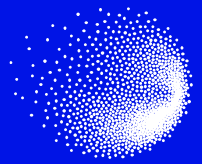

PSI Center for Scientific Computing, Theory and Data

PSI – Site update

Or: Adventures in GPFS land

Storage Scale UG Hamburg - Leonardo Sala
10 June 2025

Who I am, and what this talk is

The Speaker: Leonardo Sala, working as Systems Engineer / Group Leader at PSI since 2012. Particle and theoretical physicist in a previous life

Manages two groups: **Data** Analysis and Research Infrastructure and **Data** Curation -> Data!

First GPFS installed in 2009(?) - can’t remember version, but got serious with it in 2012 onwards (3.5)

This talk: is some rambling aggregate of our experiences (good and bad) during the last years with GPFS. Stress on *rambling*

Who we are


ETH BOARD
ETH zürich
EPFL
PAUL SCHERRER INSTITUT
PSI
WSL
Empa
Materials Science and Technology
eawag
aquatic research
PSI
Villigen
ETH Zürich
Basel
ETH Zürich / ETH-Rat
Zürich
Empa / Eawag
Dübendorf
Empa
St.Gallen
WSL
Birmensdorf
EPFL
Neuenburg
ETH-Rat
Bern
Eawag
Kastanienbaum
Empa
Thun
EPFL
Freiburg
EPFL / WSL
Lausanne
WSL
Davos
EPFL / WSL / Empa
Sitten
WSL
Cadenazzo
EPFL
Genf
ETH Zürich
Lugano
Page 2

Who we are


ETH zürich
EPFL
Swiss Neutron Sources SINQ and UCN
Swiss Light Source SLS
Swiss Free Electron Laser SwissFEL
Proton Therapy
Energy System Integration
Swiss Muon Source SμS
High-precision Particle Physics CHRISP
Hotlab
Radiopharmacy
PSI East
PSI West

What we do

- As part of **Scientific IT**, we do support the scientists at the Large Scale Facilities
- We provide:
  - High Performance Storage (~~GPFS~~ ~~Spectrum~~ Storage Scale) ~ 20 PB
  - File access services (GPFS / SMB / NFS)
  - High performance compute (>100 nodes)
  - High performance network (200G Ethernet / Infiniband)
  - Data analysis services and support
  - Monitoring and reporting
  - A shoulder for crying
- Amazing team of 4 people, spread over three clusters, taking care of everything: Alvise Dorigo, Borys Sharapov, Greta Assmann, Ivano Talamo

What we do

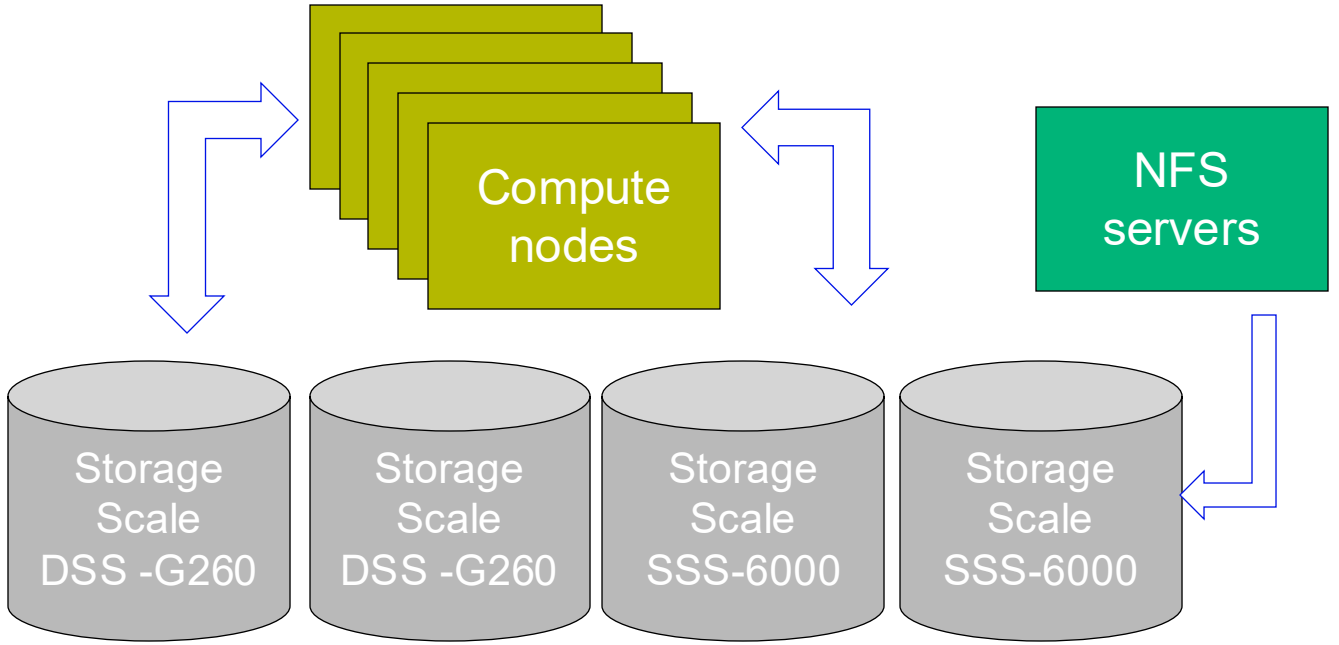
Compute nodes
NFS servers
Storage Scale DSS -G260
Storage Scale DSS -G260
Storage Scale SSS-6000
Storage Scale SSS-6000
Acquisition servers

