Defining the future scientific data flow for multi-disciplinary research data

Felix Bartusch[1], Kolja Glogowski[2], Ulrich Hahn[1], Michael Janczyk[2], Steve Kaminski[1], Jens Krüger[1], Volker Lutz[1], Gerhard Schneider[2], Mark Seifert[2], Dirk von Suchodoletz[2], Thomas Walter[1] and Bernd Wiebelt[2]

[1] Zentrum für Datenverarbeitung, University of Tübingen ;
[2] Rechenzentrum, University of Freiburg

Digital data and computerized workflows are at the core of almost every domain in science. Data is not only the base for scientific publication but can become equally important by itself. The discovery of new insights from huge amount of (unstructured) data for completely unrelated fields already have made big data a valuable asset for scientific findings. The value of the ever-increasing amounts of data for subsequent use and the requirements of funding agencies generate the need for formalized Research Data Management (RDM). Modern digital workflows involve more than one system to generate, compute or visualize ever-larger data sets. Thus, the operators of the large scale federated research infrastructures at the involved HPC computing centers in Baden-Württemberg face the challenge of providing suitable storage services. Such a Storage-for-Science (SFS) represents an essential building block for the anticipated state-wide data federation. In addition to the integration of the various pre-existing infrastructures, the long-term identification of data sets, their owners, and the definition of necessary metadata becomes a challenge. The implementation and provisioning of a RDM system needs to be organized together with the scientific communities and has to fit well into the growing Research Data Repositories landscape.

1. Introduction

Modern scientific work has become significantly digital, meaning it uses various devices, programs and tools to gather and process data in a multitude of ways. The involved digital research workflows are getting more complex and the tide of data processed or created through them is ever rising. Data intensive computing (DIC) involving big data or methods like deep learning provide a new perspective on existing data. The unstructured approaches to store and manage data of the past are not viable in the modern world of multi-disciplinary and multi-institutional research as well as over geographically distributed locations. A research data management system (RDMS) needs to answer how to properly store and present data on the long run, from short living projects and in an environment of high fluctuation of researchers as well as bridging from the existing landscape of network file systems into a world of flexible scientific workflows. The

Corresponding author:jens.krueger@uni-tuebingen.de

Felix Bartusch et al.: Defining the future scientific data flow for multi-disciplinary research data.
In: Heuveline, V., Gebhart, F. und Mohammadianbisheh, N. (Hrsg.): E-Science-Tage 2019 – Data to Knowledge, Heidelberg: heiBOOKS, 2020, 110–127. DOI: https://doi.org/10.11588/heibooks.598.c8421

long-term identification of data sets, their owners and the definition of necessary metadata presents a challenge as well as the integration of large-scale object storage concepts. Future RDMS should consider the complete data life cycle, spanning from data acquisition over the various stages of computation, visualization to long-term archiving and publication. Additional components on top of the RDMS are desired to offer added-value services like special purpose repositories, semantic search, indexing or versioning. Individual scientific communities should be enabled to provide tailored services for their specific data management tasks as well without the need to run their own RDM enabled storage systems.

To help individual researchers to adhere to the FAIR principles [1] modern data management should extend beyond the traditional data-handling performed by now. Researchers often do not standardize metadata, making interoperability and sharing difficult. Data curation, the selection of data sets of relevance, and the removal of irrelevant data are often not formalized steps in the workflow. A RDMS can help to solve these shortcomings in today's workflows by providing tools and services to support researchers in their data management tasks and automation of various workflows. To address the researchers' needs the HPC cluster sites of the BinAC[1] in Tübingen and the NEMO[2] in Freiburg complement their compute infrastructures by a holistic approach to a RDMS. The Storage-for-Science (SFS) is designed to run in a cooperated, federated way spanning both locations. The system will become an integral part of the Baden-Württemberg data federation [2, 3] offering high performance data paths to other research infrastructures like HPC clusters and cloud systems. The establishment of the Science Data Center BioDATEN[3] will provide further input from one of the core scientific communities onto the system design and intended additional services.

The following paper is structured as follows: It gives an overview of the current state of data management nationally and internationally. The requirements stemming from the data life cycle, today's and future workflows of HPC and DIC user communities will be discussed. Further, it explores options to extend and optimize existing scientific workflows. From this discussion it tries to provide an overview on the concept of the Baden-Württemberg data federation and a coherent framework for the design of a research data management aware large scale data facility. The SFS is a RDMS supported by the DFG and state funds to provide joint storage and research data management functionality for various research groups in Tübingen, Freiburg, Ulm and Stuttgart. The texts extends upon the inputs provided in a paper published by the North Rhine-Westphalia RDMS group [4], an article presented at the DFN forum in 2017 [5] and a discussion on researchers' needs in the HPC domain presented at the eScience days 2019 in Heidelberg [6].

[1] `https://www.binac.uni-tuebingen.de`
[2] `https://www.hpc.uni-freiburg.de/nemo`
[3] `http://www.biodaten.info`

2. Related Work

The Research Data Alliance (RDA), has been brought into existence to tackle the challenges of modern research data management from an infrastructure perspective [7]. From the perspective of research funding institutions there is an increasing movement to impose a good scholarly practice by requiring scholars and scientists to plan and execute sound data management.[4]

Initiatives in various countries to provide infrastructures, support and services exist for quite a while: The UK as well as the Netherlands were among the first movers in Europe to provide a range of discipline-specific data centres or to offer support for data archiving for scientists from various disciplines [8]. The Digital Curation Centre[5] is one of the most relevant national providers of expertise. In the Netherlands the Data Archiving and Network Services (DANS)[6] focuses on the data archiving of the social sciences and humanities. It started the development of cost models from the very beginning to provide a sustainable service in the long run [9]. To address the needs of data globalization and high performance access Onedata in Poland follows an approach relying on a global registry for mediating metadata synchronization and file transfer [10]. Overseas in the US, the National Science Foundation supports several projects to build sustainable infrastructure by trying to support all sides: scientists, software developers, librarians, archivists, and information scientists as well users to deeply engage in a joint process. Two prominent examples are the Science Gateway Institute USA[7] which supports domain developers to create tailored portals for their communities [11] and the US Research Software Sustainability Institute[8] helping to preserve related software. A similar institution was founded in the UK, the Software Sustainability Institute[9] which is also strongly engaged in training and education [12]. Federated infrastructures like EUDAT can provide guidance on which services are to be provided and how the several challenges are tackled [13, 14]. In Germany the developments gather momentum, one significant step forward is the development of the National Research Data Management Infrastructure [15]. Several regional, often federated initiatives work on a RDMS, like the consortium of several universities in North Rhine Westphalia lead by the university of Aachen [4]. A flexible and scalable storage system was created using a Quobyte[10] solution for the de.NBI Cloud in Tübingen, providing a broad range of functionalities. It features both a full integration with OpenStack and additionally a multipurpose S3 object storage as well [16]. The ViCE project [17] worked on different use cases for virtualization or containerization. It evaluated the various efforts to access external data stores in relation to the actual location of the virtual research environment [1]. A discussion on the design of a RDMS can be found in [5] as well as a RDMS in relation to other large scale infrastructures in [6].

