DISTRIBUTED SCIENTIFIC COMPUTING: CHALLENGES AND OUTLOOK

OXANA SMIRNOVA (NEIC/LUND)

GRID 2021, DUBNA

A BIT OF BACKGROUND: LHC AND WLCG

- Back in 1998: slow networks and many data centres, hierarchical MONARC model for LHC
  - Grid paradigm is chosen to implement it in 2000 – became WLCG

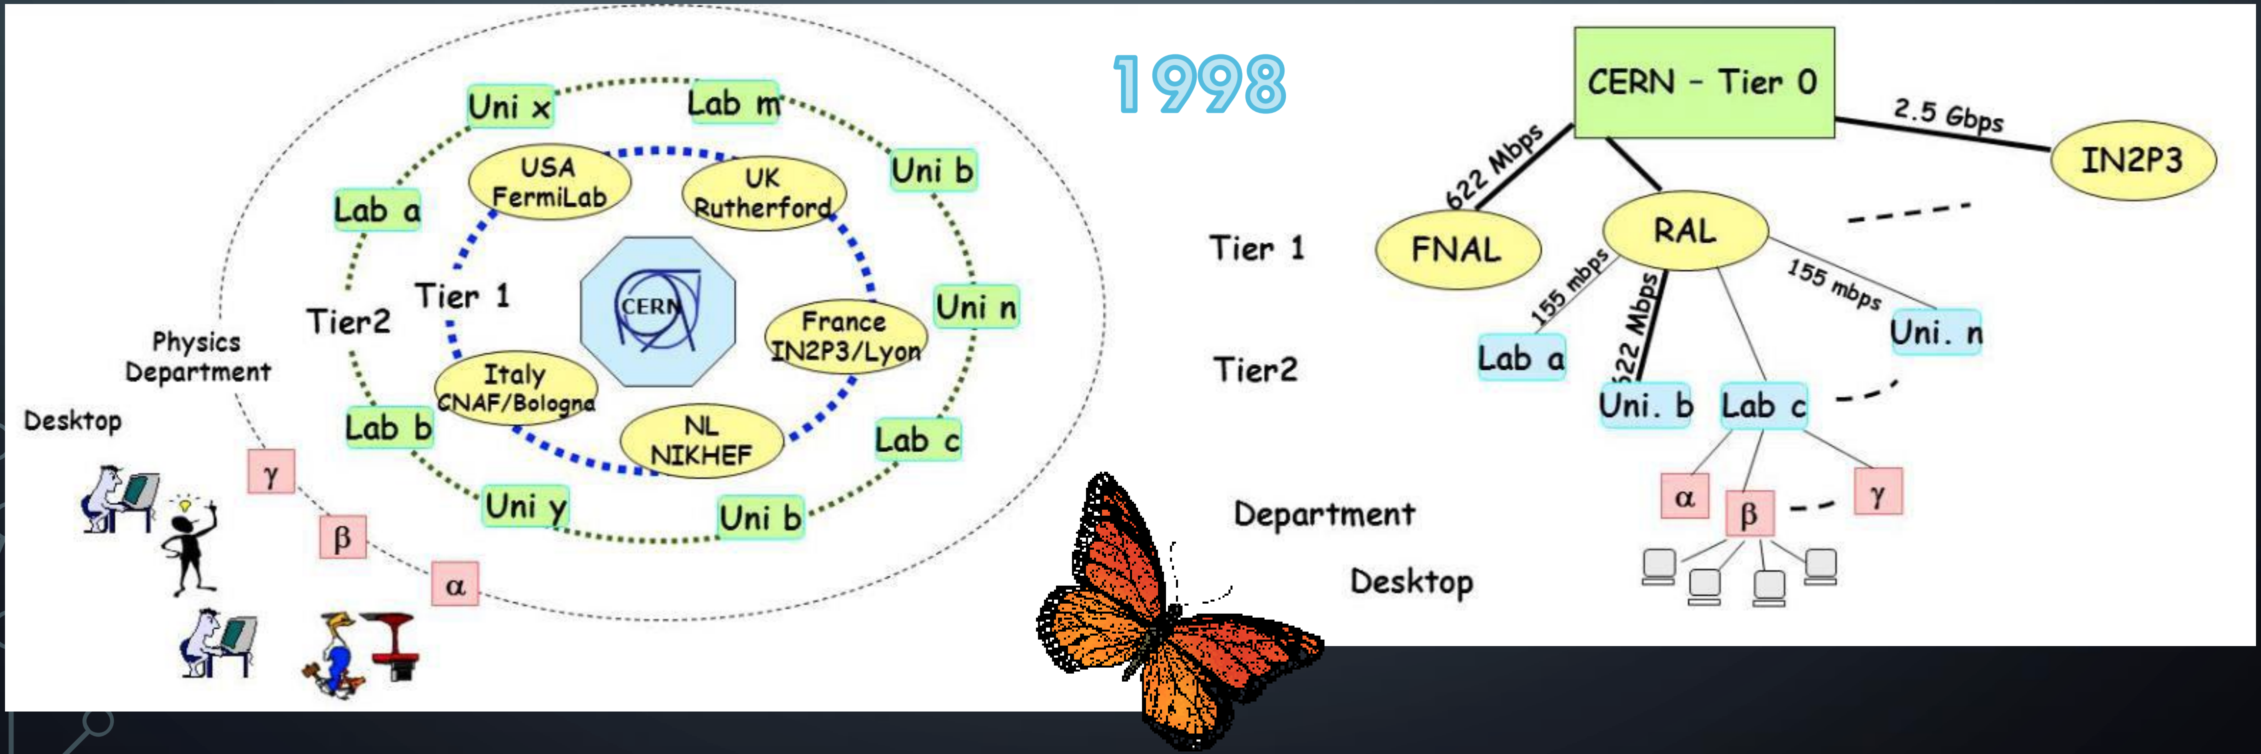

See CHEP18 talk by Ph. Charpentier

A GLIMPSE OF TECHNOLOGIES IN 1998

Best selling cell phone:
Nokia 5110

Best processor: Pentium II

DANTE

Best European research network:
TEN155 (yes, 155 Mbps)

red hat.
OFFICIAL Linux 5.1

Best Linux OS: RedHat 5.1
(kernel 2.0.34)

Top 500: ASCI RED
(First to reach 1 Tflop/s,
9120 cores)

