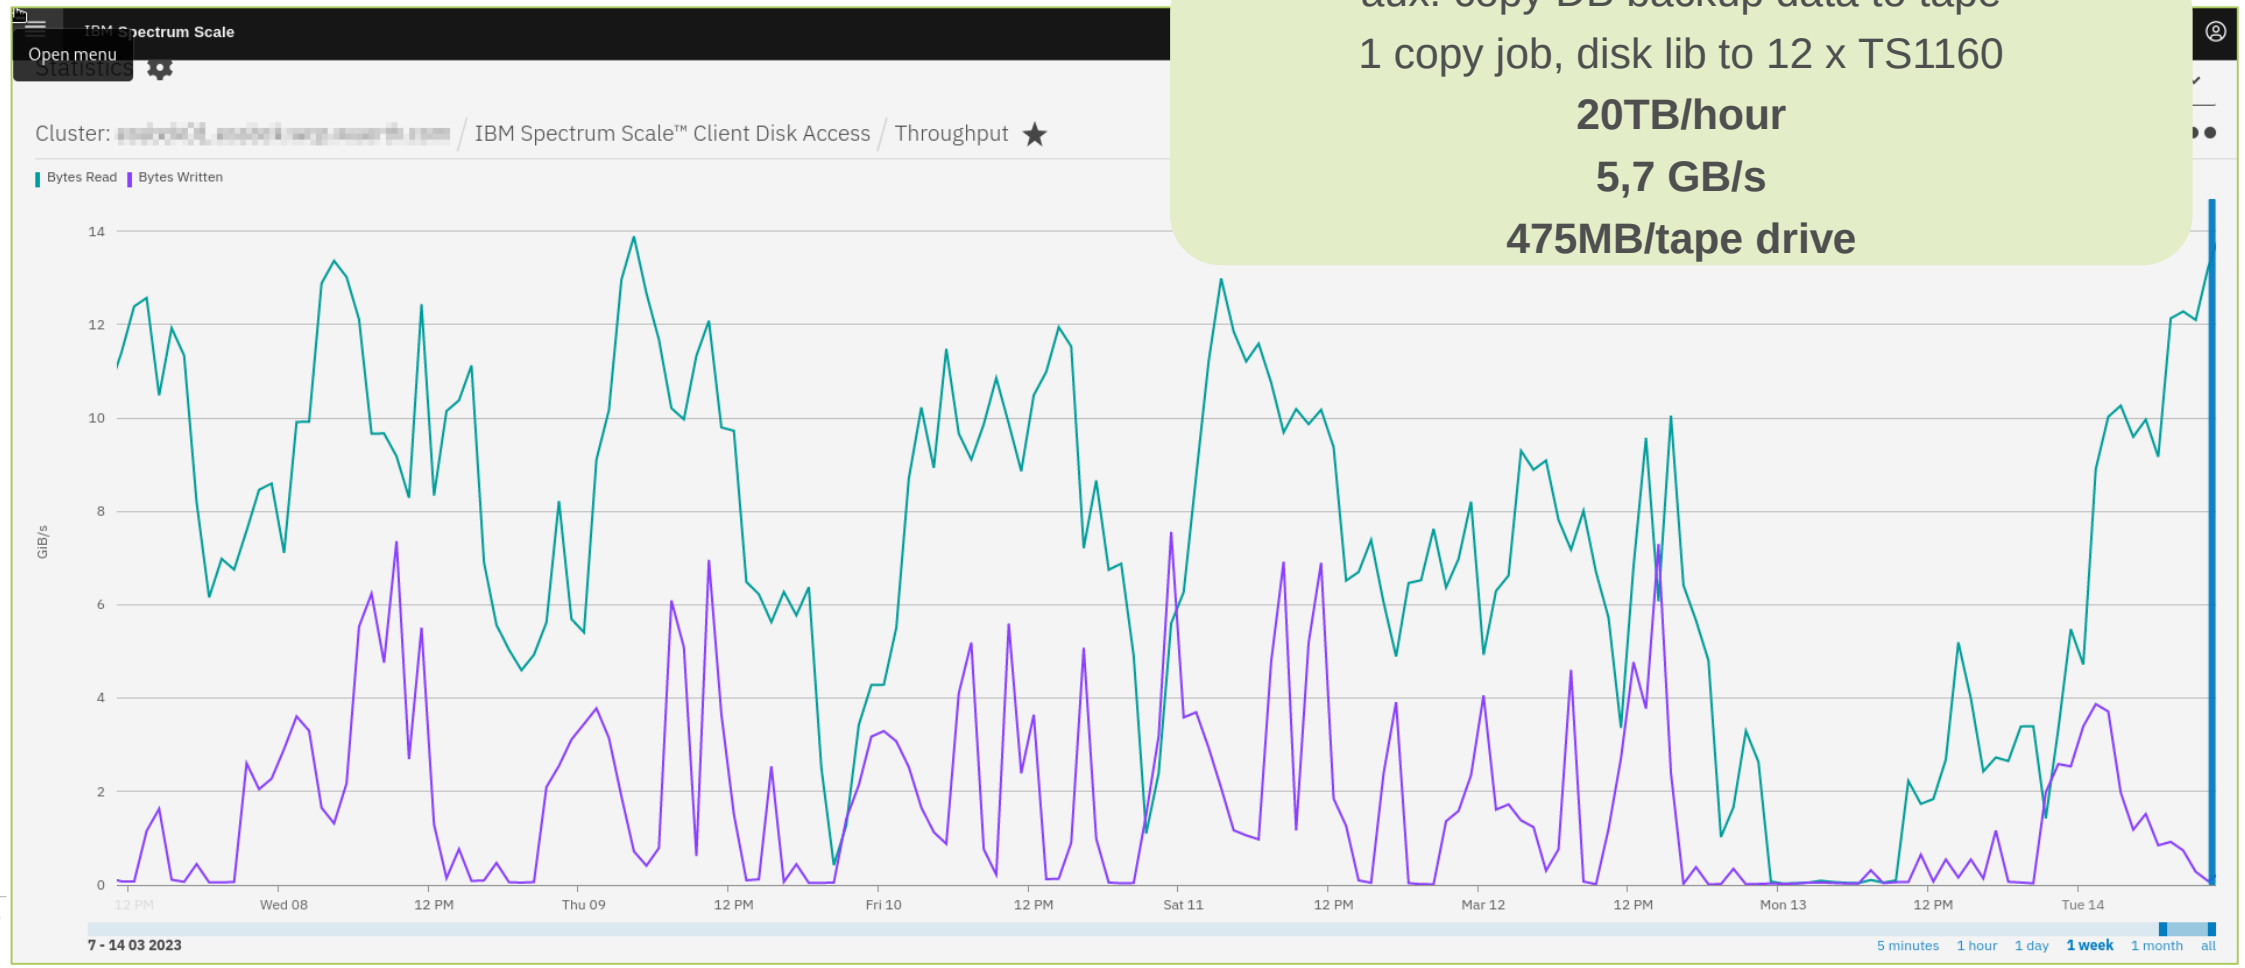

aux. copy DB backup data to tape
1 copy job, disk lib to 12 x TS1160
20TB/hour
5,7 GB/s
475MB/tape drive

/ PERFORMANCE

- Example. Copy DB backup data to tape. 1M random read. Parallel backup and other copy jobs. Last hour.