[4] See e.g. German Research Foundation [19], the Federal Ministry of Education and Research [20] or the German university rectors” conference [21, 22].

[5] For more information see, `http://www.dcc.ac.uk`

[6] For more information see, `http://www.dans.knaw.nl`

[7] For more information see, `https://sciencegateways.org`

[8] For more information see, `http://urssi.us`

[9] For more information see, `https://www.software.ac.uk`

[10] `https://www.quobyte.com/case-studies/uni-tuebingen`

The aforementioned programs, initiatives and implementations can only give a cross-section overview about the vast number of worldwide activities dealing with modern research data management in its broadest sense.

3. Data Life Cycle

Extensive data analysis has become an irreversible trend in modern science. The increasing scale of scientific data is invalidating classic methods that had previously been considered to be good enough. For example, storing large amounts of data in a single filesystem with the Data Management Plan (DMP) consisting in the proper naming of files and directories is probably one of the least scalable methods. As a popular approach, it just requires a naming convention to start with and disciplined scientists adhering to it. It might even work reasonably well for a limited time within a small group of scientists and a moderate amount of data. However, further developing this method into a proper DMP, spanning the complete life cycle of the data and extending it to larger groups of scientists, is a tedious if not impossible task.

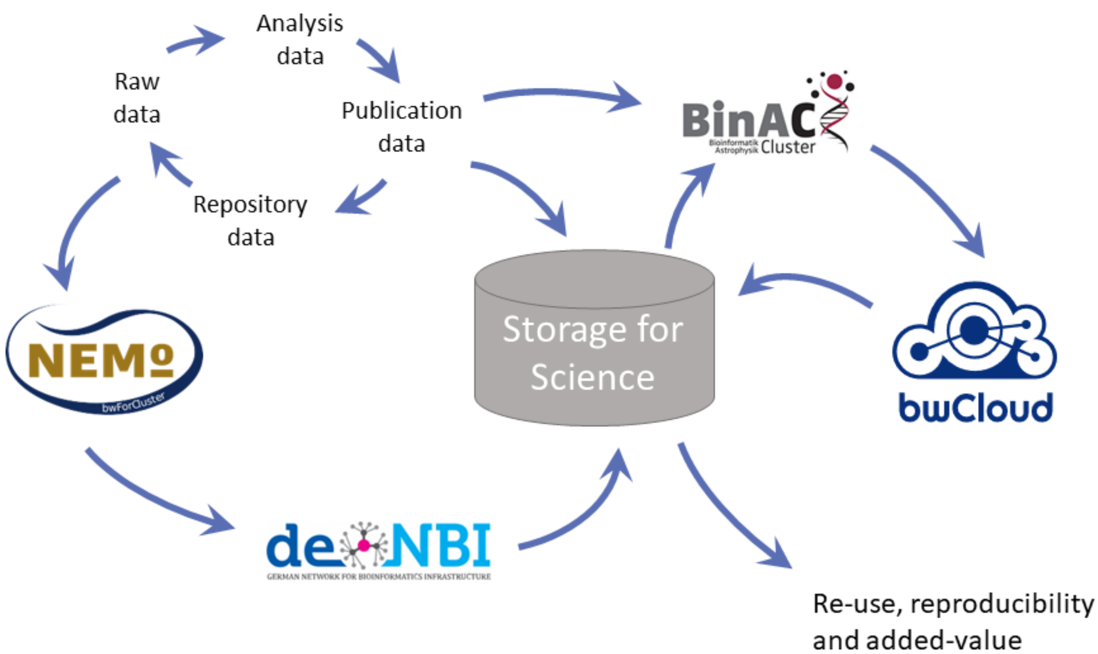

Figure 1.: Depiction of the data life cycle involving the bwForClusters NEMO and BinAC, the de.NBI Cloud in Freiburg and Tübingen as well as the bwCloud, relying the on the Storage-for-Science.

A modern DMP has to consider the complete life cycle of the data (see Figure 1) residing on multiple geographically distributed resources [23]. Typical examples for the provenance of raw scientific data include experiments and scientific simulations. In both cases, if large amounts of raw data are generated, a sufficiently large and performant storage resource is required (LSDF – large scale data facility). This storage resource needs to be either directly attached to the scientific device generating the data (e.g. microscope, compute cluster) or it needs to be connected via a high speed network and protocol. Alternatively, the scientific device could be fitted with a fast local storage that caches the data on acquisition and forwards it over time to the large scale data facility. After the raw data has been collected it has to be annotated and enriched with an initial set of qualifying metadata. This includes both, metadata for scientific processing and metadata for governance purposes. To enable scientific reproducibility, data is stored in an immutable way after acquisition and major processing and refinement steps. Changes are recorded in a way they can be tracked and possibly rolled back and reapplied. This is the core requirement for offering the data as a public repository to a wider scientific audience. Furthermore, this enables the publishing of research data respecting all four FAIR principles (findable, accessible, interoperable, re-usable [1]).

4. Annotation

Throughout a modern research data life cycle (see Figure 1), an appropriate metadata management is crucial for the immediate and long-term preservation, discovery, publication and reuse of scientific data. From the very beginning of a research project, a DMP should provide information about discipline specific metadata and related vocabulary needed for the enrichment of data objects. Especially the continuation of a provenance information chain on data objects throughout the life cycle is a key aspect of metadata management.

Research projects are usually in the need for support from infrastructure providers regarding metadata management. They need to offer an ecosystem (see Figure 2) including storage, search systems, pid services, interfaces and presentation layers for metadata management and have to take care of keeping metadata standardized (e.g. METS/PREMIS [29, 30]) and interoperable between systems (via protocols like OAI-PMH [31]). Software platforms together with rich metadata must provide an abstraction layer to enable researchers to effortlessly track, move and collect their distributed data objects in complex environments. In distributed systems, a search for data objects from rich metadata information is considered to be far superior to manual browsing and filename search.

