 https://doi.org/10.1051/epjconf/202533701326

Serving Photon Science and HEP at the same facility

Christoph Beyer[1], , *Stefan* Dietrich[1], , *Martin* Flemming[1], , *Martin* Gasthuber[1], , *Jürgen* Hannappel[1], , *Thomas* Hartmann[1], , *Yves* Kemp[1], , *Philipp* Neumann[1], , *Frank* Schluenzen[1], , *Sven* Sternberger[1], , and *Christian* Voß[1]

[1]Deutsches Elektronen-Synchrotron DESY

Abstract. DESY (Deutsches Elektronen-Synchrotron) has been a center of German high energy physics (HEP) research since 1959 with on-site experiments and theoretical research. Since the turn of the century, the research scope has evolved beyond HEP. With the PETRA-III or FLASH facilities on site and the collaboration within the European XFEL, DESY has become a major photon science laboratory as well. To be able to serve the computing and data taking needs for these various user groups and scientific backgrounds, DESY IT has developed and deployed data acquisition, storage and computing solutions.

1 Introduction: DESY as an interdisciplinary laboratory

DESY is a major research laboratory in Germany, dedicated to fundamental science. Particle accelerators were its founding cornerstone and remain essential tools for DESY scientists. Accordingly, DESY develops and operates accelerators, primarily for research involving photons and high-energy particle physics. It also pursues astroparticle physics and other scientific fields loosely linked to accelerator use. Below is a brief overview of these activities.

1.1 Developing and operating accelerators

DESY operates multiple accelerators of various sizes. The principal large systems are the PETRA III electron storage ring [1], serving as a synchrotron radiation source, and the European XFEL free-electron laser (EuXFEL) [2]. DESY also runs FLASH [3], the world's first XUV and soft X-ray free-electron laser, as well as smaller pre-accelerators and R&D accelerators. Ongoing work includes R&D for new accelerators (e.g., PETRA IV) and advanced techniques (e.g., plasma wakefield).
These accelerators are mainly operated by DESY staff, with some long-term collaborations involving other institutes or universities.
IT requirements are shared between central DESY IT and the machine groups, using both Linux and Windows systems. New facilities feature increasing numbers of sensors and actuators, generating growing volumes of telemetry data. Research and development for novel accelerators rely heavily on simulation, typically conducted on HPC systems. Both telemetry data and simulation results can also be preserved in long-term archives.

1.2 Research with photons

DESY provides more than 40 experiment stations at PETRA III and FLASH, accessible via a proposal system to scientists from DESY and worldwide. EuXFEL likewise hosts numerous users generating data at its facilities, which are largely collected and analyzed at DESY.
DESY IT collaborates with experts from PETRA III, FLASH, and EuXFEL to establish data acquisition, storage, and analysis infrastructures. Experiment teams generally run the data acquisition systems, while DESY IT integrates data transfer into the analysis chain.
DESY and EuXFEL staff, along with visiting scientists, use DESY IT-managed systems for simulations, data management, and analyses. Increasing online processing demands are met by appropriately configured analysis setups, and raw data are also archived for future use.

1.3 High energy physics

DESY participates in the LHC experiments ATLAS and CMS, the BelleII experiment at KEK, and smaller or emerging projects such as ALPS II or the ILC. Some legacy analyses still involve HERA data.
DESY IT has long provided computing resources, support, and advice for the DESY HEP community. Beyond serving local DESY staff, DESY IT contributes significantly to national and international computing efforts for these experiments. Both Zeuthen and Hamburg sites operate large grid computing and storage clusters that play substantial roles in WLCG and Belle II production. In addition, DESY hosts the National Analysis Facility (NAF) for fast user analyses and offers raw and simulation data archival on tape. Services extend beyond CPU and storage, supporting both on-site and external collaborators.

1.4 Other Activities

While accelerator-based research is centered at the DESY Hamburg site, DESY Zeuthen is heavily involved in astroparticle physics. This includes neutrino astronomy with IceCube, IceTop, Amanda, and Baikal, and gamma-ray astronomy with H.E.S.S., Magic, Veritas, CTA, and Fermi. Some of these projects have significant computing and storage needs, which DESY supports.
DESY's theory groups drive innovations in all these domains. DESY IT at both Hamburg and Zeuthen actively assists with simulation tasks and computational physics across the lab.