LHC GRID TODAY – A VERY HIGH-LEVEL OVERVIEW

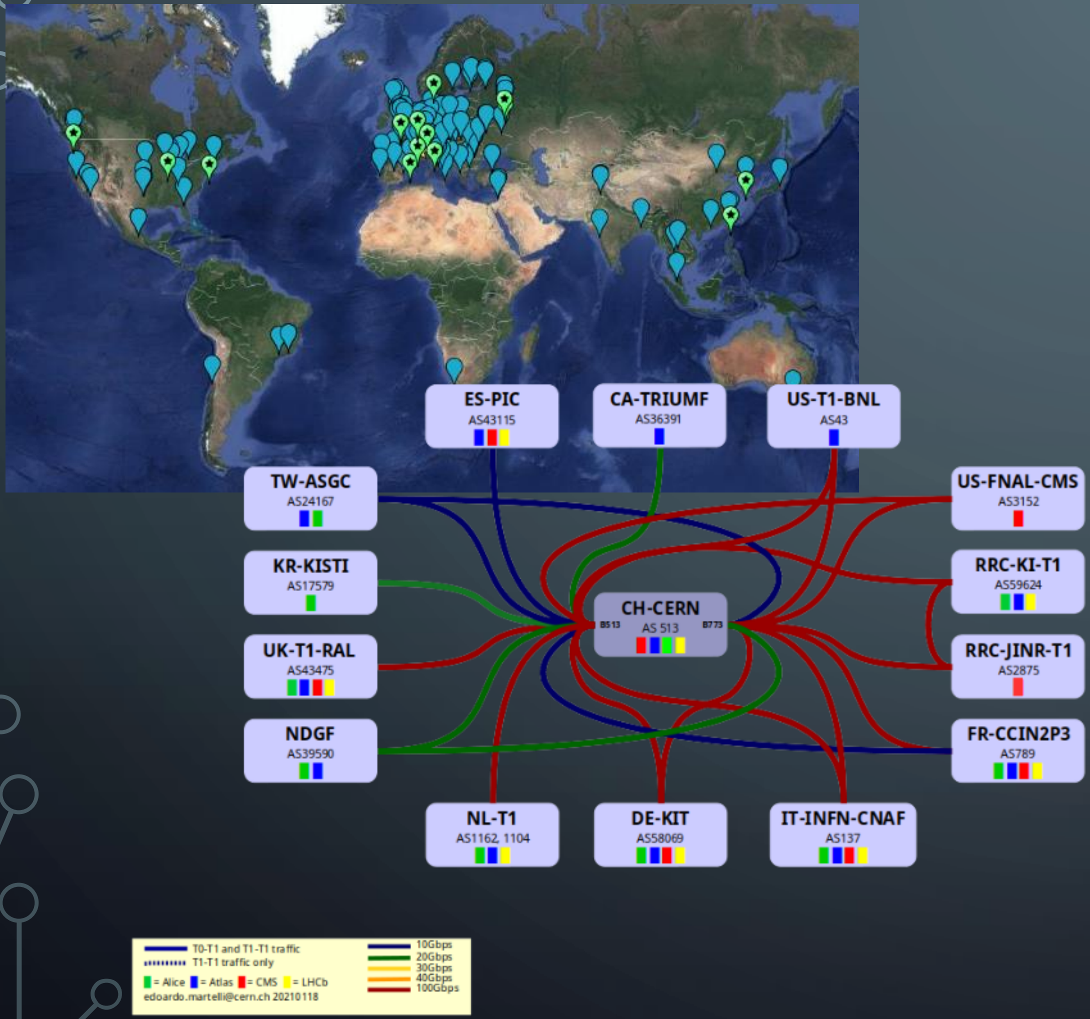

- Tier-0 at CERN
- 14 Tier-1 sites (2 distributed)
- 139 Tier-2 sites
- Up to 100 Gbps connectivity
  - 800 Gbps and 1.6 Tbps expected to become IEEE standards by 2025, pushed by Google, Facebook and such

- Mostly custom-built for LHC needs
  - Usually x86 CPUs
  - Usually CentOS7
  - Usually network connectivity for nodes
  - Usually high data throughput designs
  - Usually late binding scheduling ("pilots")