The enormous amount of digital objects created in modern HPC environments makes an entirely manual generation and management of metadata impractical. In fact, a phenomenon known as metadata bottleneck [11] must be circumvented via the usage of automated processes to support researchers and system administrators. One step further, automated parsing of basic information from e.g. running HPC jobs into metadata records, such as job-id, user-id, number of allocated cpu and so forth, could help to complete metadata information. Other processes could check for changes on data files and automatically generate and parse provenance information into related metadata records. Based on the specifics of the research project, several types of metadata information like file ownership, access rights and license information for data reuse could be conveniently configured as predefined values for all data objects within the project scope. With correspondingly higher effort, even parts of discipline specific metadata could be automatically collected via script based scanning through job output and log files. Based on automatically created metadata records available in a searchable data store, researchers must furthermore be

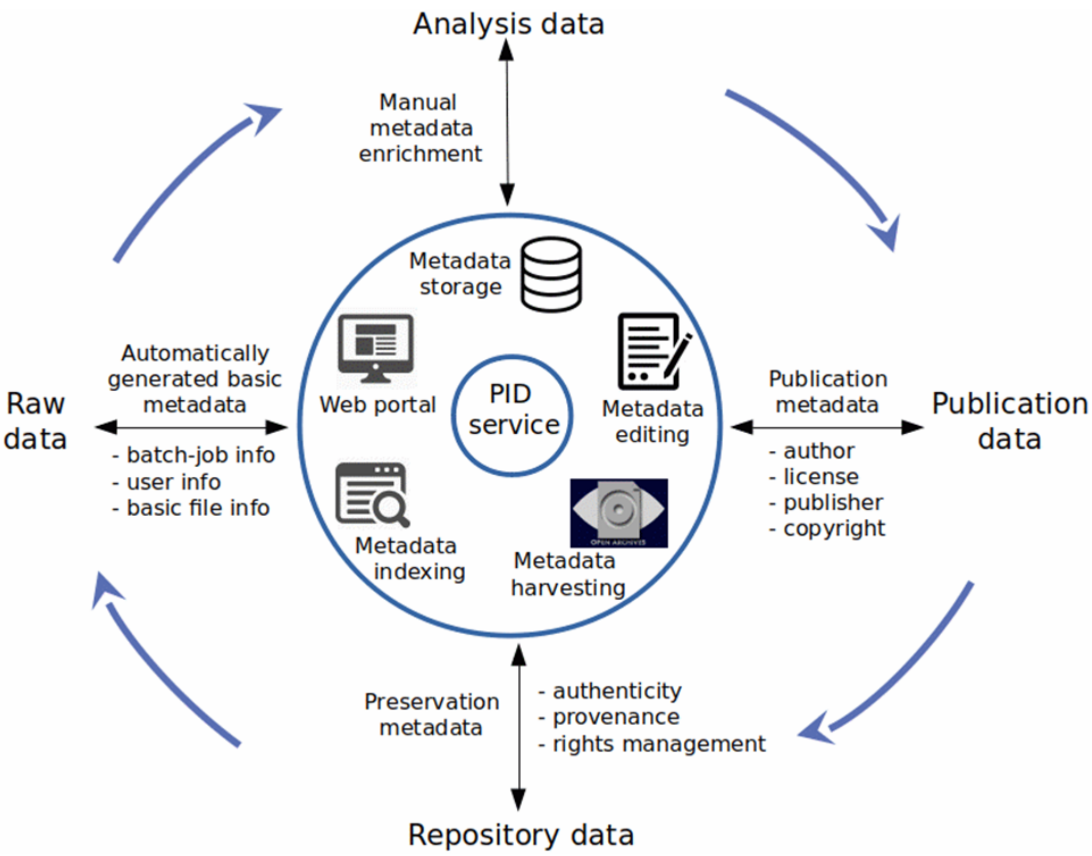

Figure 2.: Illustration of a metadata management chain. Icons where taken from [24, 25, 26, 27, 28].

enabled to manually add descriptive information via appropriate tools and standardized workflows provided by the technical infrastructure.

Metadata information can also support data staging processes on federated storage systems such as the SFS based on file status information available from the metadata. Data records with the status „cold“ could be semi-automatically moved to affordable and less performant storage media, usually until the end of a research project, whereas data annotated as „hot“ must be kept on expensive and performant storage for subsequent data analysis.

5. Data Staging

Research data are often produced on storage systems at different locations from where they are further processed or stored for future use. One example for this is a simulation that writes its results to the parallel file system of an HPC cluster. These results are then often analyzed using compute resources that have no direct access to the original parallel file system (e.g. cloud-based virtual machines (VMs), dedicated high-memory servers or visualization workstations, or external HPC clusters). Other examples are scientific instruments (e.g. telescopes, microscopes or sequencers) that create large amounts of observational data on local cache storage systems. After acquisition the data need to be moved to other storage systems, presumably at a different geographic location, where they can be stored safely, calibrated, visualized and analyzed.

In very simple cases, direct copies can be a viable solution for transferring data between different locations, but handling data manually gets quickly complicated when more complex scenarios and environments are involved. Especially when data need to be replicated at multiple locations on storage systems belonging to different organizations, or when data sets are shared between collaborators, the utilization of a dedicated RDMS becomes inevitable.

Starting from data stored in mutable files on a file system that can be directly manipulated by the user, a first step towards a managed system is to define data sets or collections of files, and ingest these data sets into such a system, creating well-defined objects that can be associated with persistent identifiers. This way data sets can be distributed asynchronously to any number of locations, without worrying about changes to local files. The transition from mutable files to well-defined, managed objects also marks a point where data sets can be augmented with additional governance metadata that define ownership, access permissions and data retention in a way that is independent of location-specific administrative domains [5], and allows collaborations to manage their data using a uniform, location-independent namespace.

With regard to the anticipated data federation in Baden-Württemberg, a suitable data management system needs to be vendor-independent, and must support heterogeneous storage solutions as well as a variety of access protocols. In order to gain high acceptance from scientific communities, it also needs to be able to support different workflows and scale well for large data sets and high numbers of files. It should ideally be an actively maintained open-source software that is based on standard software components and provides modern application programming interfaces (like REST APIs) for its subsystems. Two popular examples of existing open-source data management systems are iRODS[11] and Rucio[12] [33].

