EPiC Series in Computing

Volume 105, 2025, Pages 69–79

Proceedings of EUNIS 2024 annual congress in Athens

Migrating a Federated Educational Infrastructure to the Cloud: Lessons Learned from a National Project

Hendrik Geßner[1], Dirk Bußler[1], Duy Nguyen[2], Holger Zimmermann[3] and Ulrike Lucke[4]

[1] Center for IT and Media Management, University of Potsdam, Germany
hendrik.gessner@uni-potsdam.de
dirk.bussler@uni-potsdam.de
[2] g.a.s.t., Germany
duy.nguyen@gast.de
[3] snoopmedia, Germany
h.zimmermann@snoopmedia.com
[4] Department of Computer Science, University of Potsdam, Germany
ulrike.lucke@uni-potsdam.de

Abstract

The migration of IT services to a cloud infrastructure is complex - depending on the existing service landscape - and requires special awareness. Recommendations and guidelines can help, but must be adapted to the respective context. This article focuses on an example that is research-related and is distributed across institutional boundaries: the prototype of a national networked infrastructure for education currently being developed in Germany. We determine the associated expectations of various stakeholder groups regarding a migration to the cloud, compare these with the current state of the literature and contrast them with our practical experiences gained during the migration. This identifies gaps in existing recommendations that need to be addressed in the future.

1 Background and Driving Forces for Cloud Migration

Outsourcing the operation of IT systems to the cloud (or at least working in cloud-like infrastructures) is an established approach for providing IT resources efficiently and being able to react quickly to changing requirements (Bauer, 2023). There is also hope that the available portfolio of integrable services will make software development in the cloud easier and more advanced (Caasen, 2021). Software development is changing from comprehensive solutions to modular services with distributed states that communicate flexibly via defined interfaces; it is therefore necessary to keep an eye on cloud telemetry data in order to detect anomalies at an early stage and respond to them (Lenz, 2021). Numerous articles from the relevant professional literature report on experiences (Wallmark et al., 2015; Eifert et al., 2022) or formulate recommendations on how to successfully migrate from an on-site data center to a cloud infrastructure (Bellavista et al., 2021) or on necessary accompanying training measures (Cloud Community Group, 2021). Special reports and guidelines are also available for dedicated universities' services in research, education and administration (van Gaver et al., 2016; Eifert et al., 2021; Kasprzak, 2023). So far, however, these recommendations have concentrated on the institutional perspective (certain company or university) or on certain services (such as databases or backup) in productive mode, while distributed or federated IT landscapes currently being developed in research and their necessarily more complex requirements structure have not yet been reflected.

During the development of the prototype for a nationwide networked infrastructure for education (Knoth et al., 2022), which aims to cover all areas from primary to tertiary education, vocational training and life-long learning, the challenge was to gradually migrate the services of the distributed middleware and the access portal, which were initially hosted by the project partners, to the cloud (Lucke et al., 2023). The reasons for this were initially the need to hand over the project results to the funding German Federal Ministry for Education and Research or its commissioned development teams after the end of the funding period in order to ensure sustainable availability. In addition, however, a gradual handover of individual components or shared use of existing services between the two sides was also required during the project period, which made it necessary to closely synchronize both the development cycles and the technical infrastructure. The solution provided in the prototype was hardened into a product by the developers commissioned by the ministry, i.e. re-engineered and made scalable, robust and secure; however, fundamental decisions regarding functionality, architecture or technologies were not to be revised in the process. The cloud infrastructure thus serves both as a point of delivery and as a common working basis.

The prototype involved a variety of services (from basic IT services to complex communication and collaboration tools) along the entire education chain and across different institutions. Technically speaking, our infrastructure at that time consisted of a Kubernetes cluster (K8S) based on virtual machines with Docker containers for the individual services as well as CI/CD pipelines based on local GitLab; including load balancing, firewall and monitoring. The target infrastructure was to be a managed container infrastructure with the corresponding security and control features (self-service). As it turned out, further deviations were CI/CD pipelines based on GitHub with Helm Charts and proprietary storage connections, while no off-the-shelf components for load balancing, firewall and monitoring were available. It was therefore a change across several technologies and technology generations. Thus, the question remained open as to what extent the previous experiences and recommendations collected during the literature review could be transferred to such a national infrastructure.