AN UNUSUAL CASE: OUR DISTRIBUTED NORDIC TIER1

ARC dCache

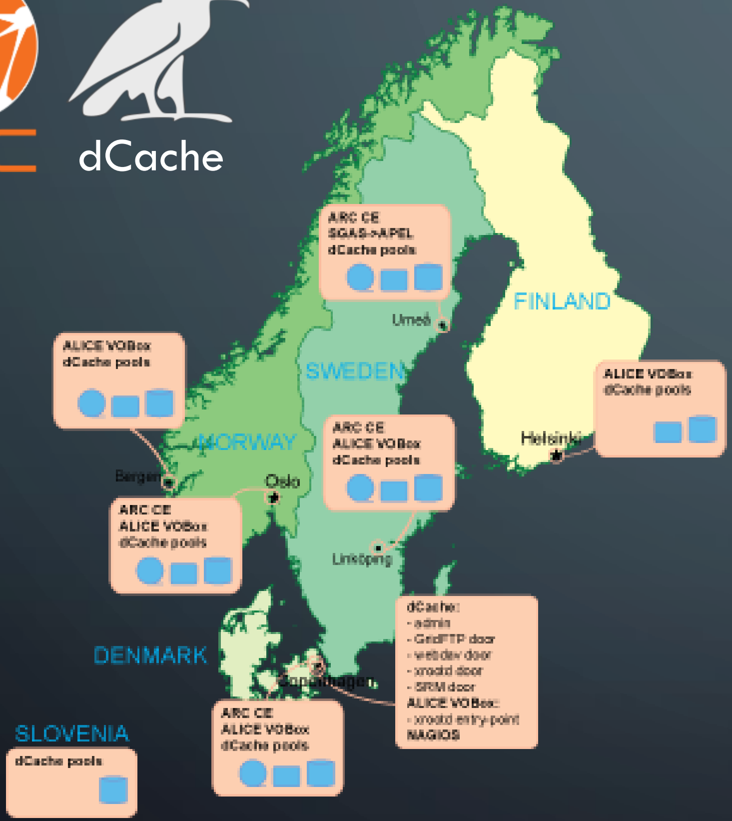

- Inspired by excellent network connectivity back in 2001
- Key requirement: use generic national computing resources
  - HPC, Private Clouds, clusters
  - Must work without outbound node connectivity
  - Must support any OS (even Ubuntu)
- ARC Grid middleware was designed to meet these requirements
  - Built-in data caching
  - Adjustable interfaces to batch systems
  - Minimal external dependencies
- dCache storage solution was extended to handle remote pools
  - Probably the first "Data Lake"

WHAT ARE THE TRENDS, 20 YEARS ON?

CHALLENGES THAT REMAIN

- Data volumes keep growing
  - Requires more simulation, too
- Data structure is unchanged
  - ROOT files
- Data analysis and simulation are largely unchanged
  - Trivial parallelism, easy to distribute
  - Limited need/support for true parallelism
- LHC still needs external computing and storage capacity
  - Also end-user analysis facilities

NEW TRENDS

- Hopes of wide take-up of Grid did not materialise
  - LHC is still a corner-case
- Processors did not become faster
  - Became more parallel instead
- Storage did not become cheaper
  - Became faster instead
- Big Data analytics methods are increasingly popular
  - Both in analysis and simulation
- Containerisation is commonplace
- Economy of scale replaces specialised solutions
- Programming skills are on decline
  - Pre-made frameworks are needed

TRENDS IN COMPUTING

10 years ago

HPC is old school

Cloud is the future

Today

Cloud is a commodity

HPC is luxury

Where are we?

- The core is still general purpose Linux clusters
  - Some with HPC capabilities
- Clouds are relatively easy to integrate
  - IaaS, private or public
  - Data movement can be an issue
- Public funding however goes to HPC
  - Shift from opportunistic usage to regular allocations

FOCUS ON HPC

- Supercomputers have been prolific data producers for decades
  - Various supercomputers in the USA, PRACE in Europe etc
- Strategic worldwide trend: massive investments into **Exascale** supercomputers
  - Leadership facilities in the USA, EuroHPC in Europe
- For most sciences, HPC offers best cost/performance balance
  - Rather than getting funding for hardware, scientists receive time allocations
  - Project of any size can be accommodated within a single Exascale HPC facility – no need to distribute the load
  - Energy efficiency
  - Human effort efficiency

SPREADING HPC KNOW-HOW

- Experience of the Nordic Tier1 and ARC approach became helpful in adding more HPC facilities to the WLCG
- Some recent examples:

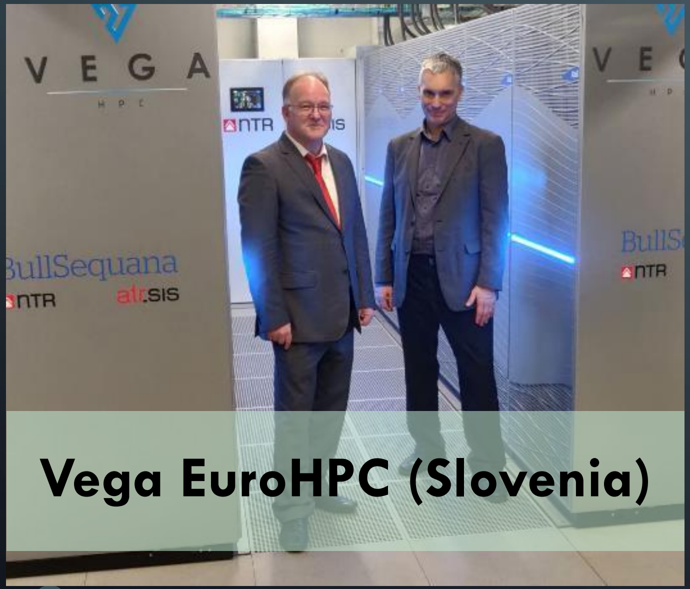

Vega EuroHPC (Slovenia)


Mare Nostrum (Spain)