6. Reproducible Research Environments

The term research environment encompasses the software stack, explicit description of workflows, custom scripts, and settings a researcher used for processing the data sets. As these components of the research environment are stored on computational resources, they should be also considered as data. As a major consequence of this, the whole data life cycle, as well as the FAIR principles, are applicable to the research environment. Leveraging the data life cycle, research environments should be versioned and archived in order to enable reproducible computations.

The explicit notation of an analytical workflow in a workflow language is good scientific practice and the publication of these workflows according to FAIR principles is essential for transparency and reusability [34]. Analogous to open data, open research environments

[11] `https://irods.org`
[12] `https://rucio.cern.ch`

increase reproducibility and – not unselfish for an author – also increases the number of citations [35]. It is obvious, that a simple listing of software used by a pipeline is not sufficient for other researchers to reproduce a complex computational analysis [36, 37]. In recent years software containerization techniques like Docker[13] and Singularity [38] were developed, allowing researchers to package a specific software stack, including an operating system and additional data like custom scripts, in a single entity called container image. This technique enables versioning and archiving of research environments and pipelines, as the environment is bundled in one image [39, 40]. It also enables researchers to share this single entity via scientific data repositories, public container hubs, or institutional repositories residing on systems such as the federated SFS.

In contrast to a static research environment installed on bare metal, the container is portable and can be deployed on arbitrary computational resources. When performing computations on big datasets, it could be faster to move the containerized research environment to the data, instead of moving the datasets to a computational resource. One advantage of a state-wide data federation would be, that researchers can decide to move their data to the containerized research environment, or vice-versa.

A disadvantage of software containers is their black-box character, as it is not obvious which software versions and scripts are bundled. As stated in Section 3, organizing the containerized environments using a simple naming scheme is not easy to implement and would not scale very well. Using a standardized metadata schema[14] to describe containers and their contents would break this black-box character to some extent and make searchers in an environment registry possible. If the containerized research environment was used for processing datasets, recording this container in the metadata of resulting datasets increases the result's provenance. Ideally, the researcher is able to create a seamless provenance chain starting from the instrument, that produced the raw data, over several processing steps using software containers to the final results.

The results of two projects, both funded by the state of Baden-Württemberg, could be used to provide researchers tools for handling containerized research environments in the proposed concept. The ViCE project created a prototype of an image registry, which can be used to manage research environments scientists tailored, based on their needs. Such an image registry can then be used to exchange images amongst platform borders, e.g. different HPC clusters [11]. Although container images allow a reproducible deployment of complex software stacks, in terms of long-time archiving and executability the container runtime is an additional point of failure. The project CiTAR (Citing and Archiving Research)[15] provides a platform for archiving virtual machine and container images. They are normalized to one standardized format (e.g. OCI image format), for which the runtime (e.g. runc) is available in the CiTAR service. A citable handle is assigned to the archived environment. A data federation, that promotes and supports the tools and workflows introduced in this section, would allow researchers to create and manage reproducible research environments.

[13] `https://www.docker.com/`

[14] `https://github.com/opencontainers/image-spec/blob/master/annotations.md`

[15] `http://citar.eaas.uni-freiburg.de/`

6.1. Virtualization

For the application of the bwForCluster „NEMO“ [41] one of the key requirements was to provide the complex computing environments for the experimental particle physics collaborations like the „Compact Muon Solenoid“ (CMS) or „A Toroidal LHC ApparatuS“ (ATLAS) at CERN. Inside this environment they had to provide services, which enable researchers to access scientific software and data. After discussion with the scientific groups the solution was to provide virtualization based on OpenStack [17]. As preparation for the bwForCluster NEMO a test cluster was created to develop this solution. The CMS groups at the KIT already had developed a resource broker „Responsive On-demand Cloud Enabled Deployment“ (ROCED) which communicates between different batch systems. This had to be enhanced with a connector to Moab/Torque for NEMO and SLURM for the ATLAS groups in Freiburg [42]. When NEMO started August 2016 the solution based on ROCED was already established and fist jobs already were computed [43].

For the upcoming replacement of the bwForCluster „BinAC“ [44] a similar strategy is anticipated, to be able to separate the environment for sensible data from the broader pool of compute and storage resources. Since virtualization encapsulates the environment and creators of this „Virtualized Research Environment“ have root access inside the virtual machine, the OpenStack environment has to be separated from the cluster. Parallel and home file systems like BeeGFS, Lustre, NFS, etc. are usually not secured for performance reasons and usually this is not necessary, since traditional clusters are black boxes and can only be accessed through special login nodes [6]. Having a closed environment without external mounts data for processing and scientific software has to be streamed or copied from external services. New approaches for caching data and software (like CVMFS, XrootD, Squid, etc.) have to be implemented and tested [45].

7. Baden-Württemberg Data Federation

The state of Baden-Württemberg has recognized the relevance and importance of RDM for sustainable and future-oriented organization of research data, its availability and publication. This is addressed in the jointly developed eScienve concept[16] and through two statewide funding programs (Research Data Management and Virtualized Research Environments).

The plan of action for development and deployment of federated research data management infrastructures follows the recommendations of the German Information Infrastructure Council.[17] Existing infrastructures (LSDF in Karlsruhe and Heidelberg) are becoming enriched with research data management capabilities. Emerging infrastructures (SFS in Freiburg, Tübingen, Stuttgart and Ulm) are already procured with research data management as a core functionality. Hence, researchers are not just simply provided with larger storage systems. Additionally, they will be given the tools they need to devise and implement DMPs suited for their field of research. To prevent isolated non-sustainable

[16] For more information see, `https://idw-online.de/de/attachmentdata37340.pdf`

[17] For more information see, `http://www.rfii.de/download/rfii-recommendations-2016-performance-through-diversity`

provisional solutions, these efforts are coordinated in a data federation. Establishment of a federation by connecting existing individual data management systems can be studied from several examples: Initiatives in the Helmholtz Association (Supercomputing and Big Data program)[18], activities at the national level (BMBF, DFG: LIS programs)[19], activities on the European scale (EUDAT)[20], and possibly EOSC[21] in the future) and finally attempts on the international level (Research Data Alliance)[22].

Access to storage systems requires a variety of protocols to meet the diverse needs of scientists. These include remote file system protocols such as NFS and CIFS/SMB as well as object store protocols such as Amazon S3. Additionally, performant network connections have to be offered for the coupling of data and HPC systems.

In many areas of application, sensible data is subject to complex and restrictive usage rules, which have to be taken into account. These restrictions are typically due to storage of personal data (e.g. medicine, social sciences, psychology and mobility research) or storage of secret data (e.g. economics and engineering). A data federation can establish, support and enforce policies to govern the adherence to these rules.