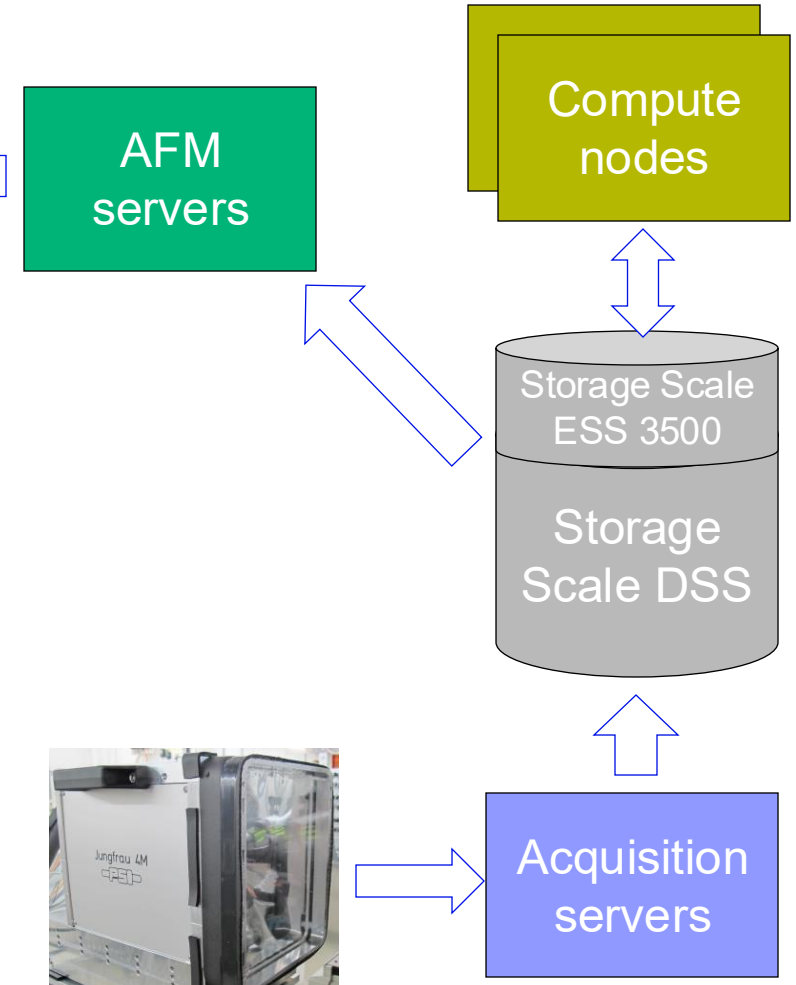
AFM servers
Compute nodes
Storage Scale ESS 3500
Storage Scale DSS
Acquisition servers