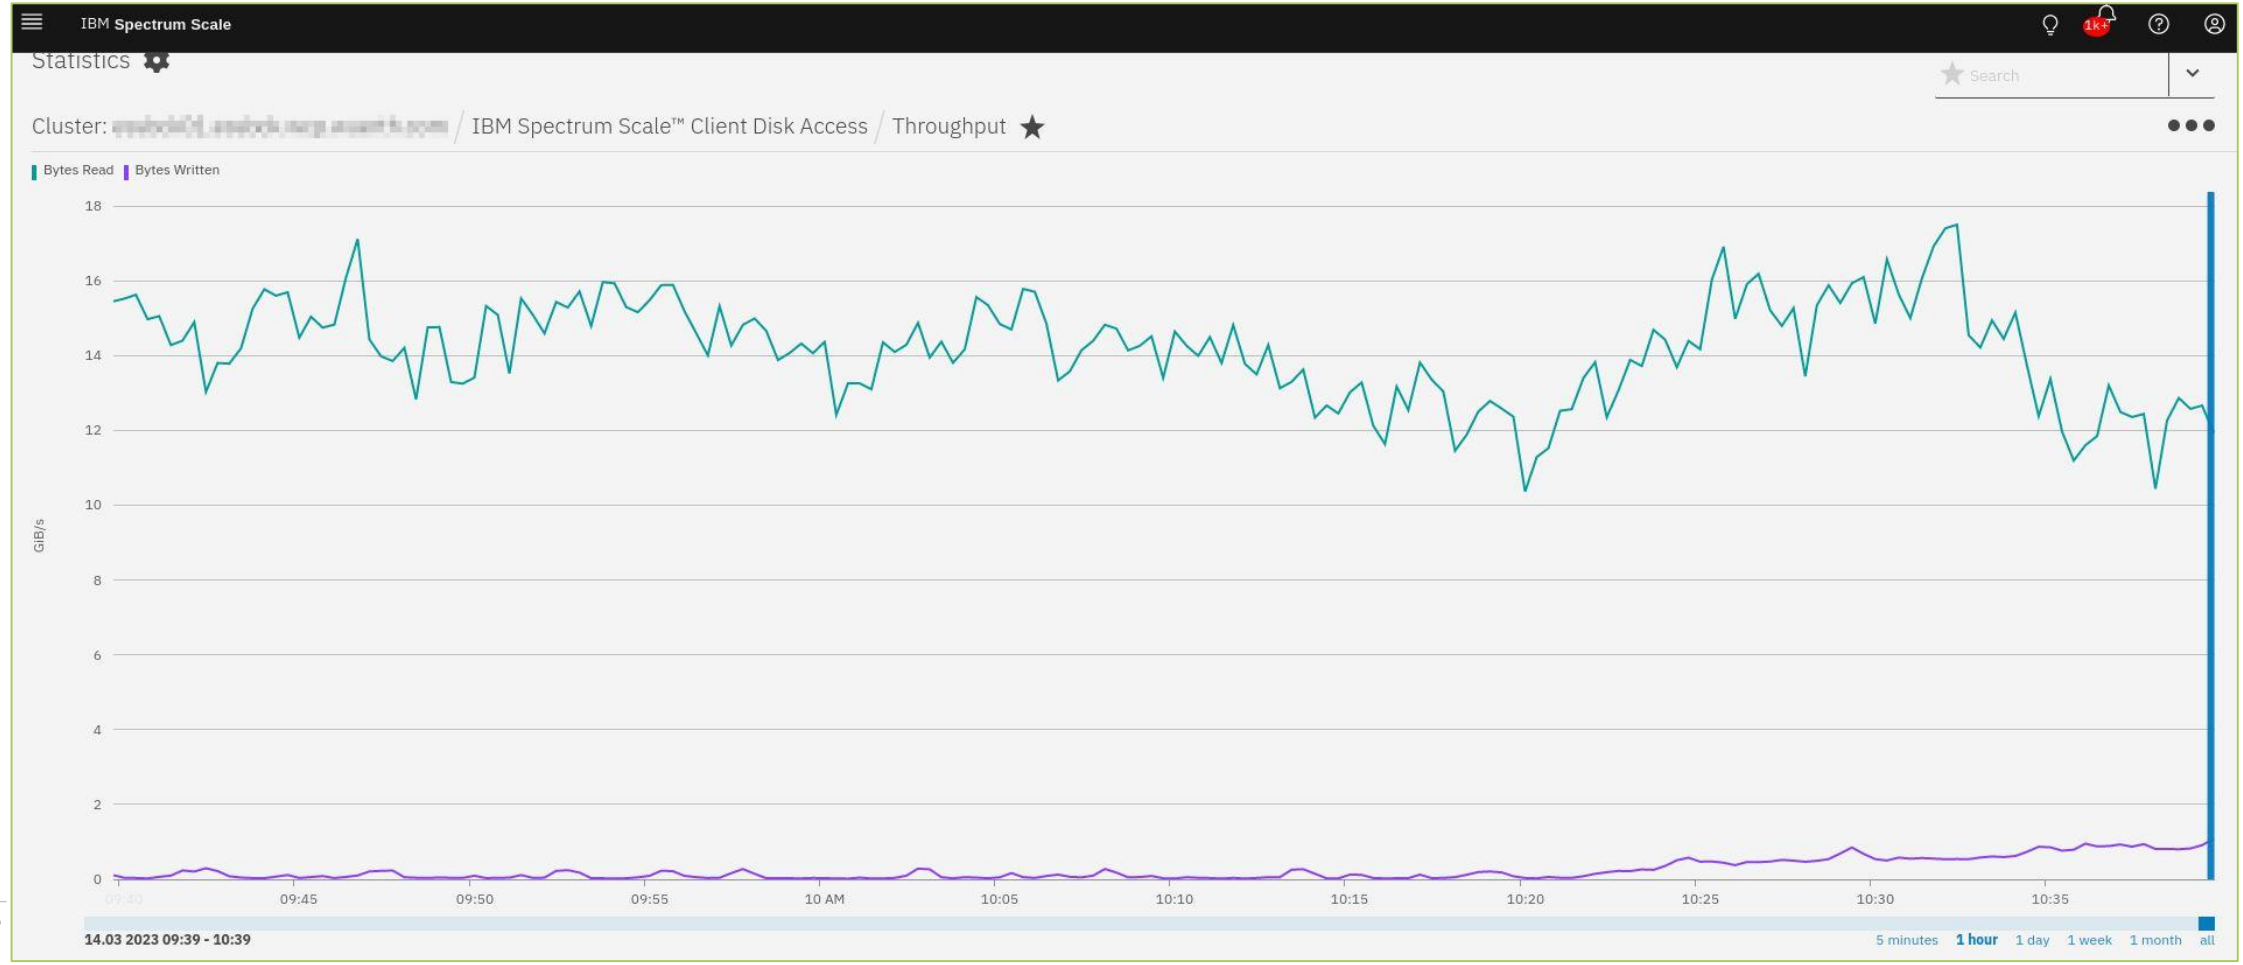

/ PERFORMANCE

- Example. 1 week after adding second ESS. One file system per ESS. Throughput per file system. 1 hour.