The basis for the Baden-Württemberg data federation are the existing and emerging infrastructures for data management and data storage in the state (see Figure 3). These include the parallel file systems of the HPC systems, LSDFs, specialized data analysis systems, repositories and archiving systems such as bwDataArchiv and bwDataDiss. In addition to these statewide storage systems and repositories, gateways need to be established to other national and international systems operated by their respective communities.

8. Storage-for-Science

The concept of SFS is a distributed federated storage system that spans over four different university locations in Baden-Württemberg. Besides offering a large storage capacity for research the focus is on additional RDM functionality to enable researchers to properly annotate their data from the start to the end of their project data life cycle. Figure 4 shows the main building blocks for the SFS. Caching systems provide access to existing data and allow data import from all kinds of measuring instruments. The main sites in Freiburg and Tübingen provide central mass storage for active data as well as a geo-replicated long-term storage infrastructure for „cold“ and archived data. Configurable data movers will move the data between the already mentioned hierarchy levels of SFS. The data movers and their APIs will enable the control of the data flows within the SFS system, to the HPC systems such as NEMO and BinAC and to other storage systems, thus spanning part of the data federation.

In addition to generic RDM methods for capturing and handling metadata, the development of additional RDM tools in collaboration with the research groups on the SFS is

[18] For more information see, `https://www.helmholtz.de/en/research/key_technologies/supercomputing_big_data`

[19] For more information see, `https://www.dfg.de/foerderung/programme/infrastruktur/lis`

[20] For more information see, `https://eudat.eu`

[21] For more information see, `https://www.eosc-portal.eu`

[22] For more information see, `https://rd-alliance.org`

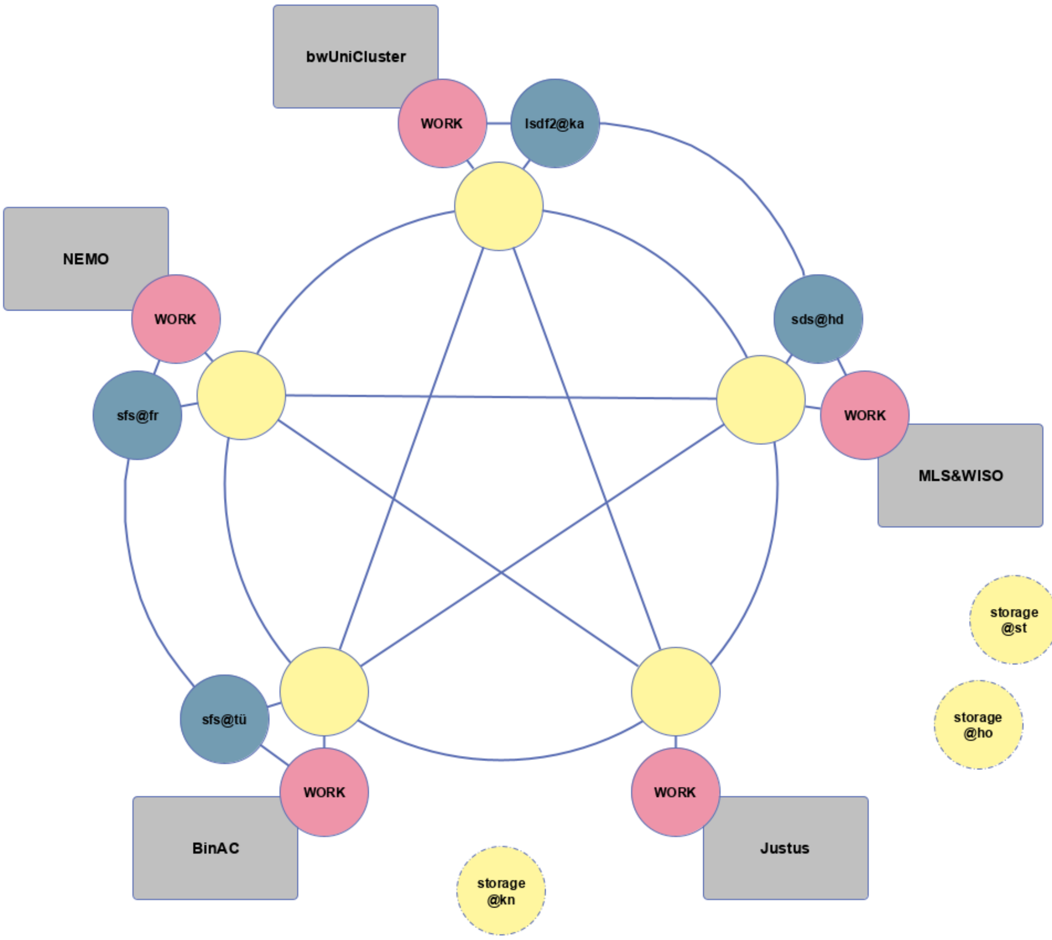

Figure 3.: Depiction of the anticipated state-wide data federation, involving Storage-for-Science (blue circles on the left).

key for the integration of community based metadata standards. Manual and automatic generation of metadata will be supported, therefore the data movers will always move metadata alongside data. Consequently, data and metadata consistency is given on every level of the SFS hierarchy thus generating an integrated research data management. Beside the benefits of getting integrated data management for project data life cycles and research data life cycles further improvements for research by utilizing the data mover capabilities to implement workflows that resemble the data analysis and project workflows in research are expected. The integration of existing workflow engines like Galaxy[23], UNICORE[24] and NextFlow[25] is envisioned.