What we do – we observe a lot

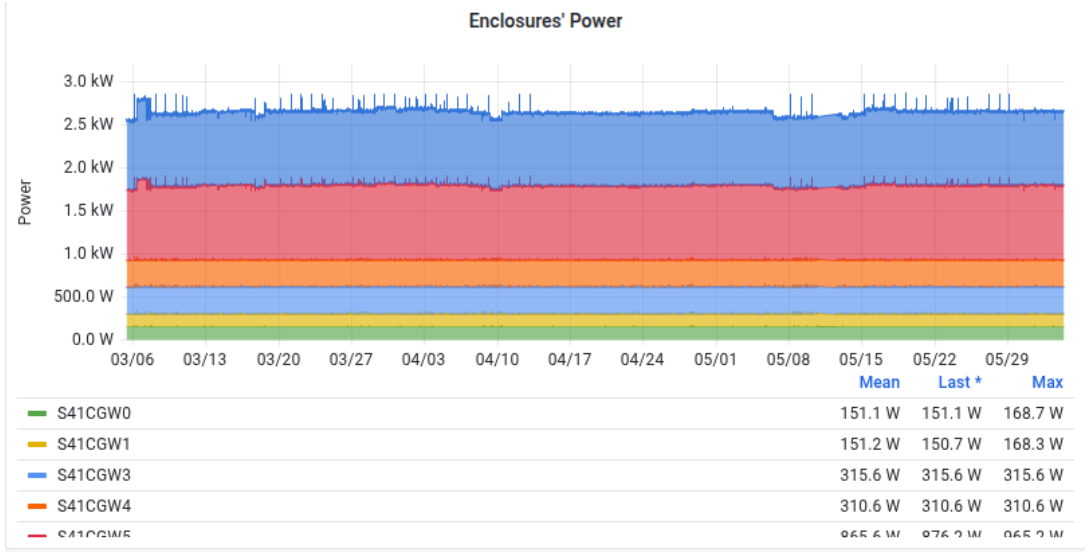

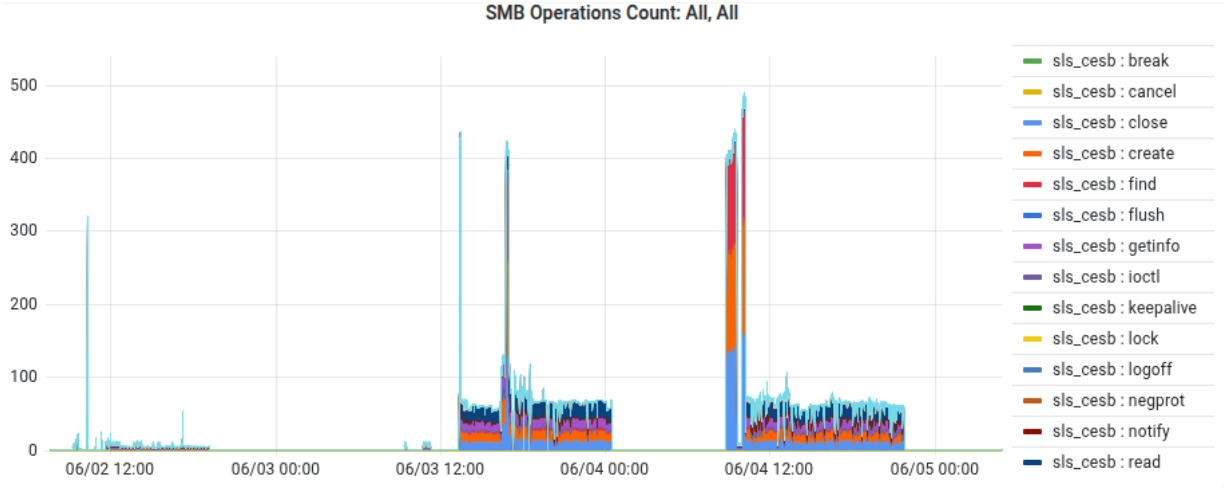

In-house python pmcollector to InfluxDB ingestor (not using Grafana bridge)