The remainder of this article is structured as follows. In Chapter 2, we provide an overview of the stakeholder expectations (from the political, application and technical domain) and analyze to which extent the initial hypothesis is valid that our migration project is associated with different requirements than others. In chapter 3, we then describe the experiences gained from the cloud migration process in our project, and contrast these experiences with the previously collected expectations. Finally, chapter 4 transforms our lessons learned from the cloud migration process to recommendations for other projects, thus extending existing guidelines beyond the institutional realm.

2 Analyzing the Expectations

We supplement the intuitive expectations of project sponsors regarding professional IT management in the cloud with a systematic survey (section 2.1). This includes perspectives of the application domain (section 2.2) and IT professionals (section 2.3), generated from project-specific surveys and literature research. The results of our analysis are summarized in section 2.4.

2.1 From Intuitive Estimate to Systematic Analysis

From the very beginning of the project, it became apparent during regular board meetings that the political stakeholders intuitively ascribed more professionalism in the operation of a scalable IT infrastructure to commercial cloud providers than to university data centers. This assessment was apparently also influenced by their desire for a transparent cost structure and simple handling after the end of the project period, including a medium to long-term operational and governance concept. In addition, a complex procurement procedure was to be expected. Both the federal ministry and the project stakeholders had to tender for their own parts of a common cloud infrastructure at the same time, with partly differing legal procurement requirements. Moreover, strict technical requirements (e.g. security and privacy level, in this case C5 certification; see BSI, 2020) had to be observed. This resulted in a very small number of potential providers at that time.

In order to delve deeper into the expectations associated with a migration to the cloud, we used a systematic approach to collect relevant points. These expectations were formulated as qualitative statements describing potential benefits or impacts of such a migration, such as "high scalability" or "savings in IT staff". Three main sources were used to collect data: an end-user oriented survey within the project network (N=21), a survey of the IT staff involved (N=4) and a literature review that took into account both specialist literature and field reports (as sketched in chapter 1). Our aim was to cover a broad spectrum of perspectives.

We have compiled the statements from these sources in a list. Only the mention of an aspect was decisive, not its valence. For example, "reduced costs in the cloud" and "increased costs in the cloud" were counted equally as "costs in the cloud". The collected expectations were then collated into groups to form thematic categories reflecting the various aspects of a cloud migration. In an iterative process, 18 thematic clusters were ultimately formed. Each individual source (e.g. an interview or a publication) was counted at most once per cluster. This allowed us to systematically analyze the expectations and identify similarities and differences between the various data sources. By summing up the expectations per data source, we were also able to determine the relative weightings and importance of the various sources for the overall assessment of a cloud migration. It should be noted that due to the small N, this count does not quantify the absolute importance of an aspect, but merely allows us to give an approximate weighting on a project-specific basis.

2.2 Users' Perspective

During an on-site workshop, we asked the members of the project consortium (with backgrounds in education, design, computer science and management) to write down their expectations of a migration to the cloud. The participants then discussed the identified aspects in pairs and were given space to formulate additional expectations.

The respondents most frequently expressed expectations regarding the following points: "decentralized access management", "scalability", "technical flexibility", "workload reduction", "support" and "migration". These topics were mentioned by 45% to 40% of internal stakeholders. We assume that these expectations are due to points of friction in the developers' daily work. The infrastructure in place at that time was based on virtual machines to which only administrators had access. There were no centralized protocols and new infrastructure was set up manually by the administrators. For the rollout of containers, separate CI/CD pipelines were created by the administrators for each container.

Expectations on the topics of "value creation", "impact on existing methods" and "prevention of shadow IT" were mentioned the least, with the proportion ranging between 5% and 0%. Our interpretation is that such terms are usually discussed at higher management levels, which were rarely represented among the participants of the workshop.

2.3 IT Professionals

The IT staff (some employed in a research group, some in the university data center and some contracted externally) are responsible for the project's infrastructure. These people jointly formulated their expectations of a migration to a cloud infrastructure in a workshop format and compared them with the experiences of the other stakeholders in a presentation to the project group. For three people, the use of a Kubernetes cloud hoster was new. Although some of them had previous knowledge of Kubernetes, they had no practical experience of it.