The technical building blocks for the SFS system will be file system based NAS and object storage systems. While the file system based technology will be used for the cache level and the largest part of the central storage level, the object storage systems will

[23] For more information see, `https://galaxyproject.org`
[24] For more information see, `https://www.unicore.eu`
[25] For more information see, `https://www.nextflow.io`

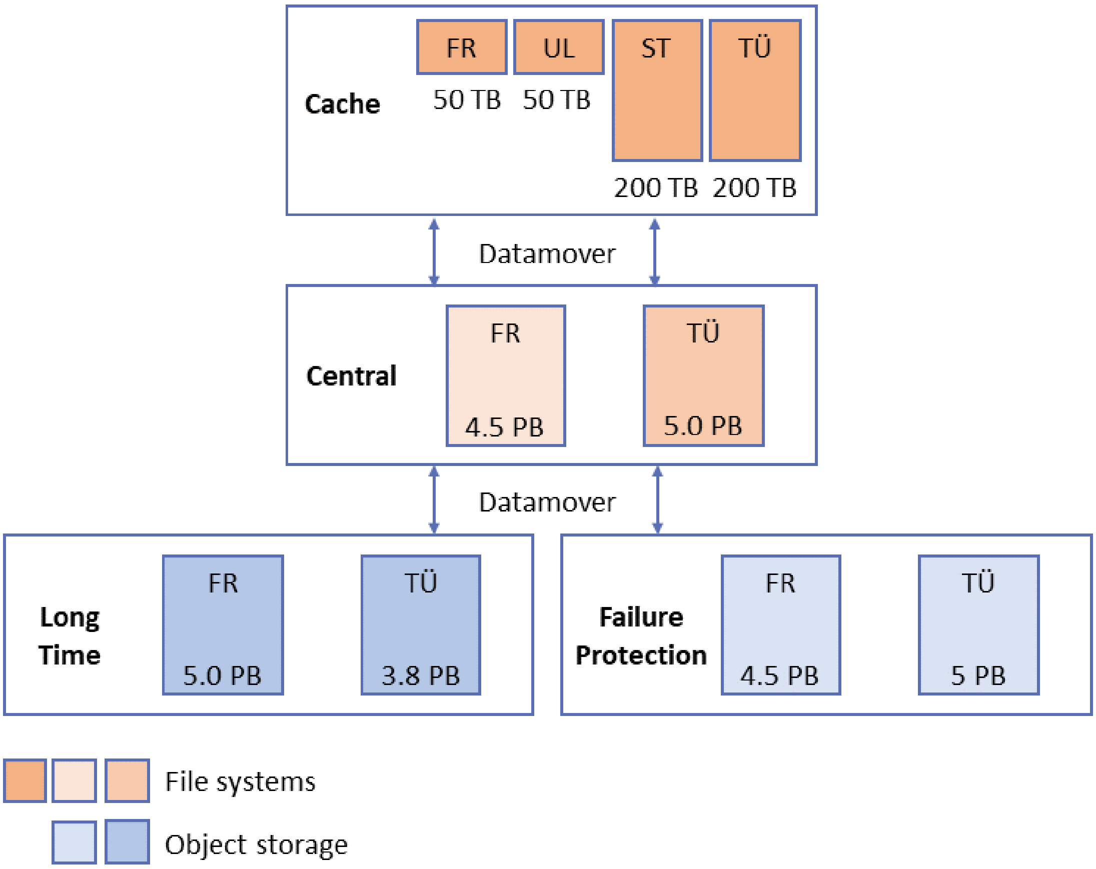

Figure 4.: Planned hierarchy levels of the distributed federated Storage-For-Science system including the anticipated size of the individual components.

be used on the long-term storage and data protection levels. The integration of object storage into SFS will increase the flexibility of data protection and data archiving, as provisioning of storage resources with dedicated levels of data protection and wide area geo-replication are already implemented. Software components such as the data movers will provide transparent access to data for researchers on all levels and additional RDM functionality. The integration of existing RDM solutions as well as interfaces for external and community repositories will provide further benefits for research. The SFS system user management needs to be only loosely coupled to the individual local IDM systems of the involved sites. It will adapt concepts from established AAI federations to ensure the usability across all participating universities, institutions and support widely cooperating scientists.

Similar to the „Central Application Site for bwHPC“[26] a self-service interface is envisioned. It would serve multiple purposes, like allowing project managers to declare their storage needs and provide information on their project, while storage administrators will also get informed about their users' needs. The descriptive and administrative metadata getting collected helps both to plan for storage volumes required and provide information

[26] For more information see, `https://www.bwhpc.de/en/zas_info_bwforcluster.php`

to university research information systems or grant providing agencies. Typically such information would include project title, containing additionally organizational names and persons involved, contact information, research grant and a short project summary to provide a context for the data, e.g. a description from a relevant ontology, an abstract of a related or planned paper, or excerpts of a relevant proposal. Further information should include the type (e.g. „hot“ or „cold“ data; file system, object storage, repository, long-term archive), redundancy expectations and the expected capacity and duration for storing the data. Especially the latter information would help for long-term planning and development of the system as well to plan for refinancing.

9. Conclusion

In the process of grant application and procurement for the federated SFS essential core characteristics and abstract features were identified together with the served scientific communities. The system will offer a good compromise of price per terabyte, performance and capacity. It features both traditional file systems as well as object storage to address the various needs of the scientists. It provides various levels of defined services including geo-redundancy. For trusted, reproducible scientific workflows the well-annotated, archived data will be immutable. It allows the automation of workflows by providing appropriate interfaces like REST APIs allowing asynchronous operation. The SFS system implements an identity mapping for users adapting concepts from established AAI federations. It thus abstracts from site-specific identity management to support the linking of long-term object identifiers to data owners and their individual requirements like embargo or retention periods. The SFS will handle the core project descriptive, administrative and technical metadata leaving the freedom to use specific scientific metadata to the respective communities.

The joint procurement and later operation of the system deepens the cooperation between the involved computing centers and communities. The ongoing process generates insights to be used for general application of research data management in the involved universities. The expertise is shared within the context of the data federation in Baden-Württemberg and the bwHPC-S5 project [3].

Acknowledgements

The research infrastructures described in this publication are part of the bwHPC project sponsored by the Ministry of Science, Research and the Arts, Baden-Württemberg (MWK) and the German Science Foundation (INST 37/935-1 FUGG and INST 39/963-1 FUGG). The same applies to the work presented from the CiTAR and ViCE projects which got supported by the MWK as part of it’s eScience initiative. Part of the work presented here was also supported through the BMBF funded project de.NBI (031 A 534A). The support is gratefully acknowledged.

Bibliography

[1] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Sant os, Philip E. Bourne, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3, 2016. doi:10.1038/sdata.2016.18.

[2] Gerhard Schneider, Vincent Heuveline, Karl-Wilhelm Horstmann, Bernhard Neumair, Petra Hätscher, Josef Kolbitsch, Simone Rehm, Michael Resch, Thomas Walter, Stefan Wesner, and Peter Castellaz. Umsetzungskonzept der Universitäten des Landes Baden-Württemberg für das High Performance Computing (HPC), Data Intensive Computing (DIC) und Large Scale Scientific Data Management (LS^{2}DM). In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors, *Proceedings of the 5th bwHPC Symposium*, pages 3–16. TLP, Tübingen, 2019. doi:10.15496/publikation-29040.