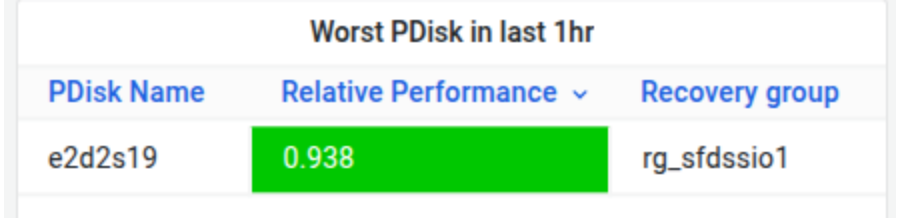

Tale 1: CES everywhere

We need to export **both** NFS and SMB3 (Windows…)

CES might be complex, but works

CES is critical for our users! No CES, no data, tons of angry phone calls at night (true story)

CES is redundant **within the cluster**, but needs downtime to upgrade – this is difficult for us

Solution: 2 clusters (A and B)!

- Playbooks to migrate IPs between clusters
- Some issues with KRB5 NFSv4, as Linux caches at odds with SPNs migration between AD objects
  - Solution: create an AD object containing all SPNs
  - Then merge its keytab with the one automatically created by GPFS
  - Add the keytab to CCR
  - More tests needed, but this black magic seems to work
- This allows also A/B testing with new versions or configuration changes

Tale 2: the Rack - Despair

- **Long story short:** SSS 6000 enclosures were not fitting our racks
- Details
  - It was difficult to get the enclosure **part number**
  - Original documentation reported **no measurements** for enclosure+rails width, got that later from development (info now updated on the website)
  - Our rack was up to standard, conformant to IBM third-party rack specs (specs now changes, good)
  - Enclosure+rails was not fitting by ~1-2 mm
  - We had to modify the racks
  - This was not caught by the preliminary checks done before the installation
- This caused ~1 month delay, and me learning way too many details about racks to prove our case

Tale 2: the Rack - surprises

PSI

Feedback on SSS 6000 (and ESS 3500) systems

- **They work**
- They cannot be customized
  - No personal / central **accounts** allowed, just root
    - Installation and configuration of SSSD resulted in the impossibility to login into the EMS node, even from the console.
  - No integration with site-specific **automation** allowed
  - No easy **integration** in Campus network
  - No **additional software** supported (e.g. monitoring agents, metrics, …)
    - Page 22 of the "IBM Storage Scale System 6.2.2.0": "*IBM Storage Scale System 6.x.x.x uses a new embedded license. It is important to know that installation of any Red Hat packages outside of the deployment upgrade flow is not supported.* " So it seems more of a **licensing issue**
- It is a bit **complex**
  - Matryoshka approach on the support node: bare metal -> VM -> container
  - A lot of **manual configuration from IBM technician** during the installation
- **Feeling: system that can be touched and upgraded only by IBM**