2 DESY involvement in HEP computing

DESY has a long history in particle physics, or high energy physics. This started with DESY and DORIS, and continued via PETRA and HERA and on premises accelerators and their experiments being the host for national and international collaborations.

2.1 Large HEP experiments

With ATLAS and CMS, DESY has joined itself collaborations based at another facility. DESY is a large contributor to these large collaborations in all aspects. DESY is also a national hub, and most prominently for the ATLAS and CMS university groups at the Berlin Humboldt University and University of Hamburg respectively, with whom DESY has strong connections in different matters.
DESY also joined the Belle II collaboration, based at the KEK accelerator facility. DESY is one of the largest member institutes of the Belle II collaboration, and provides collaborative

services to the management and organisation of the collaboration.
HERA was shut down in 2007. At this time, a subsequent electron-proton collider was not on the global agenda, so a lot of effort was put into data preservation and legacy analysis by the collaborations. These effort were successful: Still today, HERA data is accessible, and HERA analyes are being carried on.
DESY also has a strong group doing R&D for future colliders.
DESY has managed to concentrate computing for these communities to two infrastructures: The (global) Grid and the NAF. The number of additional services is small, and probably Belle II has the largest number of additional services for collaborative services, such as wiki, or user management tools.
With these infrastructures, DESY can play different roles: The Grid serves both as a WLCG Tier 2 center for ATLAS and CMS, and as a raw data center for Belle II (similar to a Tier-1 center in WLCG).
The NAF is local analysis facility for the legacy HERA and future collider R&D, a national analysis facility for the German ATLAS and CMS community, and a global analysis facility for the Belle II collaboration.
This concentration offers huge scaling potentials, but needs a good standardization. We will detail the concepts and setup in a later section.

2.2 On-site particle physics experiments

Recently, a number of smaller particle physics experiments started or are about to start taking data on dark matter searches, mostly axion-photon conversions using light-through-a-wall. The experiments are organised in different collaborations, that are similar in structure to the larger HEP collaborations, but smaller in size, usually less than 100 collaborators.
The computational needs are smaller in size, the produced data and the CPU power needed for analysis and simulation are smaller compared to larger HEP experiments.
However, since the collaborations are rather small, not much person power can be dedicated to computing organisation and operation. As such a larger operational and conceptual load has been shifted towards the central DESY IT department to allow these science groups to achieve a similar level of support compared to larger experiments.
Again, standardization and integration into existing setups is the key: ALPS II data taking is completely integrated into the DAQ system of the accelerator division, data is stored and analysed on the National Analysis Facility, making use of as many already existing workflows as possible. In essence the experiences and lessons learned by providing computing and storage infrastructure to one community can be adapted for different scientific communities, thus profiting from synergies.

2.3 The Grid and the NAF

Since 2003, DESY offers Grid resources provided through a batch system based on HT-Condor [4], which holds in 2025 about 20.000 cores providing more than 300 kHS23 in benchmark metrics, *c.f.* [5] for details on the technical setup. The high-capacity dCache storage element [6] both is accessible from the Internet, and from the Grid batch farm. These resources serve different projects, most of them with a high energy particle physics background, as multipurpose infrastructures. Most notably, the DESY contributions to WLCG for ATLAS, CMS and LHCb as well as Belle II are provided using the shared DESY Grid infrastructure. DESY is one of the largest Tier-2 sites world-wide for ATLAS and CMS as well as one of the major Raw Data Centers for Belle II.
The batch system of the Grid is a standard high-throughput computing cluster: Jobs are

placed to a scheduling system, which handles them to a pool of commodity server nodes. The network is a standard peer-to-peer Ethernet network. Inter-node-communication is not foreseen. Scheduling is based on group fairshare, and carried out on a request basis with dynamic core count and memory allocations possible per job on a node.