Czech National HPC centre IT4Innovations

SUCCESS STORIES

Vega given exclusively to ATLAS – provided more cycles than the total pledge

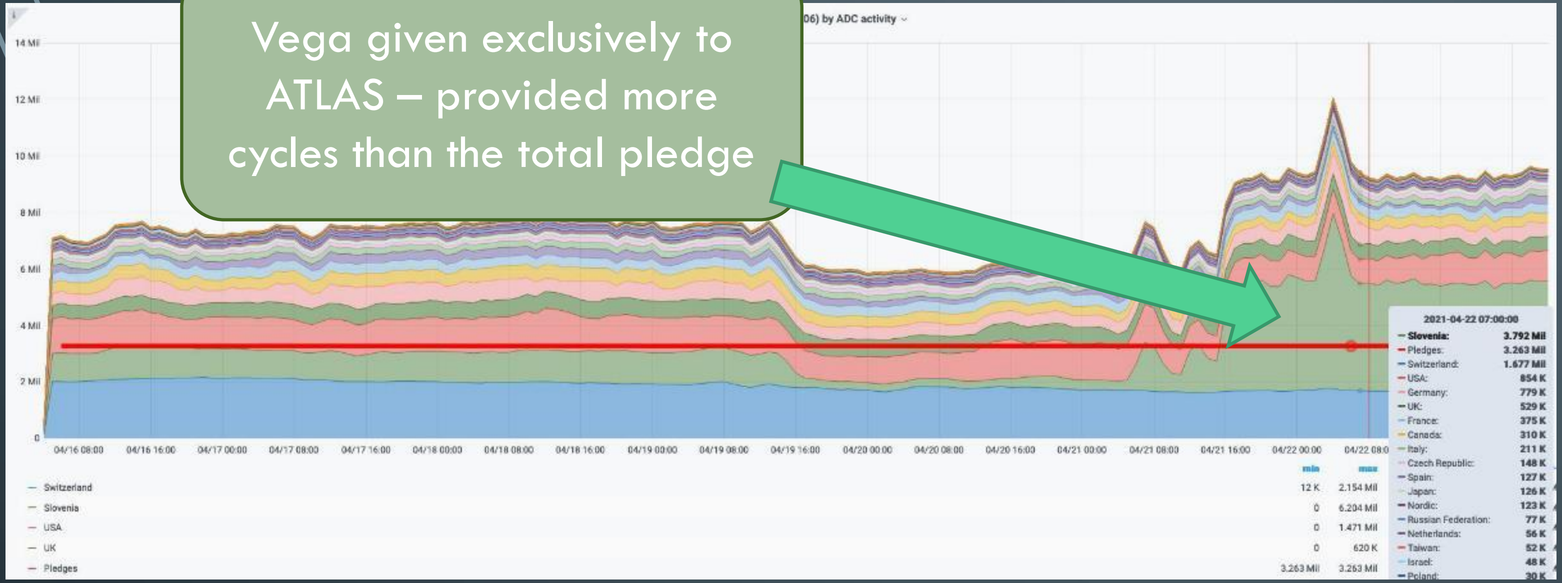

HPC share in Czech contribution

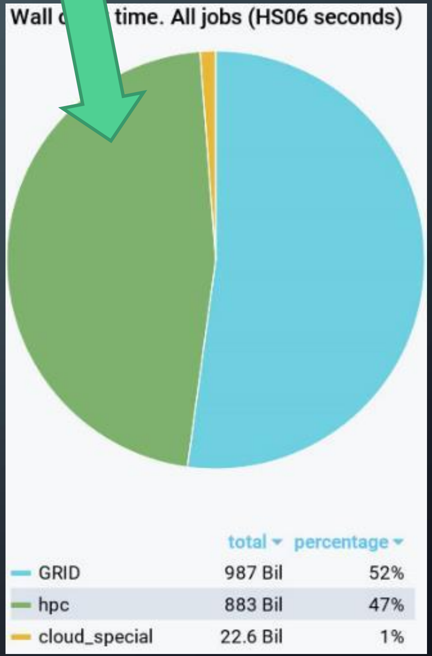

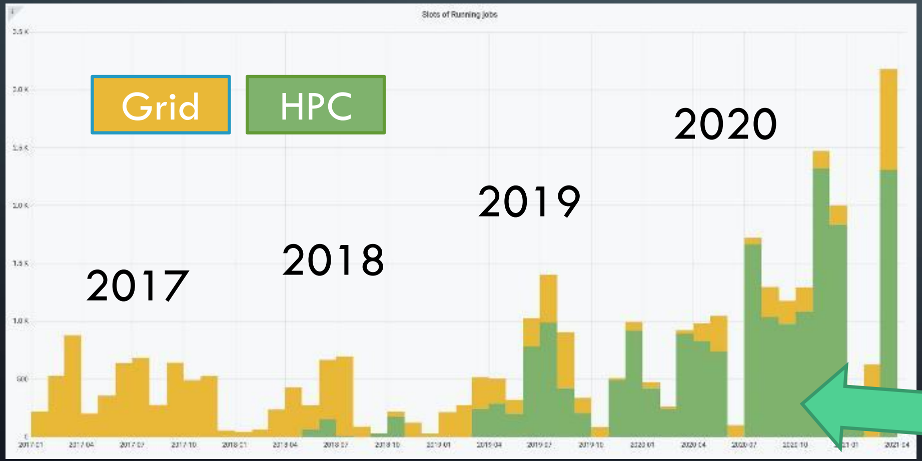

Mare Nostrum share in Spanish contribution

Graphs from talks by A. Filipcic at NeIC NT1-AHM21 and A. Pacheco Pages and M. Svatos at vCHEP21

ARE WE READY TO COMPLETELY SWITCH TO HPC?

- No. By far most tasks in these examples are still simulation jobs
- The key difference remains: HPC facilities produce data, while many sciences also need to process/reduce data
  - A conventional HPC facility is not designed for high throughput tasks
  - A conventional LHC workflow is not designed for parallel processing
- This can change, as Nordic Tier1 and Vega show
  - Painstakingly designed to meet both HPC and LHC computing needs
- Need outbound connectivity from jobs to access frequently changing detector conditions databases
- For high-throughput reproducible data analysis we need massive storage
  - This includes archiving with very active tape usage

WHAT DATA ARE THERE?