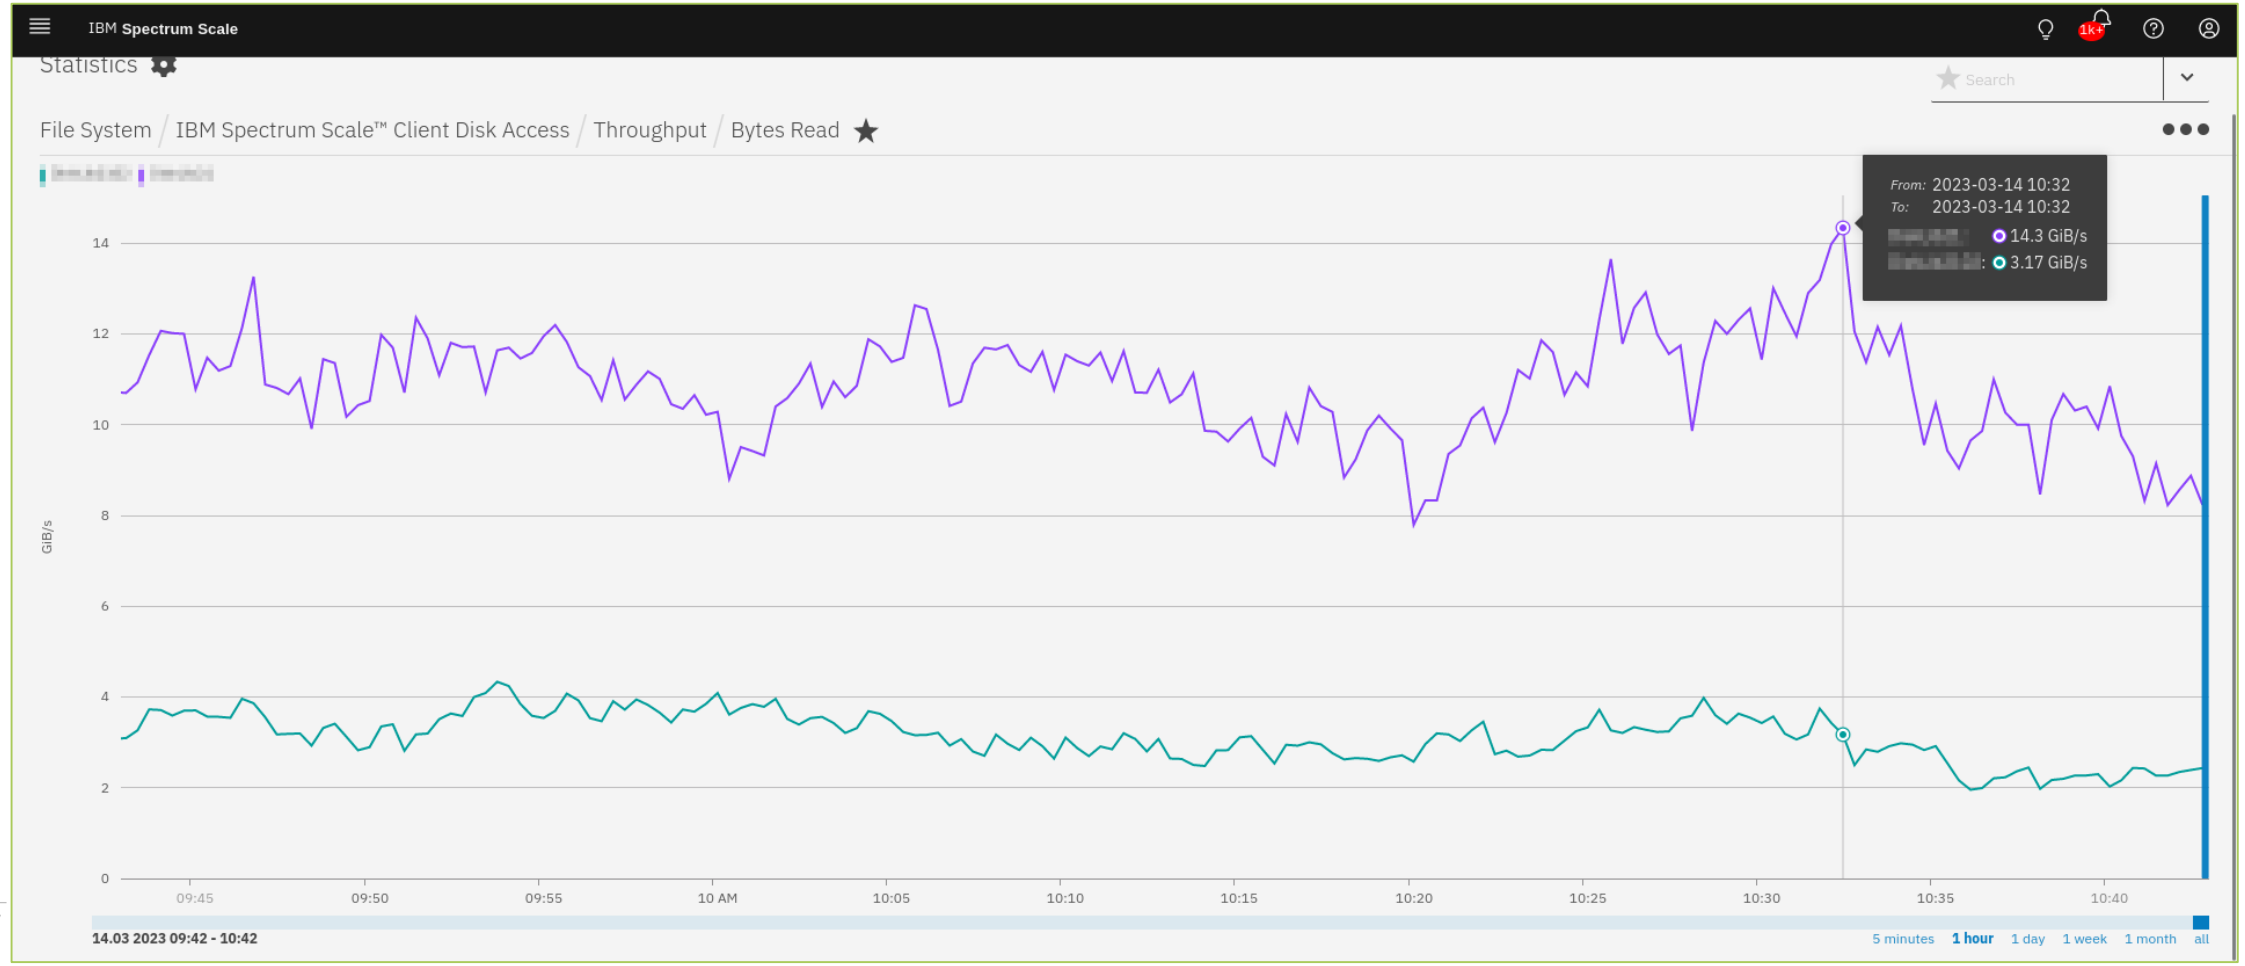

/ FINALLY AVAILABLE – BRING YOUR OWN EMS

- Awesome! After waiting for a long time!
- No more outdated, overpriced, HDD-equipped POWER9!
- Bring your own EMS absolutely with 1TB NVMe boot drives, to make the GUI fast
- Bring your own EMS, Dell, HPE, Fujitsu, Lenovo, NEC, what ever, doesn’t care
- The requirements are still very restricted, almost the 1:1 replica of a POWER9 EMS but with AMD CPU
- Running an EMS VM inside your own EMS server
- Running a podman container inside the VM
- For more information:

  https://www.ibm.com/docs/en/ess-p8/6.1.6?topic=guide-bring-your-own-ems