Nowadays, users do not directly access the DESY Grid resources. Rather, they use large, central job and data management frameworks from their collaborations that distribute their compute load and data to several Grid sites, based on policies and availability. Authentication and authorisation has been shifting away from X.509 based proxies to authorization tokens with the validation shifted to the experiments. In 2007, the NAF [7] was set up in the context of the Helmholtz Alliance "Physics at the Terascale". It complements the DESY and German Grid resources and serves similar projects. As the NAF is intended for direct interactive access and batch processing for end-users, it allows for fast- response workflows necessary for development, debugging, testing and small-scale private productions - important complements to the Grid infrastructure, which provides computing resources for a continuous massive production albeit with higher latencies. With the shift of the computing grid to primarily production workloads, the NAF has become even more crucial for end user analyses. The NAF provides around 9.000 CPU cores. The HTCondor batch system of the NAF is set up in a similar high-throughput computing scheme as the Grid batch system but with focus on integrating grid and scratch storage in a consistent namespace for POSIX-based file access. In addition we have optimized the NAF for interactive Jupyter [8] jobs and a preference for jobs with small requested resources to reflect the fast-response requirements of the NAF.
Data centricity is key to the DESY setup: The dCache storage system and all its data is accessible from both the NAF as well as the Grid. To facilitate usage, DESY has pioneered NFS as access protocol for data access [9]. This enable both industry standard applications, and generic client side caching optimization through Linux operating system tools. Nowadays, most dCache traffic stems from NFS usage. The dCache NFS mount is identical on all systems, and starts with "/pnfs/desy.de/". As the dCache storage is also integrated in the global experimental data management, users can easily stage their data-sets in and out of the NAF with the help of the experiment specific data management software.
In addition to dCache, DESY offers temporary project space to NAF users: The "DUST" storage system. This storage system currently is only available locally at DESY, since it is optimized for high performance and low latency single stream data acces, as well as supporting a broader set of features described in the POSIX standard. Again, data centricity is at the heart of the "DUST" storage system: In 2024, DUST was expanded to both serve NAF and Maxwell HPC system, and user communities of both systems. The namespace was adapted to not reflect the technology of the product ("/nfs" or "/gpfs"), but rather functionality ("/data/-dust"). This way, a transparent cross usage of NAF and Maxwell HPC is made easier since all data is available everywhere which serves as fundamental for the IDAF.
User access to the NAF had been granted by representatives of the collaborations after a request based on X.509 certificates. A special DESY account had been created, and further workflows are based on account name, password, and additional security tokens. In lockstep with the general migration away from X.509, the account creation process is being combined with AAI federated methods. Again, a local identity is created with a corresponding mapping of a user's federated identity, thus avoiding the complications of a global identity namespace.

2.4 Software for and beyond HEP: dCache development at DESY

The dCache storage system got initiated at DESY as *disk cache* software to serve as data and tape archival solution for the experiments at the HERA ring accelerator at DESY as well as

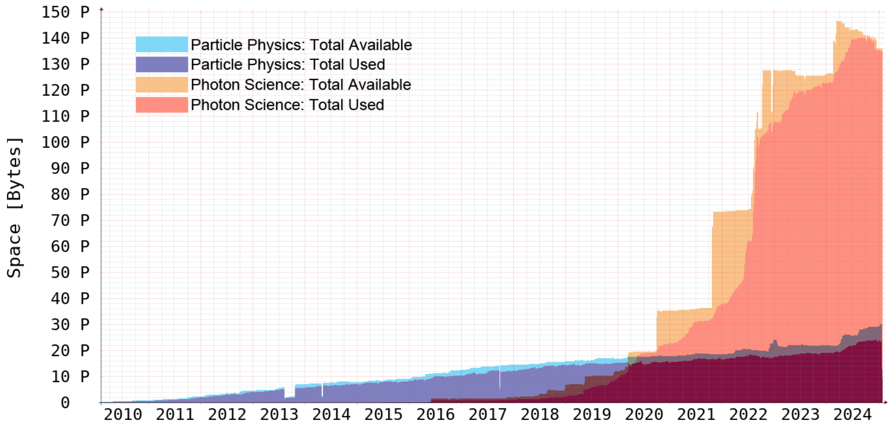

Figure 1. Time evolution of dCache storage capacity provided and its usage for HEP and photon experiments since 2010 to 2024.

the upcoming experiments at the LHC. Today it is a joint development by groups at DESY, Fermilab and NDGF serving the needs for scientific experiments from high energy physics and beyond. About 2/3 of the global WLCG first and second tier level data are hosted on dCache installations world wide. dCache with its individual components is designed as a collections of micro-services communicating over a common message system. This allows for an arbitrary scale out limited primarily by physical constraints.