- LHC raw data are, and likely will stay, the same: hierarchical, in ROOT files
  - They are a lot, but not classical "Big Data"
    - Compare with e.g. social networks: texts, images, audio, video, search histories etc
- Research data in general are very diverse
  - Different origin
    - Measured once (e.g. temperature, humidity, salinity) and non-reproducible
    - Produced (in controlled experiments, like in accelerators) and reproducible
    - Simulated (particle and nuclear physics, astrophysics, molecular dynamics, protein folding, climate models, markets etc)
  - Different size
    - And different storage solutions
  - Different structure
  - Different data taking rates
  - Different analysis workflows (traditional statistics vs Big Data analytics)

LHC AS DATA GENERATOR

- LHC is arguably the greatest controlled experiment ever
  - It produces much more data than modern IT can handle
  - It is IT limitations that force physicists to throw away “uninteresting” data
- Technically, all LHC data can be re-produced – at LHC
  - It is still cheaper and more energy-efficient to store them
- Long-term preservation of exabytes of raw data is a unique need of LHC
  - Radioastronomy is similar, only their data are non-reproducible even
- Data storage problem will become even more acute at HL-LHC
  - While processor cycles can be borrowed and re-used, storage can not
  - Data storage may even become the single limiting factor for the future LHC research
- The solution? None so far, we have to buy more storage.

TRENDS IN SCIENTIFIC DATA ANALYSIS

- End-users tend to work with derived/reduced data, accessible through various platforms
  - A very diverse picture, with clouds and analytics becoming commonplace technologies

