

# Towards a trustworthy, secure and reliable enclave for machine learning in a hospital setting: The Essen Medical Computing Platform (EMCP)

Hendrik F.R. Schmidt\*, Jörg Schlötterer\*<sup>†‡</sup>, Marcel Bargull\*, Enrico Nasca\*<sup>‡</sup>, Ryan Aydelott\*,  
Christin Seifert\*<sup>†‡§</sup>, Folker Meyer\*<sup>†§</sup>

\*Institute for Artificial Intelligence in Medicine, University Hospital Essen, Essen, Germany  
{firstname.lastname}@uk-essen.de

<sup>†</sup>University of Duisburg-Essen, Essen, Germany  
{firstname.lastname}@uni-due.de

<sup>‡</sup>Cancer Research Center Cologne Essen (CCCE), Essen, Germany

<sup>§</sup>Co-senior authors

**Abstract**—AI/Computing at scale is a difficult problem, especially in a health care setting. We outline the requirements, planning and implementation choices as well as the guiding principles that led to the implementation of our secure research computing enclave, the Essen Medical Computing Platform (EMCP), affiliated with a major German hospital. Compliance, data privacy and usability were the immutable requirements of the system. We will discuss the features of our computing enclave and we will provide our recipe for groups wishing to adopt a similar setup.<sup>1</sup>

**Index Terms**—scientific computing infrastructure, medical data

## I. INTRODUCTION

The use of scientific computing has great potential for health care; in fact currently researchers are exploring computing approaches to classify potential melanoma [1], detect breast cancer [2] and personalize treatment [3]. Machine Learning (ML) approaches are opening new venues for applied clinical research, while the advent of tooling from computer science is triggering a digital transformation of health care and health care-related research. Colocation of computational research labs and health care IT have fundamentally different technical requirements. This goes beyond traditional IT resources in health care and moreover does not allow for separation of concerns between clinical services and research. Health care IT with a myriad of different systems and many users of different skill sets cannot easily accommodate a computational research platform as they have very specific requirements along with vastly different user expectations.

Efficient computational research requires rapid evaluation of published software tools and libraries, as well as scripting on top of a heterogeneous, often changing set of dependencies. In contrast, clinical environments are not typically supportive of individual researchers installing software without the need for documentation and/or management or IT staff signoff. This is the result of the fact that health care data (particularly

patient data) is protected by a number of laws and regulations in Germany. To comply with these laws, a significant number of technical measures must be taken to insure that access is controlled and data on the system is protected. It is not an understatement to declare that hospitals take IT security and the rule of parsimonious access quite seriously.

While hospitals in most developed countries are IT rich environments with an abundance of systems, it is useful to point out that those IT infrastructures were not designed to facilitate research computing, let alone resource-intensive computing. More often than not, the research environment is a researcher’s desktop system, which typically has the least amount of restrictions and access controls and is rarely the subject of a professional audit. A department conducting large scale Machine Learning (ML) or Deep Learning (DL) experiments or resource-intensive computing clearly is not a good fit for a typical hospital. Even the largest research hospital can struggle to comply with functional requirements, access controls and audit procedures for a system focused on research, as they differ vastly both in purpose and structure from the systems that traditional IT health care playbooks were written for.

Research computing often requires access to resources significant in both quantity and novelty. This includes Graphics Processing Units (GPUs), Field Programmable Gate Arrays (FPGAs) and Application-Specific Integrated Circuits (ASICs) (including custom silicon accelerator hardware), as well as diverse computational libraries and frameworks with specific versioning requirements. Altogether this can pose a formidable challenge to even the most advanced health care IT environments when constrained by operational procedures written against non-research based systems.

We therefore conclude that a separate space is required to conduct large scale computational research. This environment is specifically crafted to facilitate research productivity. Since its user base and function is vastly different, many of the traditional operational constraints imposed by hospital security or operational requirements can be avoided. One of the largest

<sup>1</sup>The Ansible project is available from <https://github.com/IKIM-Essen/EMCP-config>

hurdles to this type of system is the principle of parsimonious access to data or systems access privileges. This hampers researchers as installing software libraries required for research is rightfully considered undesirable for highly regulated and access-controlled hospital environments. Consequently, we seek an environment that removes those barriers by enabling researchers to install software and frameworks themselves, while simultaneously maintaining a reasonable level of overall security in the infrastructure. *In summary, a dedicated ML enclave is required adjacent to the hospital IT systems.*

## II. EXAMPLE USE CASES

In this section we exemplify some computing use cases and outline their requirements on a computing infrastructure.