[3] Robert Barthel and Jürgen Salk. bwHPC-S5: Scientific Simulation and Storage Support Services – Unterstützung von Wissenschaft und Forschung beim leistungsstarken und datenintensiven Rechnen sowie großskaligem Forschungsdatenmanagement. In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors, *Proceedings of the 5th bwHPC Symposium*, pages 3–16. TLP, Tübingen, 2019. doi:10.15496/publikation-29039.

[4] Thomas Eifert, Ulrich Schilling, Hans-Jörg Bauer, Florian Krämer, and Ania Lopez. Infrastructure for Research Data Management as a Cross-University Project. In *International Conference on Human Interface and the Management of Information*, pages 493–502. Springer, 2017.

[5] Dennis Wehrle, Bernd Wiebelt, and Dirk von Suchodoletz. Design eines FDM-fähigen Speichersystems. In *10. DFN-Forum Kommunikationstechnologien, 30.-31. Mai 2017, Berlin*, pages 115–124. Gesellschaft für Informatik eV (GI), 2017. `https://dl.gi.de/bitstream/handle/20.500.12116/470/paper10.pdf`.

[6] Dirk von Suchodoletz, Ulrich Hahn, Bernd Wiebelt, Kolja Glogowski, and Mark Seifert. Storage infrastructures to support advanced scientific workflows. In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors, *Proceedings of the 5th bwHPC Symposium*, pages 263–279. TLP, Tübingen, 2019. doi:10.15496/publikation-29058.

[7] Andrew Treloar. The Research Data Alliance: globally co-ordinated action against barriers to data publishing and sharing. *Learned Publishing*, 27(5):S9–S13, 2014.

[8] Peter Doorn and Heiko Tjalsma. Introduction: archiving research data. *Archival science*, 7(1):1–20, 2007.

[9] Anna S. Palaiologk, Anastasios A. Economides, Heiko D. Tjalsma, and Laurents B. Sesink. An activity-based costing model for long-term preservation and dissemination of digital research data: the case of DANS. *International journal on digital libraries*, 12(4):195–214, 2012.

[10] Łukasz Dutka, Michal Wrzeszcz, Tomasz Lichoń, Rafał Słota, Konrad Zemek, Krzysztof Trzepla, Łukasz Opioła, Renata Słota and Jacek Kitowski. Onedata – A Step Forward towards Globalization of Data Access for Computing Infrastructures. *International Conference On Computational Science, ICCS 2015*, 51:2843–2847, 2015. doi:10.1016/j.procs.2015.05.445.

[11] Sandra Gesing, Nancy Wilkins-Diehr, Maytal Dahan, Katherine Lawrence, Michael Zentner, Marlon Pierce, Linda Hayden and Suresh Marru. Science Gateways: The Long Road to the Birth of an Institute. *Hawaii International Conference on System Sciences (HICSS)*, 2017. doi:10.24251/HICSS.2017.755.

[12] Stephen Crouch, Neil Chue Hong, Simon Hettrick, Mike Jackson, Aleksandra Pawlik, Shoaib Sufi, Les Carr, David De Roure, Carole Goble and Mark Parsons. The Software Sustainability Institute: Changing Research Software Attitudes and Practices. *Computing in Science and Engineering*, 15(6):74–80, 2013. doi:10.1109/MCSE.2013.133.

[13] Sarah Berenji Ardestani, Carl Johan Hakansson, Erwin Laure, Ilja Livenson, Pavel Stranák, Emanuel Dima, Dennis Blommesteijn, and Mark van de Sanden. B2share: An open escience data sharing platform. In *2015 IEEE 11th International Conference on e-Science (e-Science)*, pages 448–453. IEEE, 2015.

[14] Damien Lecarpentier, Peter Wittenburg, Willem Elbers, Alberto Michelini, Riam Kanso, Peter Coveney, and Rob Baxter. EUDAT: a new cross-disciplinary data infrastructure for science. *International Journal of Digital Curation*, 8(1):279–287, 2013. doi:10.2218/ijdc.v8i1.260.

[15] Rat für Informationsinfrastrukturen. Leistung aus Vielfalt. Empfehlungen zu Strukturen, Prozessen und Finanzierung des Forschungsdatenmanagements in Deutschland. `http://www.rfii.de/download/rfii-empfehlungen-2016`, 2016.

[16] Benjamin Gläßle, Maximilian Hanussek, Felix Bartusch, Volker Lutz, Ulrich Hahn, Werner Dilling, Thomas Walter, and Jens Krüger. de.NBI Cloud Storage Tübingen. In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors, *Proceedings of the 5th bwHPC Symposium*, pages 201–215. TLP, Tübingen, 2019. doi:10.15496/publikation-29054.

[17] Konrad Meier, Björn Grüning, Clemens Blank, Michael Janczyk, and Dirk von Suchodoletz. Virtualisierte wissenschaftliche Forschungsumgebungen und die zukünftige Rolle der Rechenzentren. In *10. DFN-Forum Kommunikationstechnologien, 30.-31. Mai 2017, Berlin, Gesellschaft für Informatik eV (GI)*, pages 145–154, 2017. `https://dl.gi.de/bitstream/handle/20.500.12116/473/paper13.pdf`.

[18] Jonathan Bauer, Dirk von Suchodoletz, Jeannette Vollmer, and Helena Rasche. Game of Templates. In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors, *Proceedings of the 5th bwHPC Symposium*, pages 245–262. TLP, Tübingen, 2019. doi:10.15496/publikation-29057.

[19] Deutsche Forschungsgemeinschaft (DFG). Leitlinien zum Umgang mit Forschungsdaten. `https://www.dfg.de/download/pdf/foerderung/antragstellung/forschungsdaten/richtlinien_forschungsdaten.pdf`, September 2015.

[20] Bundesministerium für Bildung und Forschung (BMBF). Leitlinien für das neue EU-Rahmenprogramm für Forschung und Innovation. `https://www.bmbf.de/files/Bundesregierung_FP9_Leitlinienpapier_September_2017.pdf`, September 2017.

[21] Hochschulrektorenkonferenz (HRK). Management von Forschungsdaten – eine zentrale strategische Herausforderung für Hochschulleitungen – Empfehlung der 16. HRK-Mitgliederversammlung. `https://www.hrk.de/positionen/beschluss/detail/management-von-forschungsdaten-eine-zentrale-strategische-herausforderung-fuer-hochschulleitungen`, May 2014.

