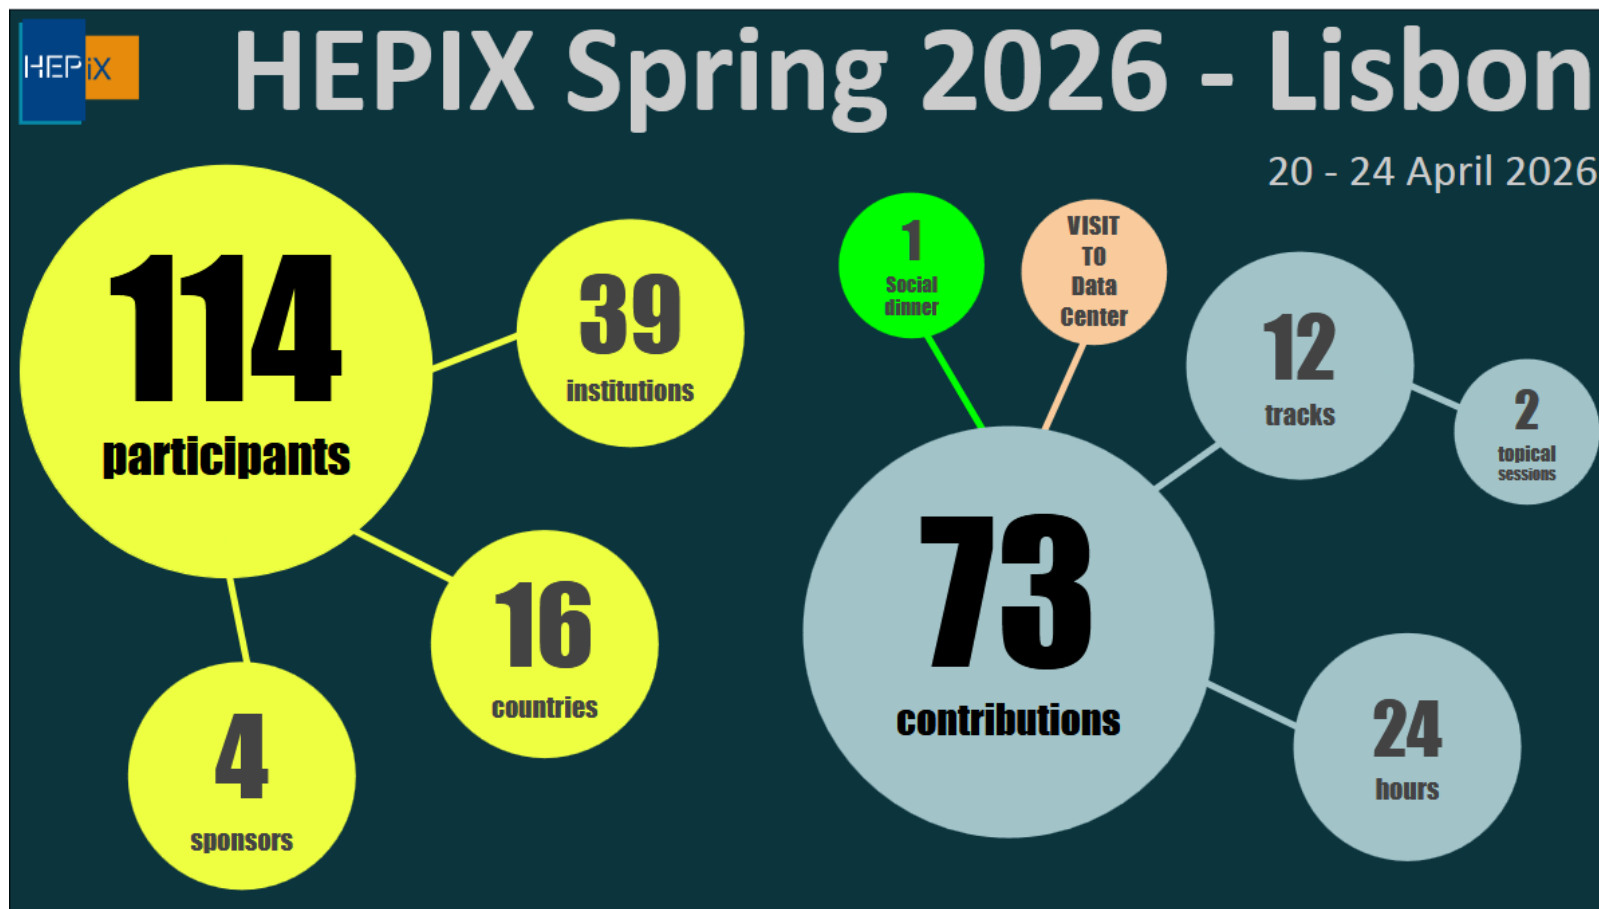


HEPiX Report


30 april 2026

Workshop in numbers



Workshop in numbers

Count of Country



~40 institutions
16 countries

Most of the participants European
7 participants from North America
22% participation from Asia

Next meeting



- Following:
 - spring 2027 – Germany
 - fall 2027 – uncertain
 - spring 2028 @Biodola

Site report: CERN

- InfluxDB deprecated in favour of the Prometheus ecosystem.
- Metric-oriented workloads moved to Mimir.
- General time-series or analytical data migrated to TimescaleDB.
- Migrating from ISC DHCP servers to the modern Kea DHCP standard; expected to finish by Q3 2026.
- AlmaLinux 10 and RHEL 10 released at CERN in October 2025 – migrations planned during LS3