The IT professionals expressed the following expectations regarding a move to the cloud: "Cost reduction compared to VMs", "Reduced maintenance and administration effort for the infrastructure", "Acceleration of the software lifecycle, esp. for developers", "Consulting during the migration, consulting and optimization during ongoing operations by cloud hosters", "Migrating the VM infrastructure to Kubernetes as a service", "Scalability of the infrastructure", "Complete migration to Kubernetes within 6 months" and "Incremental move with hybrid use of old and new infrastructure".

2.4 Analysis of Aggregated Data

Figure 1 shows the relative frequency of citations, separated by main source and sorted by literature. The IT professionals and the political stakeholders are included as a block because of the small number of involved persons, so that the frequency here is either 0% or 100%.

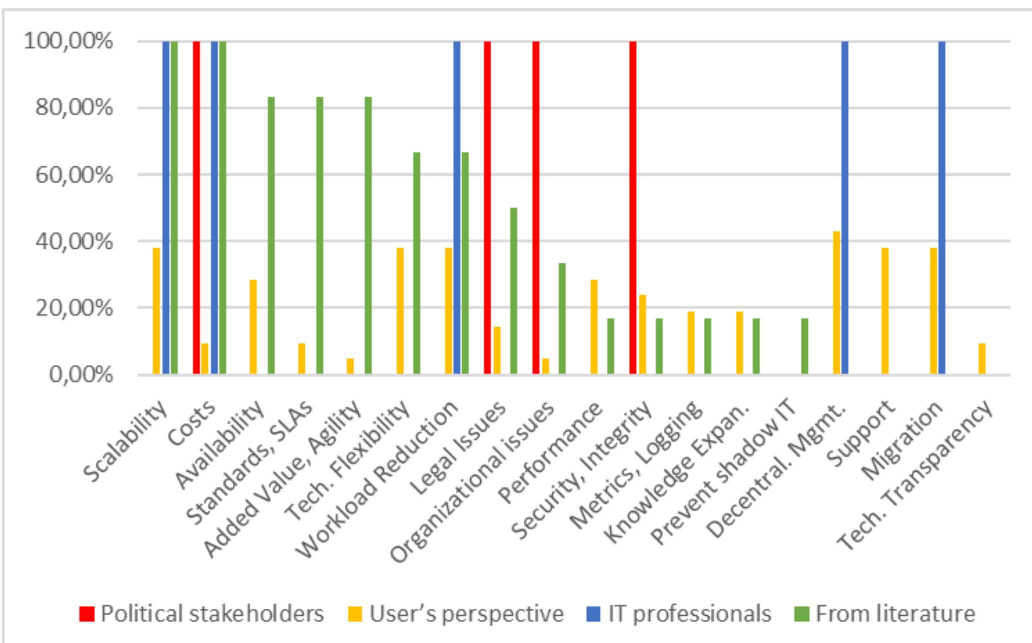

Figure 1: The expectations on cloud migration vary significantly among surveyed stakeholders.

It is apparent that the assessment of the political decision-makers relates more to the high-level management aspects and therefore does not take many other points into account. Also, it is noteworthy that the aspects of "scalability" and "costs" were considered important both by IT professionals and in all literature sources (100%), but were only mentioned by some of the project stakeholders (40%). The topics of "availability", "SLAs", "added value" and "technical flexibility" also only played a subordinate role compared to the literature (70-85%) within the project stakeholders (5-40%) and IT professionals (0%).

We explain the discrepancies in the assessments between project actors and IT professionals by their different understanding of the terms "availability" and "technical flexibility" in the context of a non-productive system. We attribute the remaining discrepancy with the literature to the fact that the prototype development in the project (unlike the product development commissioned by the ministry at the same time) is located in a research context, which is why there is no productive data or specific requirements for productive operation (like BSI, 2020) when using the cloud. On the contrary, the prioritization of expectations found in the literature seems to reflect the real-world requirements and challenges associated with the adoption of cloud services in production environments. These differences emphasize the importance of a comprehensive understanding of the specific contexts and needs of the parties involved when planning and implementing a cloud migration. In the subsequent analysis and discussion, it emerged that certain functions such as DDOS detection, web application firewall, load balancing and protocol-based attack pattern recognition were missing from the offer. Some of these essential security mechanisms had to be provided by the contractor at system level, as they could not be obtained directly from the cloud provider. This required individual implementation and integration of security functions in order to ensure the integrity and security of the system.