In addition to hosting one of the dCache developer groups, DESY IT operates some of the largest dCache instances worldwide - either with respect to the amount of stored data, number of stored files or transfer volume.

As shown in Figure 1, following the start of data taking at the EuXFEL and PETRA-II facilities, the amount of data stored on photon science dCache instances at DESY has overtaken HEP data in HEP instances in 2020 in an exponential growth. The dCache instances are tightly integrated in the data acquisition at the photon beamlines, where data are first staged into a dedicated beamline file system distributed at the EuXFEL site and the DESY computing center. After initial processing on the Maxwell cluster, these data are written to the dCache instances for long time archival including at least two copies on tape media.

This mirrors the DAQ requirement at the HEP off-site experiments and photon science has truly emerged as a contender for HEP with respect to storage and processing needs. Based on the experiences gained by HEP experiments, DESY IT has been able to adapt and scale to the photon experiments data deluge.

3 Inter-Disciplinary Analysis Facility

DESY has been developing the NAF further in the form of the Inter-Disciplinary Analysis Facility (IDAF) [12]. In its final form, the IDAF will serve as the umbrella over all user facing computing and storage needs including the NAF, Grid and Maxwell clusters and their corresponding storage systems, as can be seen in Figure 2.

We predict the scientific communities to keep their different requirements concerning cost-effective high throughput computing with embarrassingly parallel workloads versus high-performance workloads for simulations, AI and ML training, or GPU heavy data analyses. Still, we still expect, that the overwhelming number of workloads remain data centric rather

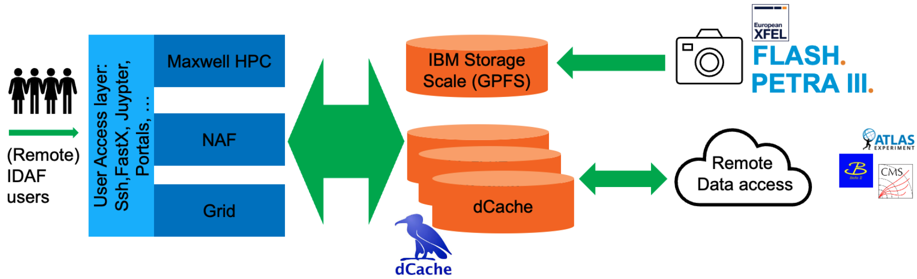

Figure 2. Schematic view of the IDAF, its components and integration.

than primarily CPU and inter-process heavy. However, we also expect the current strict separation between both compute approaches to diminish. With the move towards Python based and more interactive analyses as well as the increased usage of AI and ML techniques within HEP and the changes to experimental workflows for the Photon Science experiments, we expect the analyses to become more similar in their resource needs.
By redesigning the existing compute clusters as components of the IDAF, users will be able to use the same overarching infrastructure for their specific computing workloads. As the common storage infrastructure is at the core of the IDAF with the compute clusters enclosing these, we have already established consistent mount namespaces in the compute clusters and thus removing a hurdle from the past for our users. Work is ongoing to further consolidate the shared components of the compute clusters, that in the future HTC and HPC computing will be merely options for users to select when using the IDAF rather than distinct different systems.
Due to the inter-disciplinary nature of the NAF and with it the IDAF, the design of the computing and storage infrastructure has to be generic usable to all communities at DESY. Nevertheless, each group brings their specific needs and requirements. Thus, we have introduced a third pillar in the form of auxiliary services in addition to the fundamental batch computing and mass storage. Such auxiliary services will be deployed concurrently to the IDAF compute and storage clusters. All compute and data heavy workloads will be scaled out to the underlying HPC and HTC clusters with auxiliary services operating as optional interface layer to these. The Jupyterlab integration can be seen as blueprint for such an auxiliary services as it merely scales out its Jupyter notebooks to the different clusters. Future envisioned applications might be transformation services for serializing complex ROOT ntuples into data formats more suitable to columnar analyses or experiment specific data management services to translate local namespaces to global namespaces.
This approach allows us to avoid lock-in to highly specific experiment requirements with our complex infrastructure while also offering the necessary flexibility to react to the evolving computing and analysis landscape.