Tale 2: the Rack - spellbound

PSI

IBM design

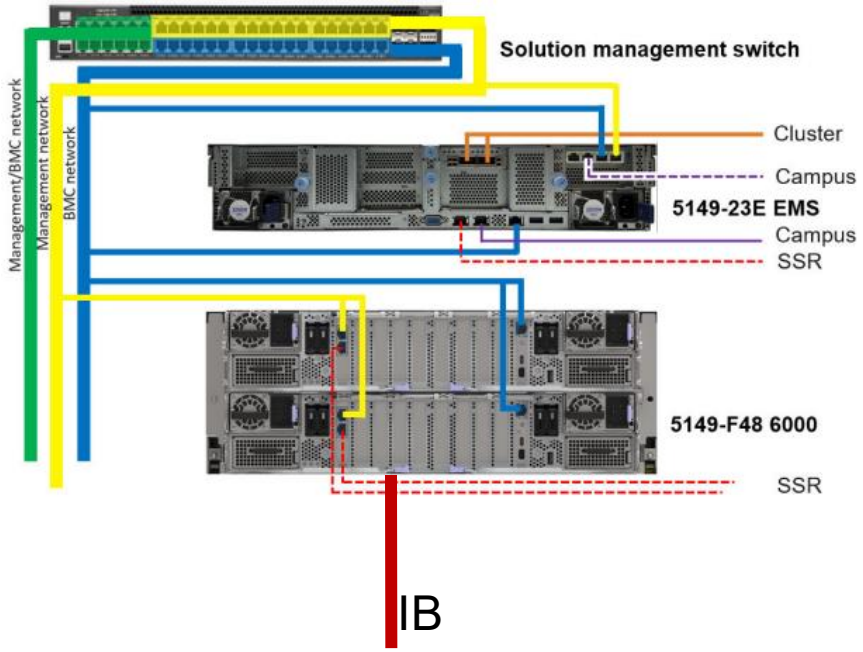

Our world / wishes

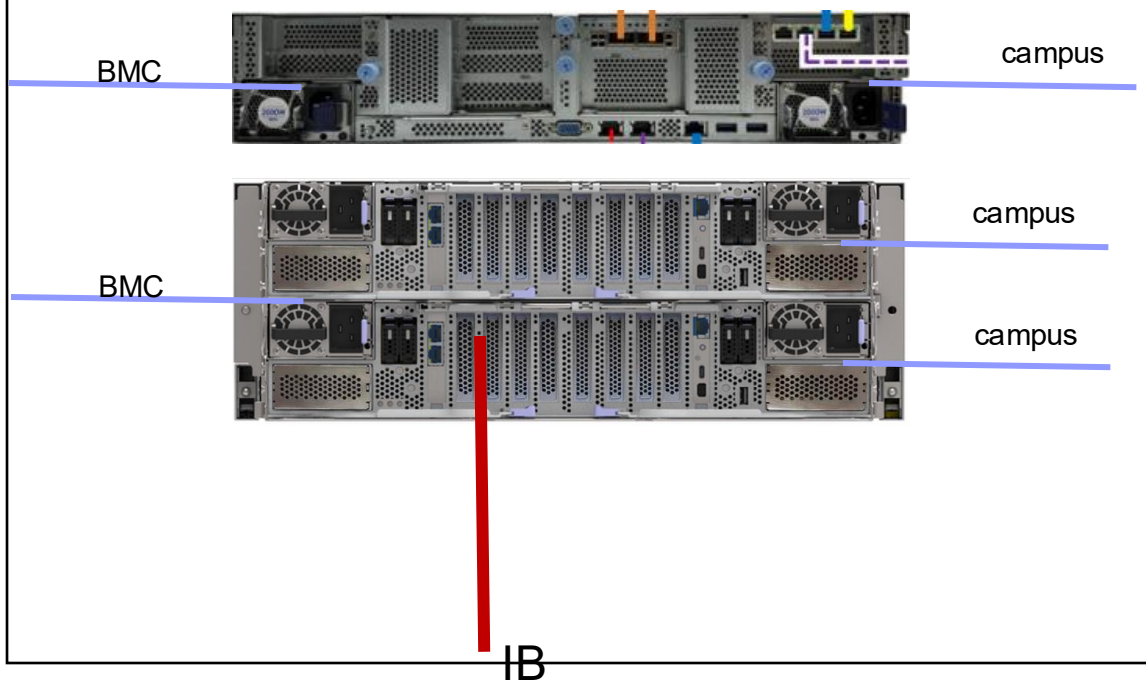

We remote-cluster mount to other networks, need to contact nodes over main IP – are we the only ones?
Ended up using a high-speed connection for campus (which feels a waste)

Tale 3: Joining Kingdoms

Managing both Lenovo and IBM storage in one cluster. Started with smaller "cache" cluster

- ***4.2 PB HDD SAS Lenovo G242 for data blocks***
- ***184 TB SSD SAS Lenovo G242 for metadata***
- ***180 TB NVMe IBM ESS3500 for datablocks fast cache (with policy for flushing to HDD)***
  - effectively helping to reduce the load on spindles (performance boost) in specific periods when online analysis, data taking, and AFM read (to send data to HOME) occur in parallel
  - Merging the two different system was easy (even with many mistakes because of no previous experience on this task) as allowed us to save some good money:
    - Old (pure Lenono) system was still in production, then I had plenty of time to do trial and error on the multi-brand system to merge in un single BB
    - Could start from scratch, so no data to pay attention to (because of previous point)
    - Administering a multi-brand cluster is easy and without nuances in our experience
  - Special care is required when coordinating upgrades

Now merging the main cluster Ra (18 PB spindles)

Tale 4: to above, and beyond!