We were surprised to find that none of the literature sources formulated expectations on the topics of "support", "migration" and "decentralized access management", while these aspects played a significant role in our project. It should be noted that, similar to many productive systems, our prototype also involved migrating an existing infrastructure to the cloud. This finding suggests that the practical challenges and requirements associated with migrating existing systems to the cloud may not be adequately addressed in the current literature. We have two possible interpretations here: Either these topics are so individualized that no written expectations are expressed on this, or these topics are a blind spot within the literature. Both interpretations offer valuable insights into the discrepancy between the expectations found in the literature and the expectations formulated in our project. The first interpretation, that these topics are so individual that no general statements can be made about them, suggests that the complexity and diversity of the challenges related to support, migration and decentralized access management may lead to them not being adequately addressed in the literature. The second interpretation, that these topics represent a blind spot in the literature, highlights the need to investigate these aspects in more detail and contribute to the visibility of these important topics. By documenting and analyzing practical experiences and challenges in the context of cloud migrations, we can help to fill the gaps and provide a more comprehensive view of the requirements and best practices related to the adoption of cloud technologies in research projects.

The results of the analysis emphasize the importance of our work, which can help advance the discussion on these topics and sensitize the scientific community to potential blind spots.

3 Experiences Gained during the Migration Process

The project had an existing infrastructure that was modest compared to typical university IT, comprising around 20 virtual machines that had been in place for 1-3 years. Nevertheless, we repeatedly faced unexpected challenges during the migration to the cloud infrastructure, which meant that a number of schedules and milestones could not be met. During this time, we went through a classic process of problem-based learning (Hmelo-Silver, 2004), in which we continuously compared our experiences with our expectations, set new common goals and thus gradually came closer to the desired solution. Through group reflection, we were able to gain distance and transfer knowledge, which we then added to our existing corpus of recommendations.

Due to the small number of expectations on cloud migration in the literature, we will highlight certain findings and the pitfalls that arise from our project in the following. We have sorted the expectations presented in section 2 in ascending order of frequency in the literature, starting with the rarest and ending with the most frequently mentioned. The extent to which expectations could be met (–/o/+) is symbolically marked.

1. Complete migration to Kubernetes within 6 months (–).
   In April 2022, the project started with the initial consultancy for Kubernetes cloud hosting, in which a reference architecture for the overall system was derived from the technical requirements. From June 2023, a functional version of our rolled-out research services was available that offered comparable capabilities to the old infrastructure. The migration therefore took 15 months. An analysis of the delays revealed that the complicated procurement process accounted for the largest share of this duration (9 months), followed by the connection of the SSO login functionality (4 months), the provision of the cluster by the service provider (2 months) and the development of Helm charts by the IT staff (2 months). During the development of the Helm charts, it became clear that various assumptions regarding the interoperability of the cloud infrastructure had not been met. Neither S3 nor NFS were available as storage technologies, which is why the team had to find and implement a workaround (1.5 months). In the background, the hoster also switched from T-Cloud to Google. In addition, there was no managed firewall on the cluster, which is why firewall capacities had to be added by IT staff via the Helm charts (0.75 months). *We would therefore like to emphasize the importance of allowing sufficient time for the conceptual steps (technology comparison, provider comparison, contract and data protection review, payment modalities).*

2. Migration of the VM infrastructure to Kubernetes as a service (–).
   The IT staff expected that the migration of the VM infrastructure to Kubernetes would be carried out as a service by the cloud hoster. During the onboarding phase, however, it became clear that the contract was imprecisely phrased in this area and that the contractually guaranteed relocation support was so concise that the cloud provider only had an advisory role. The structure of the Helm charts was agreed between the IT staff and the cloud hoster, whereby the implementation of the Helm charts was the responsibility of the IT staff in the project, who came into practical contact with Kubernetes for the first time. Although the IT staff had prepared for the migration in advance, they lacked sufficient experience with Kubernetes, which delayed the implementation of the Helm charts significantly. In addition, the local K8S test installation and the cloud K8S version differed in several areas, for example in terms of available storage mechanisms and controllers, which further hindered the implementation of the prepared solution. *We would therefore recommend that clear acceptance criteria are agreed for all contract areas and that all management and worker levels check these criteria before the contract is signed.*