Site report: RAL

- “Dell-only” site
 - New WNs: 160C, 1.5TB, 8x3.2TB nvme, 100Gbps
 - GPU nodes: 2x RTX PRO 6000
- Opportunistic use of cloud GPUs in farm
- Migration from vmware to proxmox
- Upgrade of the 16 Super Spine router hardware from 32 port 100Gbps SN2700 routers to new 64 port 400Gbps Dell Z9664 routers
- Oracle SL8500 Tape library almost very nearly at end of service
 - water leaking from offices
- Very interesting presentation on CEPH update

Virtualization

- BNL migrated from RHEV to OpenShift
 - Several issues due to new platform, slowly solved
 - Direct migration but with a reboot
- NIKHEF migrated from VMWare to Xen commercial version
- Proxmox not presented by any site (but many are adopting it)

Benchmarking

- Ecosystem: Workloads, HEP Score, Suite
- HEP GPU: workloads available in containers
 - GPUs are used in a variable way compared to CPU
- Power consumption
 - Figures start emerging in the DB
 - DGX Spark rules the list
- Still much work to do
 - in the end the most useful thing will be the DB

Power efficiency monitoring

- Enable reliable data-driven power predictions that move beyond static defaults
 - provide transparent, job-specific energy consumption metrics for all WLCG compute resources
- Leverage the HEP Benchmark Suite and HammerCloud to generate a rich dataset of electric power, performance and utilization metrics
 - detailed profiling and modeling of compute resources
- Steps
 - Increase site adoption of power collectors
 - Improving fitted models and data validation phase
 - Automated report generation
 - Anomaly detection to streamline site operations

Tape: Fuji

- Increase of tape demand up to 40% until 27
- LTO-10 → 30/40TB - 2025
- LTO-11 → 70TB, LTO-12 → 120TB
- BaFe → SrFe
 - grants up to 50 years
 - Reliability is also increased

Tape: client pow

- Market sees only a few brands involved
- Capacity will grow faster than throughput
 - Tapes are for capacity, if you look for throughput you will buy something else (e.g. disks)
- Trend is from scale-up to scale-out libraries
- HW interface: from FC to SAS
 - Cheaper but need to place server closer
- As everything else, drives are going up (30%)

Tech watch - CPUs

- AMD
 - zen6 (Venice) up to 256 p-cores
 - 3D L3 cache in Venice-x (rumor)
 - Zen7 (Verano)
- Intel
 - Xeon 6+ up to 288 e-cores
 - No GPUs expected soon
- ARM – several announces
 - Nvidia Vera 2x performance over Grace
 - CPU from ARM itself (AGI CPU)
 - Ampere Altra up to 128 cores
 - Fujitsu Monaka in 2027

Tech watch - GPUs

- AMD
 - Different models keener on HPC-oriented apps than NVIDIA
- NVIDIA
 - RTX 6000 better suited for HPC (PCI slot)
 - Trending towards rack-scale integration
 - Rubin will increase low prec performance

Tech watch – memory and storage

- DRAM
 - DDR6 only in 2029
 - GDDR7 low cost alternative to HBM for GPUs
- NAND Flash
 - 321-layers QLC, 500 expected by 2027
 - SSDs all sold out up to 2027
- Storage
 - 40TB (WD) and 44TB (Seagate) start to be available
 - Interesting technologies emerging that can speed up standard disk speed

Tech watch: Cooling

- 1 phase DLC – the standard one
 - Every vendor has it in portfolio
- 2 phase DLC – uses boiling
 - Not every vendor has it
- Immersive cooling
 - Still an option
- Vendors: CPU 500W, GPU pci-e 600W, GPU sxm 1400W, worse for the future

Tech watch: cooling at nikhef

- 1P-DLC not suitable for many reasons
- 2P-DLC tested with different server vendors
 - Zutacore retrofit
 - Showed interesting results with 2 different bios settings

WLCG heter. Computing summary

- No experiment is ready to formally request or accept pledges of GPUs: still only R&D
- Requesting GPUs in only one side – the other is provisioning
- LHC exp. Computing models leverage het. Resources on run 4, but no significant speed up has been noticed
- Lot of work to do with respect to information systems and accounting
- Relative costs of GPU vs CPU processing capacity will be essential in making decision on offloading

GPUs at KIT

- Set up a cluster with several NVIDIA GPUs, provided via the gridka cloud
- Implemented MIG to virtualize the GPU
- Very big GPU cluster provided to users, all NVIDIA GPUs.
- Users require only NVIDIA GPUs, we are risking a lot on this side.

CERN procurement

- Procuring at CERN is getting hard, like for every other WLCG site, due to market uncertainty.
- We as a community must try and share experiences, validation processes and every info that could help others

HTCondor workshop summary

- New features
 - New authenticated docker support
 - Support for het. GPUs in one server
 - No job networking
- Next meeting at Lyon 29-9 – 2-10
 - All local organization covered by IN2P3

AUDITOR

- The need to validate collected data from this tool is emerging
 - Have chance to verify them
 - Mainly for CO2 data

Secrets at CERN

- Old tool was discontinued: selected hashicorp Vault as the new alternative
 - OpenBAO interesting alternative but too new
- Integration with puppet, Kerberos
 - Enabled several services (e.g. ssh, k8s)

New track on AI

- A Centralized LLM and Agentic AI Infrastructure for CERN
 - A new, scalable ML-service designed to securely support LLMs and Agentic AI by combining an OpenAI-compatible proxy, a hybrid model catalog, and an agent orchestration platform to balance innovative research with strict operational security
- Integrating Spack, MLflow, and AgenticAI at DESY (Maxwell Cluster)
 - A unified ecosystem combining Spack, MLflow, and AgenticAI to automate complex HPC operations, such as SLURM job management, while significantly advancing the reproducibility, performance, and manageability of scientific workflows