4 DESY involvement in photon science computing

4.1 Photon Science Users and Communities

The term *photon science* covers numerous disciplines that use photons to probe and understand matter. These photons can come from advanced light sources such as synchrotrons, free-electron lasers (FELs), or high-intensity lasers. By scattering, absorbing, or emitting photons in controlled ways, researchers can reveal internal structures and dynamic processes of different materials.

Techniques like X-ray crystallography, X-ray diffraction, small-angle scattering, and advanced spectroscopies allow scientists to determine atomic structures or track molecular motions and interactions. Other key methods include X-ray and optical microscopy, coherent diffractive imaging, and various photon-based approaches that generate detailed images of structures ranging from large biological samples down to single atoms—often in real time. These imaging techniques are critical in life sciences, materials research, environmental science, and many other fields.

Under the broad umbrella of photon science, researchers come from diverse backgrounds. In some contexts, photon science might be naively equated with only laser physics or nonlinear optics. However, here it refers to almost all scientific research at modern light sources, excluding (astro-)particle physics. Within this community, there are:

- Materials scientists, working on both fundamental research and commercial applications;
- Chemists, who investigate catalytic processes, solvation mechanisms, and inorganic molecules for quantum technologies;
- Biologists, studying molecular and cellular biology, complex protein structures, and interactions;
- Medical researchers, such as those who conducted SARS-CoV-2 ("Corona") research in 2020. Notably, BioNTech had beamtime at PETRA III during that period;
- Experts in nonlinear optics and high-intensity lasers;
- Researchers exploring novel applications of free-electron lasers.

These communities use many different experimental techniques, especially various flavors of diffraction. Their targets may range from crystals to streams of randomly oriented molecules in the gas phase. In addition, photon science often involves commercial users who utilize cutting-edge light sources for proprietary research and development.

4.2 The Light Sources at DESY

DESY operates and hosts several major light sources, including PETRA III, FLASH, and the European XFEL. These facilities produce high-brilliance X-ray and laser-like beams that drive an array of experiments in fields such as materials science, biology, chemistry, and physics. Each beamline is tailored to certain techniques and detectors, resulting in a wide range of data rates and formats.

4.3 Tools Needed Beyond Computing

Beyond the sheer computational and data-storage requirements, a range of supporting services is crucial. These include:

- Proposal management systems, which organize the scientific proposal process for beamtime allocation;
- User management systems, for handling the thousands of international users who visit or access DESY beamlines remotely;
- Metadata management and cataloging tools, which track information about beamline settings, sample properties, detector configurations, and other essential details of each experiment.

4.4 A Brief Introduction to Photon Science Data Analysis

Photon science data analysis methods are as diverse as the scientific fields and experimental setup. There are essentially no commonalities for example between X-ray tomography and Grazing Incidence Small Angle X-ray Scattering (GISAXS) experiments, except that both are producing a large stack of 2-D images - not necessarily in the same data format or representation (files or streams). There is consequently no standard analysis workflow covering all - or just two - different experimental techniques, and it's not unusual that very different workflows are used even for identical experimental methods.
X-ray tomography of biological specimens for example results in a stack of 2D-images measuring the X-ray transmission/absorption at different angles typically covering 180 degrees. The 2D-images can be combined to result in a 3D-Volume representing the elemental composition of the specimen. Well contrasting materials like bone and tissue are very easy to distinguish and can be automatically separated. However different types of tissues, or bone-mimicking Magnesium-based implants are very difficult to distinguish from their environment. Obtaining an accurate high-resolution segmented model from the 3D-Volume is - albeit heavily supported by machine learning - is still a manual, laborious task requiring excellent remote 3D visualization capabilities.