Or: to physically migrate clusters while in production

Requirements:

- Migrate cluster within 600m – original server room inside major infrastructure restructuring
- With only one new storage system at destination
- Minimize downtime (1-2 days overall, for all compute / storage / network
- Lose no data
- Have a look at middle-eastern situation if time

Original plan:

- Add new storage (Lenovo DSS-G260) as new pool of old fs
- Connect everythi
- Play with replicas
- No downtime, be

MOFED releases differed, nobody could guarantee no data losses or issues
"Cannot see a reason why it should not work, but..."

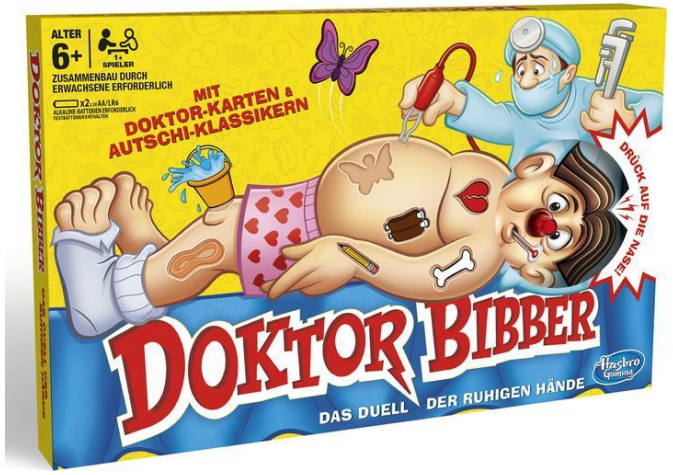

Tale 4: to above, and beyond!

New plan, codename: **Tetris**

- Free up some space, empty up one storage system
- Move storage system to new server room, together with new DSS-G260
- Create a new fs (1 of 3), copy data in background
- Short downtime, final sync and switch mountpoints
- Free up another storage system, move it to new server room
- Add to new cluster
- Copy etc etc

Using fpart, timers, policies, exorcists when needed

Tale 4: to above, and beyond!