### A. Understanding Hospital Patient Flows

One key indicator for quality of health care is the quality of patient journeys. A patient's journey through the hospital might deviate from the guidelines for multiple reasons, such as unforeseeable events and the complexity of the organization. Electronic Health Records (EHRs) contain information about events related to diagnostics and treatment of patients and can serve as input for retracing a patient's journey [4]. The application of tools to compute a patient's journey from observational data is called process mining. Process mining tools require event types (e.g., "an MRI was performed") and an associated time stamp. Such an application does not require personal data, but the data integrated from multiple hospital departments and their subsystems (e.g., pathology and radiology departments), facilitate the need for specific computing environments to install and run these process mining tools in addition to providing a data anonymization toolchain.

### B. Data-driven Treatment Effect Prediction in Oncology

Oncological research on treatment effects relies on a structured, summarized patient history (e.g., [5]), that contains treatment events from the first diagnosis until patient discharge. Currently, these structured reports are created by clinicians by manually coalescing information from multiple IT systems (this may contain, but shall not be limited to: computertomography reports, chemotherapy information, treatment plans, other clinical reports) and resolving contradictions from multiple sources. While the current process builds on abstractions on multiple levels, a purely data-driven approach would need to integrate raw data from various sources, ranging from images (magnetic resonance imaging, ultrasound, digital histopathology, ...) over structured data such as lab results, to natural language contained in reports. The challenges posed by these data sources comprise variety, volume and computational effort. In particular whole slide images, as used in digital histopathology require vast amounts of storage capacity, whereas the training time of state of the art natural language processing models is well within thousands of GPU hours. These traditional "Big Data" problems to derive both context and value from diverse sources are significantly hampered by the need for patient privacy, which requires a much more crafted approach.

### C. Using Microbiome Data to Inform the Selection of Antibiotics to Treat Infection

With metagenomic sequencing (see e.g., [6]) characterizing a patient's microbiome and thus any potential antibiotics resistance genes is well within the capabilities of modern science. However, characterizing the myriad of bacterial genomic fragments that metagenomic sequencing produces requires significant computational resources. While workflows exist already (e.g., [7]) their execution requires either containerized execution environments at scale or the ability to install bioinformatics software (e.g. via Bioconda [8]) neither of these are a good fit for the highly regulated hospital environment. In addition, few of these environments can provide access to enough systems (scale) needed to process current microbiome data in a timely manner. Due to the relative novelty and the lack of mature software systems, outsourcing the computational work can create its own challenges as it would create yet another black box introducing yet another set of obstacles for scientific discovery. Performing metagenome analysis to analyze microbiome data will require about 400 CPU core hours per gigabyte [7] of data assuming the least resource intensive workflows. Additionally this requires tens of thousands of files to be stored, many of them temporarily, all while requiring the installation of several dozen bioinformatics analysis tools, many of which have significant dependencies in the form of libraries etc.

## III. THE ARCHETYPICAL USER

Scientific computing users are comprised of a rather diverse group of (PhD) students, software developers, computer scientists, clinical scientists, bio-medical researchers and various (advanced) users. In an attempt to introduce the prototypical user we identified diverging requirements as the most common characteristic – not only across user groups, but also for the same single user. For example, when a machine learning researchers start prototyping, they want to run state of the art ML software and typically either are not proficient, nor should they spend their time dealing with 'low-level' setup details (e.g., GPU drivers, Compute Unified Device Architecture (CUDA) versions) that are traditionally the realm of system architects/integrators. In light of this, they should be provided with a ready-to-use environment (e.g., as provided by Google Colab [9] – just enter PyTorch, Tensorflow, or code of their favorite framework in a Jupyter Notebook and hit run). As the ML prototype models become more advanced, they may want to optimize model distribution and communication over multiple computation nodes, requiring access to low-level routines. In addition, our typical ML user has several dozen terabytes of image data, as well as text and structured data that needs to be extracted from a myriad of hospital business systems. We discovered that there is *not a single archetypical user* but rather a whole *host of usage patterns* that we need to support. Consequently, our user model includes a number of perspectives and also a number of services required to meet our researchers' needs.