GISAXS is an important tool to develop and characterize for example photo-conductive or catalytic nanomaterials. Interpretation of the 2D-images obtained in GISAXS experiments is however an ill-posed inverse problem. It's not directly possible to obtains the structural composition of the material from the 2D-images. Typically the 3D-structures are determined by simulations of possible topological arrangements, and use the simulation best matching the series of 2D images.

The large difference in methods, workflows, and software stacks is accompanied by large variations in detectors, data formats and data transport. Some of the differences are being abstracted by control systems like TANGO/SARDANA, developed by DESY and several other large Photon Science centers, KARABO for the EuXFEL or a combination of DESY internally developed software like DOOCS, EPICS and TINE, utilizing different messaging systems like 0MQ, ASAP::O, JMS or RabbitMQ.

In most cases, like GISAXS or X-ray tomography, a real-time interpretation of the measurement data is very hard to achieve, and usually not possible. In such case the analysis has largely to be performed offline and file-based rather than on data streams. So the experiment setup, the measurement data, interpretations, simulations and various external data sources have to managed in conjunction.

4.5 High-Energy Physics vs. Photon Science Approaches

The photon science workflow contrasts sharply with that of high-energy physics (HEP). In HEP, the standard procedure involves reconstructing physics objects from event data, looping over those events, and then applying counting or filtering to build probability distributions. In photon science, however, images are interdependent and cannot be processed individually in an event loop. The problem is typically an inverse one: reconstructing structures from measured images via Fourier transforms, tomography, phase-retrieval algorithms, or other specialized methods. These analysis strategies do not neatly map onto the parallelization approaches common to HEP (i.e., splitting data into independent events).

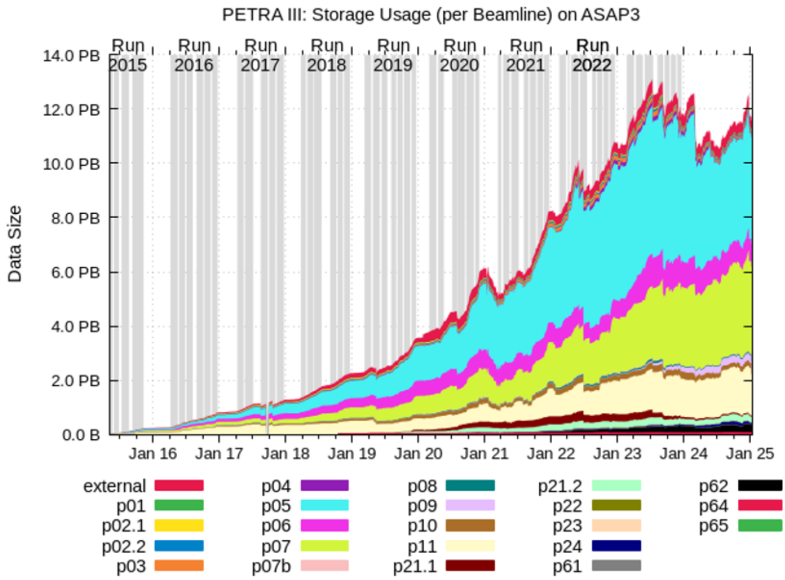

Figure 3. PETRA III storage usage per beamline on the primary ASAP3 GPFS. Please note, rather than capacity or overall amount of data stored, data under active analyses are shown. After 180 days, the beamline data are archived to tape and removed from the ASAP3 GPFS system causing the dips in usage.

Commonalities do exist at different stages, e.g. photon science benefits from the online data acquisition and filtering knowledge. Also, several algorithms, e.g. for calibration can be used in common. These similarities are exploited by scientists at DESY. Relevant for the IDAF and these proceedings are commonalities e.g. in analysis facility concepts and workflows. For example, both disciplines share a data centric approach, infrastructure access methods and workflows. Thus, synergies within the data management and analysis infrastructure increase the overall efficiency.

4.6 Data Lifecycle: From Experiment to Offline

With the advent of PETRA III, data storage and computing demands at DESY's light facilities increased drastically. A scalable, centralized approach was needed to handle both the volume of data and the compute power required for real-time or near-real-time feedback and subsequent offline analysis. Systems like ASAP3 [10] (for PETRA III) and the XFEL infrastructure exemplify how data reduction, triggering, and tiered storage solutions (including GPFS and dCache) streamline the data flow from the beamlines to central tape archives. The evolution of the data stored and analysed within the primary ASAP3 GPFS storage system at a given time is shown in Fig 3.