[22] Hochschulrektorenkonferenz (HRK). Wie Hochschulleitungen die Entwicklung des Forschungsdatenmanagements steuern können. Orientierungspfade, Handlungsoptionen, Szenarien – Empfehlung der 19. HRK-Mitgliederversammlung. `https://www.hrk.de/positionen/beschluss/detail/wie-hochschulleitungen-die-entwicklung-des-forschungsdatenmanagements-steuern-koennen-orientierungsp`, November 2015.

[23] Alex Ball. *Review of data management lifecycle models*. University of Bath, IDMRC, 2012. `https://purehost.bath.ac.uk/ws/portalfiles/portal/206543/redm1rep120110ab10.pdf`.

[24] `http://chittagongit.com/icon/webpage-icon-10.html`.

[25] `http://chittagongit.com/icon/online-form-icon-11.html`.

[26] `http://chittagongit.com/icon/website-search-icon-17.html`.

[27] Gregor Cresnar. `https://www.flaticon.com/free-icon/database_159252#term=database&page=1&position=2`. Icon made by [Gregor Cresnar] from `www.flaticon.com`.

[28] `https://www.openarchives.org/images/OA200.gif`.

[29] Linda Cantara. METS: The metadata encoding and transmission standard. *Cataloging & classification quarterly*, 40(3-4):237–253, 2005.

[30] PREMIS Working Group and others. PREservation Metadata: Implementation Strategies, 2004.

[31] Herbert Van de Sompel, Michael L. Nelson, Carl Lagoze, and Simeon Warner. Resource harvesting within the OAI-PMH framework. *D-lib magazine*, 10(12), 2004.

[32] Christopher B. Hauser and Jörg Domaschka. ViCE Registry : An Image Registry for Virtual Collaborative Environments. In *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 82–89, 2017. doi:10.1109/CloudCom.2017.11.

[33] Martin Barisits, Thomas Beermann, Frank Berghaus, Brian Bockelman, Joaquin Bogado, David Cameron, Dimitrios Christidis, Diego Ciangottini, Gancho Dimitrov, Markus Elsing, Vincent Garonne, Alessandro di Girolamo, Luc Goossens, Wen Guan, Jaroslav Guenther, Tomas Javurek, Dietmar Kuhn, Mario Lassnig, Fernando Lopez, Nicolo Magini, Angelos Molfetas, Armin Nairz, Farid Ould-Saada, Stefan Prenner, Cedric Serfon, Graeme Stewart, Eric Vaand ering, Petya Vasileva, Ralph Vigne, and Tobias Wegner. Rucio – Scientific Data Management. *arXiv e-prints*, page arXiv:1902.09857, February 2019. `https://arxiv.org/abs/1902.09857`.

[34] FAIR principles for data stewardship. *Nature Genetics*, 48(4):343–343, 2016. doi:10.1038/ng.3544.

[35] Heather A. Piwowar, Roger S. Day, and Douglas B. Fridsma. Sharing detailed research data is associated with increased citation rate. *PLoS ONE*, 2(3), 2007. doi:10.1371/journal.pone.0000308.

[36] Daniel Garijo, Sarah Kinnings, Li Xie, Lei Xie, Yinliang Zhang, Philip E. Bourne, and Yolanda Gil. Quantifying reproducibility in computational biology: The case of the tuberculosis drugome. *PLoS ONE*, 8(11):1–11, 2013. doi:10.1371/journal.pone.0080278.

[37] Jun Zhao, Jose Manuel Gomez-Perez, Khalid Belhajjame, Graham Klyne, Esteban Garcia-Cuesta, Aleix Garrido, Kristina Hettne, Marco Roos, David de Roure, and Carole Goble. Why Workflows Break – Understanding and Combating Decay in Taverna Workflows. In *2012 IEEE 8th International Conference on E-Science*, pages 1–9, 2012. doi:10.1109/eScience.2012.6404482.

[38] Gregory M. Kurtzer, Vanessa Sochat, and Michael W. Bauer. Singularity: Scientific containers for mobility of compute. *PLOS ONE*, 12(5):1–20, 2017.

[39] Carl Boettiger. An introduction to Docker for reproducible research. *SIGOPS Operating Systems Review*, 49(1):71–79, 2015. doi:10.1145/2723872.2723882.

[40] Felix Bartusch, Maximilian Hanussek, and Jens Krüger. Containerization of Galaxy Workflows increases Reproducibility. In *Proceedings of the bwHPC Symposium*, pages 16–19, 2017. doi:10.15496/publikation-25200.

[41] Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt. bwForCluster NEMO. In Michael Janczyk, Dirk von Suchodoletz, and Bernd Wiebelt, editors,

Proceedings of the 5^{th} bwHPC Symposium, pages 29–50. TLP, Tübingen, 2019. doi:10.15496/publikation-29041.

[42] Felix Bührer, Anton J. Gamel, Benoît Roland, Benjamin Rottler, Markus Schumacher, and Ulrike Schnoor. Integration of NEMO into an existing particle physics environment through virtualization. In *Proceedings of the 5^{th} bwHPC Symposium*, pages 187–200. TLP, Tübingen, 2019. doi:10.15496/publikation-29053.

[43] Felix Bührer, Frank Fischer, Georg Fleig, Anton Gamel, Manuel Giffels, Thomas Hauth, Michael Janczyk, Konrad Meier, Günter Quast, Benoît Roland, Ulrike Schnoor, Markus Schumacher, Dirk von Suchodoletz, and Bernd Wiebelt. Dynamic Virtualized Deployment of Particle Physics Environments on a High Performance Computing Cluster. *Computing and Software for Big Science*, 2018.

[44] Jens Krüger, Volker Lutz, Felix Bartusch, Werner Dilling, Anna Gorska, Christoph Schäfer, and Thomas Walter. Bioinformatics and Astrophysics Cluster (BinAC). In Sabine Richling, Martin Baumann, and Vincent Heuveline editors, *Proceedings of the 3^{rd} bwHPC Symposium*, pages 91–95. Heidelberg: heiBOOKS, 2017. doi:https://doi.org/10.11588/heibooks.308.418.

[45] Christoph Heidecker, Matthias J. Schnepf, Florian von Cube, Manuel Giffels, and Günter Quast. Dynamic Resource Extension for Data Intensive Computing with Specialized Software Environments on HPC systems. In *Proceedings of the 5^{th} bwHPC Symposium*, pages 161–172. TLP, Tübingen, 2019. doi:10.15496/publikation-29051.