3. Incremental migration with hybrid use of old and new infrastructure (–).
   During the planning phase, it was agreed that the move to the cloud would be incremental. The following migration steps were planned: (1) deployment to the Kubernetes cluster, (2) migration to a hosted container registry and (3) migration from GitLab to GitHub, including the attached CI/CD pipelines. The goal was for IT staff to quickly share the benefits of the new cloud environment with other teams. However, when the cluster was ready, access was only possible from the preconfigured GitHub Actions, and reconfiguration was ruled out by the

hoster. Step (3) was therefore mandatory in order to achieve steps (1) and (2). The incremental move could therefore not be implemented, the move had to take place in a block and results were only available to other teams once all components had been moved. *Based on this experience, we would recommend considering the results of preliminary discussions as binding, recording them in writing and only allowing replanning with the mutual agreement of both parties.*

4. Advice on migration; advice and optimization during operation by the cloud hoster (–).
   A K8S architecture was developed as part of the preliminary consultation for the contract. It turned out that the calculated resources (CPU and RAM) were significantly overprovisioned and therefore caused excessive running costs. In addition, the architecture lacked a suitable firewall, and a virus scanner that was added very late did not offer the possibility of forwarding reports to the IT team. We had expected that such anomalies would be addressed by the service provider as part of the regular support and consulting work. During operation, the support team was available but often remained superficial in its responses, which resulted in an increased research effort on the part of the IT team. *Based on our experience, we would therefore recommend clarifying expectations of consulting and support in advance. This includes, in particular, areas of consultancy as well as the amount and scope of consulting in ongoing operations.*

5. Decentralized Management (+).
   The use of Kubernetes accelerates the software life cycle for developers in particular. Thanks to faster deployment processes, features can be tested quickly in several harmonized environments. The cloud offering also enables independent access to log files, which speeds up the debugging of errors - as long as the log environment is provided by the cloud hoster. The platform also allows the simple implementation of complex deployment structures, such as Rasa (Bocklisch et. al., 2017), and therefore offers a flexible deployment solution. With the help of Helm charts, developers and administrators can better coordinate and divide their work. However, the cloud platform has its limits when it comes to user management: Accounts and rights are managed by the service provider and transferred via Excel spreadsheets. *Based on our experience, we recommend placing great value on decentralized access management. In addition, the provision of centralized developer tools, such as for logging and deployments, has proven itself in our practice.*

6. Professional organization of external sub-contractors (o).
   There was a pronounced expectation, particularly on the political side, but also to some extent on the user side, that external service providers would be associated with a higher degree of professionalism than accessing the services of local data centers in universities. The IT staff also had the expectation that an external service provider would be far ahead of them in terms of working methods and that, based on their own preliminary work, cooperation with or advice from external players would lead to further professionalization in the data center. However, the experiences described above showed that a high degree of professionalism already existed in the IT departments of the universities. This is not surprising given the operation of federated services for several hundred thousand users, in some cases across the borders of federal states. *We therefore recommend that stakeholders in university data centers should be more self-confident in dealing with both public authorities and external service providers and openly demonstrate their existing expertise.*

7. Decreasing Workload Reduction (+).

In terms of personnel resources, we currently require one full-time position for administration and an external service provider with 30 hours per month. We assume that local administration support for cluster, monitoring, logging, databases, provision of specialized GitHub actions, storage, networks, firewall and associated updates would require an additional full-time position. Our expectation of reduced workload was met. *Based on our experience, we recommend opting for included managed components wherever possible.*

8. Decreasing Infrastructure Costs (–).
   Reduced costs are one of the most common expectations when moving to a hosted cloud offering. These costs are usually made up of two items: non-personnel resources and personnel resources. A few months after the migration, our monthly costs for cloud hosting amount to 4500€. We compare these costs with the previous costs for on-site infrastructure. For the non-personnel resources, we have determined that the same performance in the university data center on virtual machines would cost around €850. An analysis of cloud costs has shown that a combination of heavily overprovisioned resources (CPU and RAM) and a reserved workload model is driving up costs. By adjusting the limits, we expect a cost reduction of 50%. However, even with this adjustment to the limits, our costs would still be 260% of the VM costs. *Based on our experience, we recommend agreeing a consulting mode for cost analysis and cost optimization between the cloud hoster and the client at the start of the contract. We suggest arranging an appointment during the migration phase, another one after completion of the migration and then regular appointments, for example every six months.*