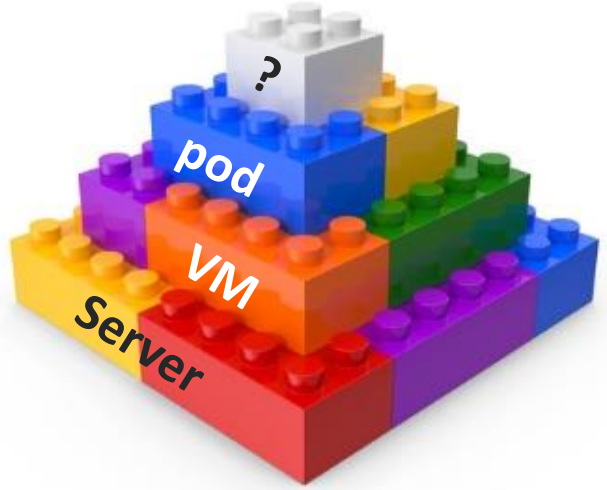

/ SUMMARY

- IBM ESS and Commvault are a perfect match
- IBM ESS provides a huge and fast disk lib to all Media Agents
- One huge disk lib reduces storage and Commvault administration effort and makes both more simple
- Works very well even with small block sizes, like the default 128K
- Improves backup, auxiliary copy (tape copy) and restore times
- Oh yeah, did I mention that an ESS is really fast?

The good thing about Spectrum Scale is that you can design and adjust a lot.

The bad thing about Spectrum Scale is that you can design and adjust a lot.

/ THE END!