4.7 Offline Analysis: The Maxwell HPC Cluster

DESY implements a data-centric computing approach to meet the high-performance computing (HPC) needs of photon science experiments. All major computing tasks—online reconstruction, offline analysis, simulation, or preparation for future experiments—center on efficient data handling and high-speed access to storage.

Cluster Organization and Capabilities.

The Maxwell HPC cluster is designed with several key criteria in mind:

- Interactive, scalable access for many users concurrently;
- High-performance connectivity to all relevant storage systems;
- Centralized management of system and application software, with user support;
- Flexibility to accommodate a wide array of photon science workloads and workflows.

Originally, Maxwell [11] was a smaller HPC setup optimized for parallel computing tasks in accelerator research, using InfiniBand networking for both MPI scalability and fast data access to GPFS. Over time, this has expanded significantly. Each node in the cluster also provides 10 Gbit Ethernet connections for compatibility with dCache and other resources.

Photon science workflows frequently need large memory allocations and do not always parallelize well, because the images in a dataset can be correlated. Since some workflows require significant single-thread performance, DESY typically chooses CPU models that balance many cores with high clock frequencies. Maxwell nodes also carry larger-than-typical amounts of RAM, accommodating memory-intensive tasks such as 3D segmentation or advanced simulations.

GPU Resources.

GPU usage on Maxwell is substantial yet heterogeneous, reflecting different scientific demands:

- Plasma wakefield acceleration simulations may efficiently scale to multiple GPUs with moderate memory demands;
- Popular deep learning applications like AlphaFold require only a single GPU but with a large memory footprint;
- Large Language Model (LLM) inference, increasingly used in accelerator control, tolerates slower GPUs yet needs ample memory;
- 3D rendering of images or real-time visualizations also benefits from GPU acceleration, though often with different resource requirements than AI/ML.

Accordingly, Maxwell integrates multiple generations and types of NVIDIA GPUs to satisfy these varied computational profiles.

Network and Storage Topology.

The Maxwell cluster uses a three-layer fat-tree InfiniBand fabric. Different parts of the cluster have varying generations of InfiniBand; while not common in HPC, this arrangement reflects a staged procurement process and a distributed ownership model among DESY groups. The blocking factor across the network is generally below 6–8, and the GPFS storage systems connect at multiple leaf-level switches to maximize bandwidth and provide redundancy. Dedicated nodes with 100 Gbit Ethernet links import detector data from beamlines in real time, injecting it into GPFS for immediate availability to the cluster.

Ease of Use and Scheduling.

Because many photon science users have limited HPC or even Linux experience, the Maxwell environment is designed to be as user-friendly as possible. Jupyter notebooks are a popular entry point for data analysis and are employed both during beamtime for online analysis and offline for post-experimental studies. A custom JupyterHub service integrates with the SLURM scheduler [8][14], allowing users to launch interactive sessions across various cluster partitions with minimal overhead. This arrangement also supports reserved or priority resources for experiments that need guaranteed computational capacity during their beamtime.

User and Software Management.

User management on Maxwell and DESY's other HPC resources must cater to a large, international user base. Many scientists use beamlines periodically, so their workflow demands might spike around experimental campaigns. In addition, specialized commercial or in-house software packages—often supporting only local file access—must be made available centrally and efficiently. Fine-grained access control lists (ACLs) on GPFS ensure that proprietary datasets remain secure, while beamline scientists can still share data as needed for collaboration or post-analysis with the broader photon science community.

5 Summary and future role(s) of DESY

DESY has a long history offering compute infrastructures for HEP and photon science [13]. Interdisciplinary synergies between domains are strategic elements for the design of the infrastructure. Furthermore, DESY has been a center of German high energy physics (HEP) research since 1959 with on-site experiments and theoretical research. Since the turn of the century, the research scope has evolved beyond HEP. With the PETRA-III or FLASH facilities on site and the collaboration within the European XFEL, DESY has become a major photon science laboratory as well. To be able to serve the computing and data taking needs for these various user groups and scientific backgrounds, DESY IT has developed and deployed data acquisition, storage and computing solutions
The scientific landscape is constantly evolving, and so does DESY and the IT offerings and technology. Some challenges we see are summarized in the following.