9. Scalability of the Infrastructure (–).
   To make the most of the advantages of the cloud, it is important to scale resources based on the current workload and to create them dynamically. We are already benefiting from the advantages that Kubernetes offers. However, we are currently constrained by the reserved workload model, which means we have fixed CPU and RAM ressources and are limited in both scaling down and scaling up. *Our recommendation is therefore that these advantages should be passed on to the customer by the service provider.*

Despite the fact that the degree of realization of the existing expectations with only two fulfilled and one partially fulfilled aspect may at a first glance paint a sobering picture, the learning effect from the migration project is significant. The expertise gained on the added value that can actually be achieved, the prerequisites required for this and the associated side effects is of great value for future IT projects. The devil is literally in the technical details here. In this respect, when it comes to large infrastructures, it is important to carefully consider whether the cloud expertise that needs to be built up in-house in advance outweighs the achievable added value - or whether it would be more expedient to build an own cloud infrastructure on-site.

4 Lessons Learned and Recommendations

Reflecting on our expectations for migrating to the cloud, we have to summarize that six out of nine expectations were *not* met. Time goals were significantly exceeded, migration had to be carried out by existing IT personnel, the migration plan was not feasible, consultation by the cloud hoster did not meet our expectations, and the scalability is greatly limited due to the Reserved Workload model. Two other aspects, decentralized access to developer tools and reduced workload after migration, were fulfilled. An increase in professionalism at the data center was not observed. Although the costs are not yet in line with our expectations, we are confident that we will be able to make progress here in the future. Figure 2 provides a visual summary.

Figure 2: Fulfillment of project-specific expectations from the cloud migration at a glance.

Based on the retrospective alignment between our previous expectations and actual observations, we can summarize our recommendations as follows:

- Allow sufficient time for the conceptual steps (technology comparison, provider comparison, contract and data protection review, payment modalities).
- Define clear acceptance criteria for all contract areas and ensure that all management and worker levels check these criteria before the contract is signed.
- Record the results of preliminary discussions in writing, agree to them as binding, and only allow replanning with the mutual agreement of both parties.
- Clarify expectations of consulting and support in advance and precisely. This includes, in particular, the areas of consultancy as well as the amount and scope of consulting in ongoing operations.
- Place great value on decentralized access management. Ensure the provision of centralized developer tools, such as for logging and deployments.
- Opt for included managed components wherever possible.
- Agree a consulting mode for cost analysis and cost optimization between the cloud hoster and the client at the start of the contract. We suggest arranging an appointment during the migration phase, another one after completion of the migration and then regular appointments, for example every six months.
- Ensure that dynamic resources management is available as self-service.
- As stakeholders in university data centers, be more self-confident in dealing with both public authorities and external service providers and openly demonstrate your existing expertise.

In addition, through reflection with members of the project consortium, we learned several points that were not covered in the initial expectations:

- When selecting a cloud solution, a mature standard product is necessary. This applies even to the migration of research prototypes, even if their demands on the degree of maturity of technical solutions may often be regarded as lower. Infrastructure that is still under development should preferably not be purchased, which is however not easy to assess on the customer side. Note that this also includes products from experienced standard providers.
- Building in-house expertise is essential, even if a third-party provider takes on large portions of the work. Where in-house expertise is still lacking, it needs to be built up, since otherwise neither the specification of contract terms nor the monitoring of their compliance is feasible.
- Caution should be exercised with intermediary layers introduced by service providers managing access to cloud offerings and implementing their own management solutions. Here, we would instead recommend purchasing a cloud solution directly and involving an external service provider in an advisory and management capacity.
- The perspective shift from service provider to service customer for large infrastructure environments was educational for the data center staff.

Retrospectively, the recommendations formulated above may seem obvious at first glance. However, in view of the wealth of existing guidelines, checklists, standards and experience reports, it is difficult to identify those points that can be applied to a specific migration project in advance. As shown in this article, there are some gaps in existing recommendations on migration aspects. In some cases, they are also not immediately transferable to other situations. It would therefore be helpful to contextualize existing recommendations (e.g. to what extent they were intended for research or productive systems, for smaller or larger, for homogeneous or heterogeneous, for institutional or federated ... systems) in order to better assess their transferability.