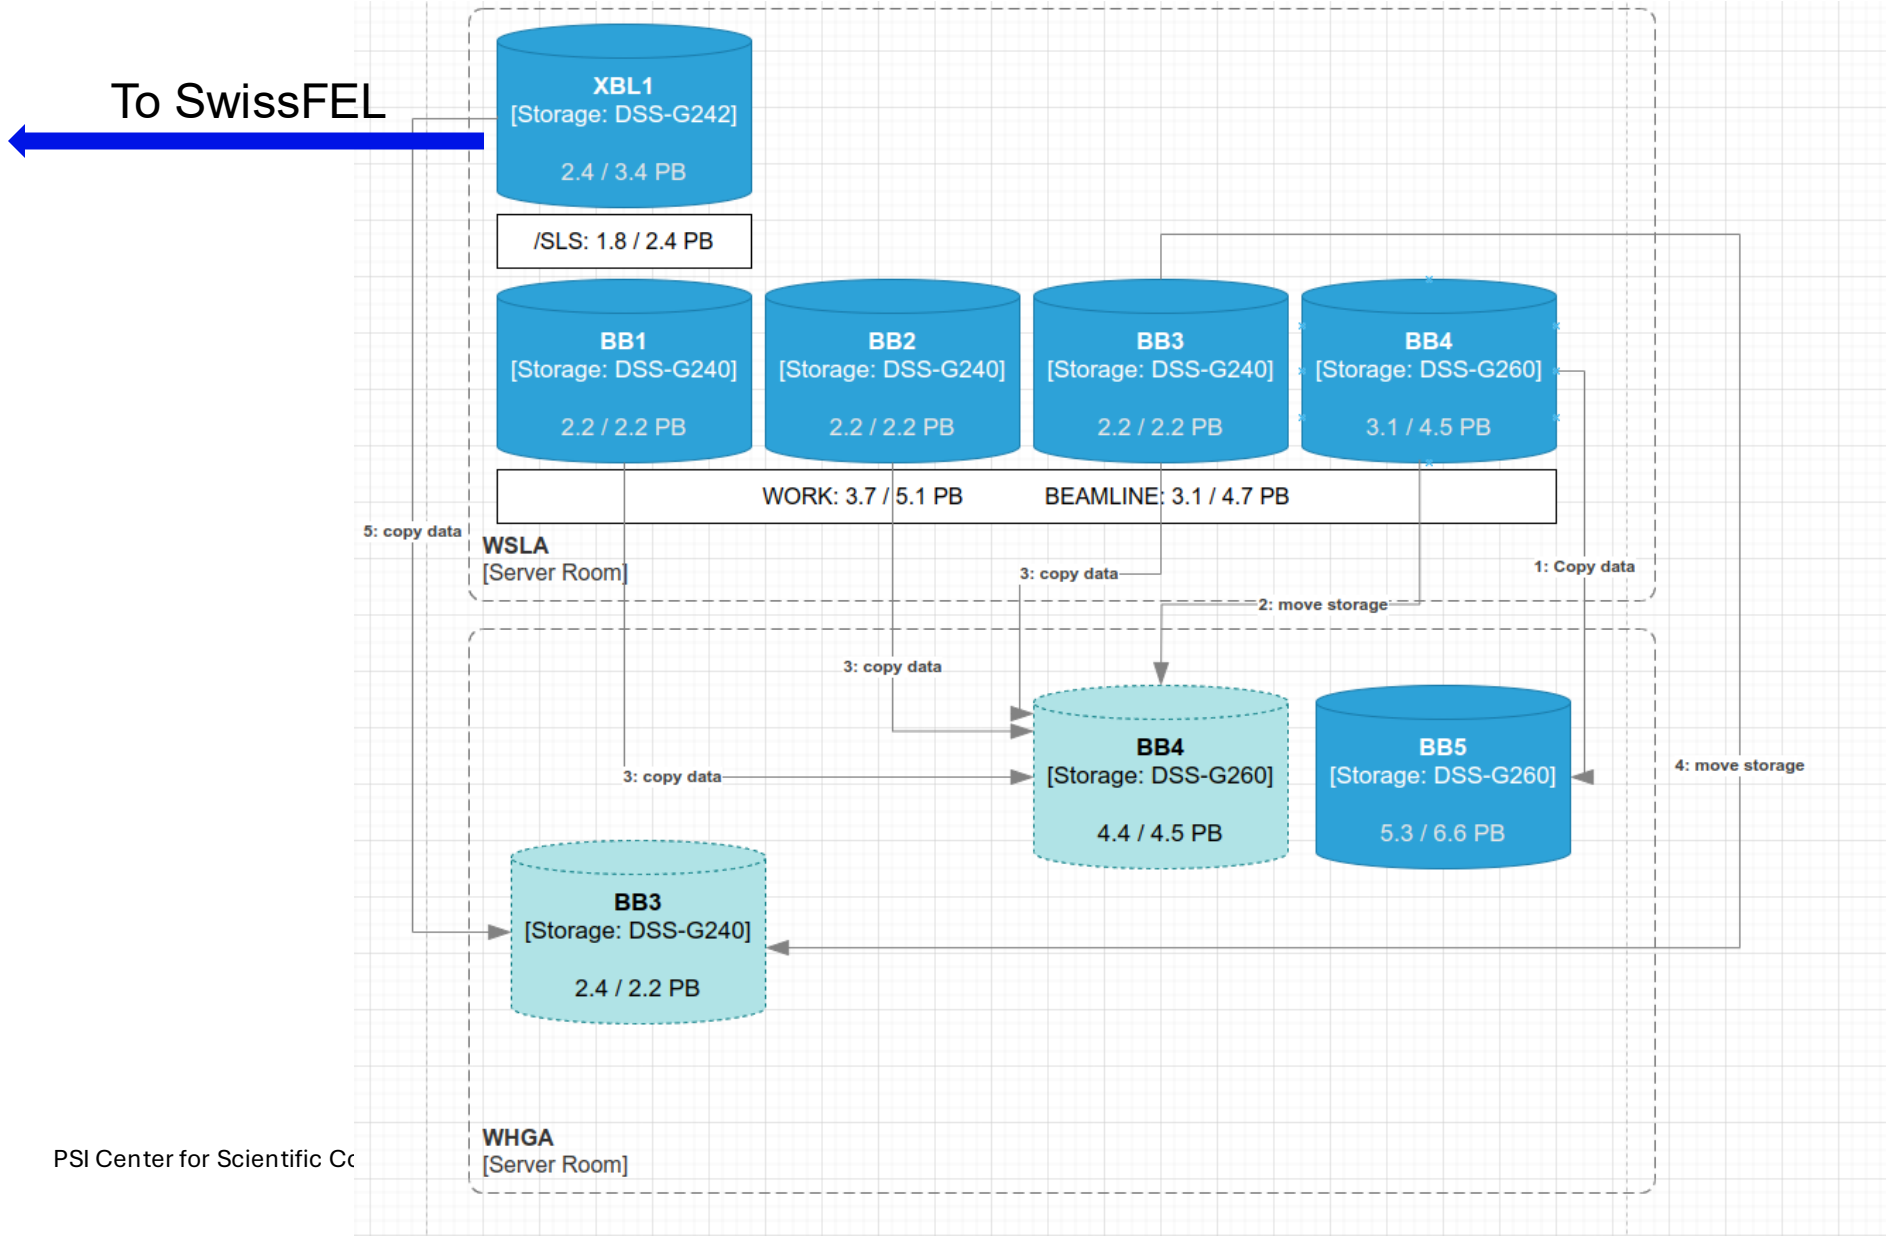

Tale 5: the book of secrets

- Very happy about availability of Storage Scale courses on **Coursera** – we thought this could ease new members onboarding
  - Plus: available with a simple Coursera Plus subscription
  - Plus: **on-demand**, self paced training
  - Plus: hands-on exercises
  - Plus: good lab infrastructure for exercises
  - Minus: it feels an **actor is playing a role**, just reading a script. Happens he says: for the details read the slides
  - Minus: no slides available (***but can download videos***)
  - Minus: some exercises are really manual (see: installation)
  - **General**: very nice start, would like to see something more like the e.g. AWS architect courses
- IBM documentation tends to be very complete, but also a bit sparse in different chapters and verbose

Tale 6: fast and away

AFM is still **critical** to us (see A. Dorigo's talk in 2019 at SSUK)

- 2x Gateways (1CPU Epyc/128GB RAM / node good enough for 2-3GB/s transfer cache > home) - NFS coupling CACHE-HOME
- SW mode: Home is much like an RO backup copy (from which data are in turn archived to TAPE at CSCS in Lugano), but users also use it as offline storage for post-factum data analysis
- No parallel transfer: it was, time ago, not very stable… some WriteSplit every now and then (IBM AFM support very responsive, kind and effective in helping to find a workaround). Now it might be stable but, by now we are happy with parallel mount, which provides the xfer speed we need.
- System can cope with ~3GB/s are no user complained so far (observed peak: 4.5 GB/s)

Tale 6: fast and away

AFM behaved well (and simplified things) when we migrated storage between server rooms (see Table 4). It was as simple as

- Put CACHE in RO
- Resync everything (just in case) to Home
- Dump CACHE's filesets configuration
- Decommission CACHE storage
- Put new storage (IBM/Lenovo) into production
- Create CACHE's filesets and apply saved configuration; Done! Tried several prefetches (around 15-20% of the total amount of data) and no problem at all

Epilogue

PSI

Storage Scale is a critical and important part of our infrastructure:

- performance
- reasonable maintenance effort
- advanced features when needed.

We find it good for our use cases, where customers cannot plan detailed requirements with great notice.

Looking forward its modernization: new API, new CLI tools, rootless mode, fun with NVMe-OF / UStore /…

Thanks for your attention!