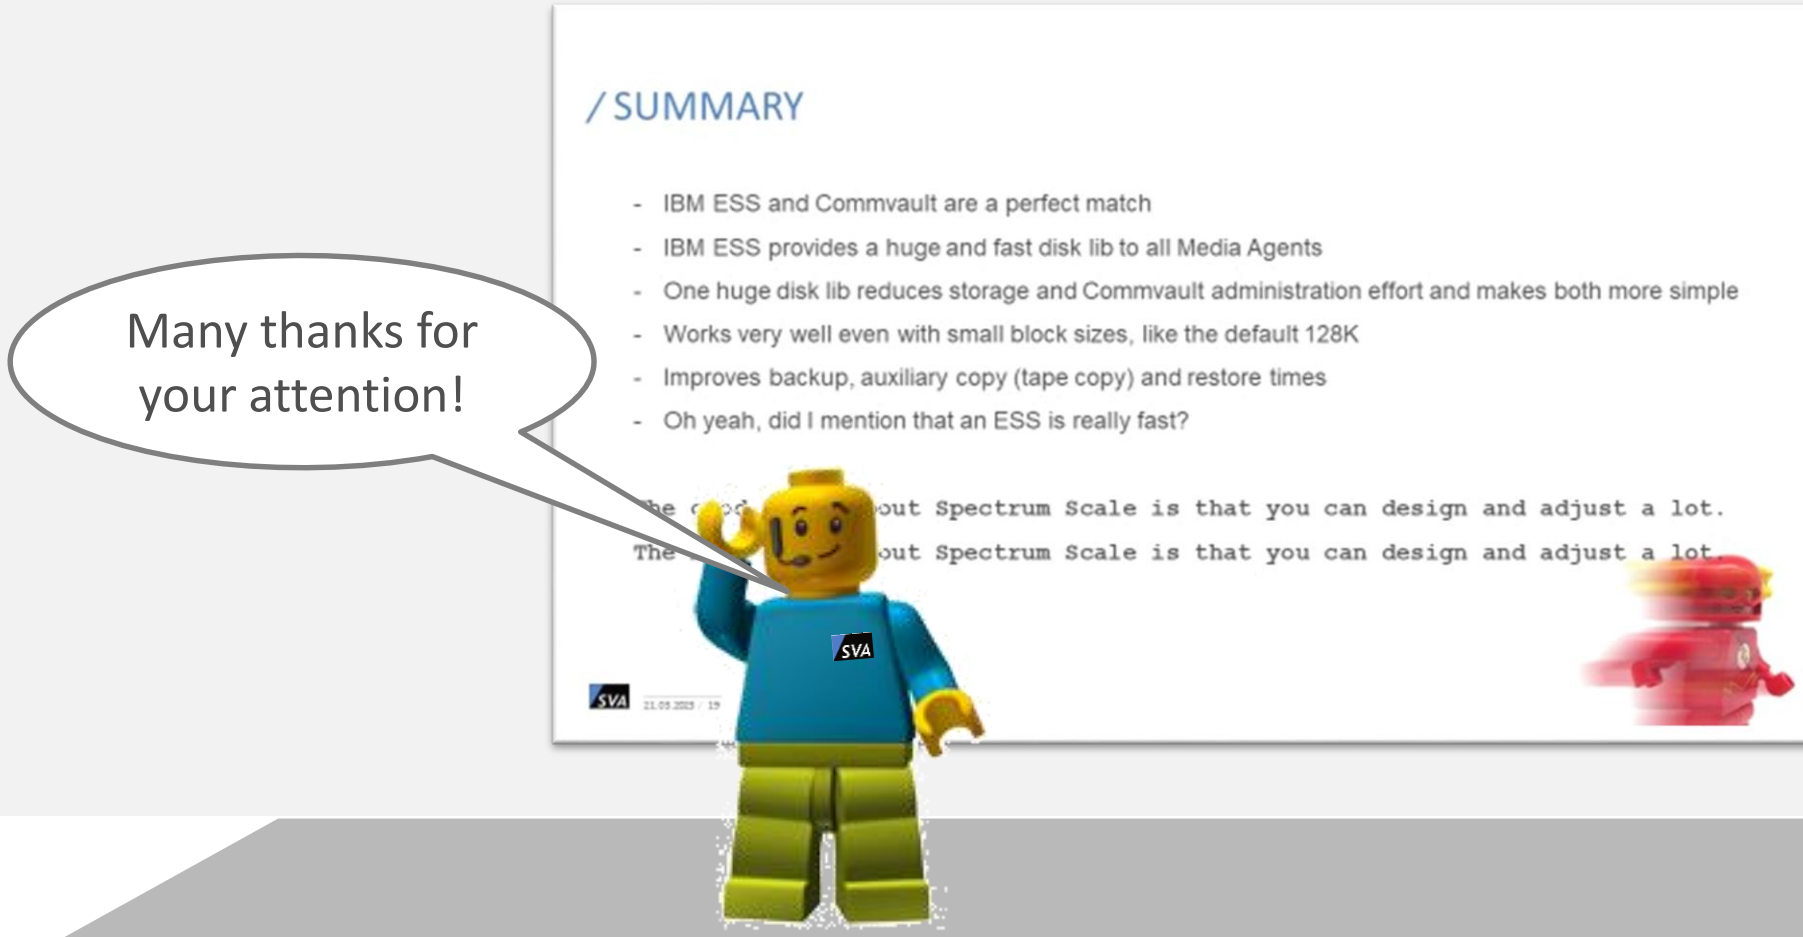
/ SUMMARY
- IBM ESS and Commvault are a perfect match
- IBM ESS provides a huge and fast disk lib to all Media Agents
- One huge disk lib reduces storage and Commvault administration effort and makes both more simple
- Works very well even with small block sizes, like the default 128K
- Improves backup, auxiliary copy (tape copy) and restore times
- Oh yeah, did I mention that an ESS is really fast?
out Spectrum Scale is that you can design and adjust a lot.
out Spectrum Scale is that you can design and adjust a lot.
Many thanks for your attention!

/ CONTACT

JOCHEN ZELLER

IT Architect

Technical Leader IBM Spectrum Scale

Phone.: +49 151 180 256 77

Mail: jochen.zeller@sva.de