Acknowledgements

This work was financed by the European Union - NextGenerationEU through the German Federal Ministry of Education and Research (BMBF) as part of the Digital Education initiative under grant number 16NB001. We are deeply grateful to our partners in the project for the valuable exchange.

References

Bauer, G. (2023). Warum Datenbanken in die Cloud wandern. In: *Linux Magazin*, No. 01. https://www.linux-magazin.de/ausgaben/2023/01/cloud-db-grundlagen/

Bellavista, P.; Corradi, A.; Edmonds, A.; Foschini, L.; Zanni, A.; Bohnert, T.M. (2021). Elastic Provisioning of Stateful Telco Services in Mobile Cloud Networking. In: *IEEE Transactions on Services Computing*, 14/03, pp. 710-723. https://doi.org/10.1109/TSC.2018.2826003

Bocklisch, T.; Faulkner, J.; Pawlowski, N.; Nichol, A. (2017): Rasa: Open Source Language Understanding and Dialogue Management. https://arxiv.org/abs/1712.05181

BSI - Bundesamt für Sicherheit in der Informationstechnik (2020). Kriterienkatalog Cloud Computing C5. https://www.bsi.bund.de/DE/Themen/Unternehmen-und-Organisationen/Informationen-und-Empfehlungen/Empfehlungen-nach-Angriffszielen/Cloud-Computing/Kriterienkatalog-C5/kriterienkatalog-c5_node.html

Caasen, J. (2021). Software: Life in the Cloud. In: *windows.developer*, No. 10. https://entwickler.de/azure/life-in-the-cloud

Cloud Community Group (2021). Survey Results / Cloud Management report. E*uropean University Information Systems.* https://www.eunis.org/wp-content/uploads/2021/01/EUNIS_cloudcommunity_survey_2020_Results.pdf

Eifert, T.; Dittrich, D.; Gündogan, A. (2021). University's Core Business - How digitalization, cooperation, and cloud effect IT's value proposition and metrics. In: *Proceedings European University Information Systems* - EUNIS '21. EPiC Series in Computing 78, 41-48. https://doi.org/10.29007/t9bc

Eifert, T.: Filla, N. (2022). Co-operative Service Design for a Cross-University Data Backup Service. In: *Proceedings European University Information Systems* - EUNIS '22. https://www.eunis.org/download/2022/EUNIS_2022_paper_22.pdf

van Gaver, S.; Fiorentino, A.; Slitni, M.; Bedouin, T.; Alehyane, A. (2016). *A Very Real Cloud: Best PRactices for Cloud Adoption in Higher Education*. Paris Ile-de-France Digital University. https://eunis.org/wp-content/uploads/2016/05/Guide-Cloud_UK.pdf

Hmelo-Silver, C.E. (2004). Problem-Based Learning: What and How Do Students Learn?. *Educational Psychology Review*, 16, 235–266. https://doi.org/10.1023/B:EDPR.0000034022.16470.f3

Kasprzak, P. (2023). GWDG Trusted Research Environment (TRE): Empowering European Research through AWS Cloud Services. *EOSC In-Practice Story No. 32*. Zenodo. https://doi.org/10.5281/zenodo.8340057

Knoth, A.; Blum, F.; Soldo, E.; Lucke U. (2022). Structural Challenges in the Educational System Meet a Federated IT-Infrastructure for Education: Insights into a Real Lab. In: *Proceedings of the 14th International Conference on Computer Supported Education - CSEDU* (pp. 369-375). SciTePress. https://doi.org/10.5220/0011085800003182

Lenz, F. (2021). Architekturen für Cloud-Lösungen. In: *windows.developer*, No. 03. https://entwickler.de/software-architektur/architekturen-fur-cloud-losungen

Lucke, U.; Knoth, A.; Wilhelm-Weidner, A. (2023). Perspektiven von Wissenschaft und Praxis auf die digitale Vernetzungsinfrastruktur für die Bildung. In *eleed – e-learning and education*, Iss. 15. https://nbn-resolving.de/urn:nbn:de:0009-5-57929

Wallmark, M.; Berglund. S. (2015). Moving a Decentralized Public University to the Cloud. In: *Proceedings European University Information Systems* - EUNIS ‘15. https://www.eunis.org/download/2015/papers/EUNIS2015_submission_7.pdf