5.1 Analysis Facilities: generic solution for all

As well as within WLCG, the access to resources for scientific data analysis is being discussed. DESY has a long history of providing such resources for a vast number of users, stemming from diverse backgrounds. DESY will continue to work towards generic solutions for all scientists and science domains.
Science at DESY is data centric. All technology choices must be integrated with the data, and the way data is stored and accessed. Together with users and experts from the communities, we perform R&D to adapt or advance infrastructure technology. We define, build and operate the analysis facility of the future: An innovation factory for HEP- and photon science-related IT solutions as well as a reliable IT center for data-centric services.

5.2 New role in German WLCG computing approach

The German WLCG environment will change such that the university Tier-2 centers will disappear before the start of the HL-LHC. Their CPU pledges will be fulfilled by the German Nationales Hochleistungsrechnen (NHR) centers, a collection of large HPC sites. The disk pledges will be delivered by the two Helmholtz centers KIT and DESY. The additional data storage might be relatively small, however, the data access will be decentralized and federated. New R&D is planned to optimize for efficiency and reliability.

We are working on the upcoming challenges in the form of distributed compute resources at remote HPC centers for which DESY will serve as data lake hosting the input and output data together with KIT. Thus, the demands on a stable and scalable storage operated and developed at DESY will grow not only with respect to the internal IDAF but also due to the external dependencies on DESY as a Germany-wide data lake.

5.3 PETRA IV

For PETRA IV, several new challenges will require additional R&D and new concepts. The data rates will drastically increase even beyond the current DAQ needs, making near real-time analyses and data reduction a necessity. For PETRA IV we plan to offer services for a broader variety of science domains as "photon science as a service" - which will call for easier and more standardized access to compute, storage and workflows. Most importantly, access to commercial users will gain in importance, posing not only technical challenges. The preparation for these challenges has already started.

5.4 Sustainability

As final global challenge, DESY IT has to strive for an even more sustainable operations model avoiding unnecessary waste of precious resources - again, R&D is ongoing.

References

[1] The Petra-III storage ring, https://petra3.desy.de
[2] The European XFEL, http://www.xfel.eu/
[3] The Free-electron Laser, https://flash.desy.de
[4] Distributed Computing in Practice: The Condor Experience. Thain D., Tannenbaum T. and Livny M., Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005.
[5] A. Haupt et al. The DESY Grid Centre Journal of Physics: Conference Series **396** (2012) 042026
[6] The dCache project website http://www.dcache.org/
[7] A. Haupt, Y. Kemp and F. Nowak, J. Phys. Conf. Ser. **513** (2014), 032072 https://doi.org/10.1088/1742-6596/513/3/032072
[8] Reppin J., Beyer C., Hartmann T., *et al.* Interactive analysis notebooks on DESY batch resources., Comput Softw Big Sci 5, 16 (2021).
[9] Elmsheuser J., Fuhrmann P., Kemp Y., Mkrtchyan T., Ozerov D. and Stadie H., Journal of Physics: Conference Series vol **331** (IOP Publishing) p 052010 (2011)
[10] ASAP3 - New Data Taking and Analysis Infrastructure for PETRA III, M. Gasthuber et al., Journal of Physics Conference Series **664(4)**:042053 (2015)
[11] DESY Maxwell HPC Cluster website https://maxwell.desy.de/
[12] DESY IDAF website https://idaf.desy.de/
[13] C. Beyer, S. Bujack, S. Dietrich, T. Finnern, M. Flemming, P. Fuhrmann, M. Gasthuber, A. Gellrich, V. Guelzow and T. Hartmann, *et al.* EPJ Web Conf. **245** (2020), 07036 https://doi.org/10.1051/epjconf/202024507036
[14] SLURM: Simple Linux Utility for Resource Management. Jette M. and Grondona M., United States: N. p., 2002. Web.