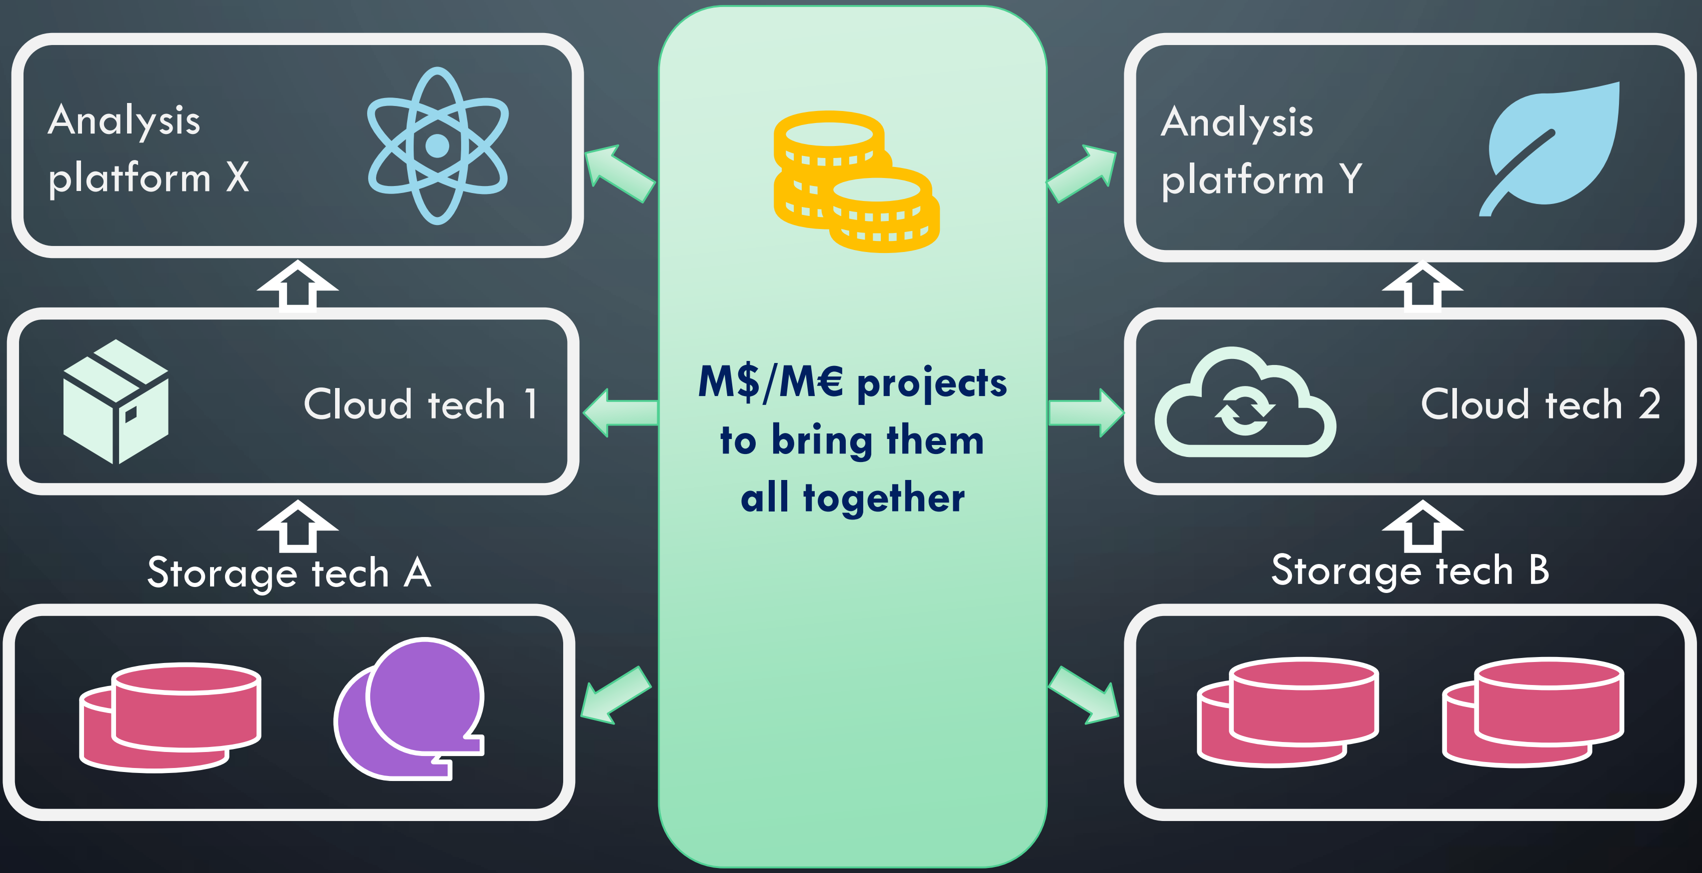

EMERGING POLICIES AND REGULATIONS

- The initial Wild West in digital data handling is getting more regulated

  - Both legally and through the market
- Tailored approaches are not sustainable, replaced with industry standards
  - GridFTP is obsoleted by HTTP(s)
  - X509 proxies give way to OIDC tokens
  - Scientific Linux OS is gone
- There are laws and policies to respect
  - GDPR seemed to be irrelevant to LHC, but data about scientists are also data
  - FAIR is targeting much smaller data, but we have to comply, too

EMERGING TECHNOLOGIES

A personal view on the current order of priorities

Hardware acceleration: GPGPUs, FPGAs, ASICs etc

- Not really new, but no established coding standards yet: CUDA, ROCm/HIP, oneAPI/SYCL; wait for C++20 and beyond
- For LHC, human effort needed for code re-writing needs to be compared to benefits

Low-energy architectures in supercomputing, e.g. ARM

- In Top500 now

In-memory processing

- Not new either, but still not common in LHC workflows

Quantum computing

- May sound like a SciFi still, but it would be sad to miss that train

SUMMARY

- Future of computing: large service providers, decline of small clusters
  - Service providers: clouds and supercomputers, depending on tasks
  - Personal computers are good enough for many smaller tasks – if one can handle data and software
- Future of data: still uncertain
  - No standardisation effort took wing
  - Long-term data preservation is rather expensive
    - Big Data of social networks generate profit, scientific data don’t

- Future of technologies: industry-driven
  - “Classical” CPUs will soon become a rarity
  - Artificial intelligence will be applied everywhere, from simulation to data acquisition to analysis
  - Energy efficiency will gain even more importance