#### IV. REGULATORY CONSTRAINTS

Data access in a hospital environment in Germany is regulated by a number of laws and regulations. IT security in larger hospitals in Germany is regulated by KRITIS (§8a IT-SiG, §6 BSI-KritisV), the German social code (§75b SGB V, §75c SGB V), national data privacy law (§22 section 2. BDSG), and finally the General Data Protection Regulation (GDPR) [10] or its German implementation. There are also additional codes covering the data protection requirements that vary by federal state (LDSG and LKHG).

While it is not appropriate to summarize the legal aspects here, the basic take home is parsimonious access, requirements for authentication and a need for documentation of procedures and access is an important part of any system. The authors deem data privacy to be a very important component of civil rights. However, it can be argued that compliance with a complex set of laws (see above) has not led to a hospital IT landscape conducive to a research-first approach. Certain components of the governing rules lead to unforeseen side-effects now that computational science is trying to get a foothold in the hospital system: e.g., the need to obtain an ethics vote prior to any use of data with Protected/Personal Health Information (PHI) ultimately leads to a de-facto ban on explorative studies.

Working with our organizational data privacy officer we devised a plan to create an environment that allows the application of modern computing tools to (bio-)medical data in a hospital environment. By removing the PHI from the data sets, data is considered no longer protected by the above mentioned laws and regulations. We therefore have reduced the regulatory burden and unlocked the full repertoire of scientific computing.

While we have removed PHI from the data, we nevertheless acknowledge the potential for misuse of the data and therefore implement robust IT security policies for the enclave.

#### V. SYSTEM REQUIREMENTS

The system needs to provide state of the art IT security and minimize compliance constraints as much as possible. Any future solution needs to comply with the various legal requirements (in our case EU GDPR, German federal §22 Abs 2 BDSG and state LDSG)<sup>2</sup>. However, in addition to complying with the legal framework, we need to ensure that the solution is deemed socially acceptable by both our peers inside the hospital as well as the general public, thus processing of PHI needs to be limited to an absolute minimum or avoided altogether. Social acceptance has another facet: developers need to find the working environment acceptable to achieve high performance. Any solution must make available modern computational abstractions that render developers and scientists productive as well as limit the amount of oversight required, e.g., for PhD students when working with data. Modern computing abstractions (e.g., container registries, Helm Charts,

<sup>2</sup>For an international audience: The rule set is similar to the HIPAA regulatory framework in the U.S.

existing ML/DL systems) are critical aspects of performance in an organisation that both develops and maintains software environments. The system needs to provide data ingress and egress solutions. And finally, ideally the system itself will utilize the expertise of our scientists and developers and allow us to evolve the system over time.

As our users' requirements change over time, we need to devise a platform for service provisioning rather than a fixed set of services. The list of services is certain to grow over time, but initially we included the following services in our plans:

- **File services**, i.e., secure Network File System (NFS) services inside the enclave and Server Message Block (SMB) services to specific medical research devices outside patient care.
- **Object store**, i.e., secure S3 services inside and outside the enclave.
- **Scheduling and resource management** for various classes of computational resources, e.g., specific GPU families.
- **Web-based user interfaces**, such as Jupyter [11], an interactive data science platform, requiring little or no command-line experience.
- **Package management** for scientific software. A myriad of scientific software systems exist as pre-packaged Conda packages.
- **Version control software**, such as Git in the form of CI/CD enabled tools such as GitLab [12].
- **Team collaboration software**, such as Rocket.Chat [13] or Mattermost [14].

#### VI. SYSTEM DESIGN

##### A. Solution Sketch

We use Linux/UNIX as a scientific computing platform. As international *cloud* solutions, providing Infrastructure as a Service (IaaS) [15] are deemed not acceptable for patient data both socially and legally in Germany, an on-premise solution was chosen.

The systems are deployed in a dedicated network segment separate from the hospital and the internet (cf. Figure 1), with incoming access possible via SSH into a bastion host with internet connectivity. Multi-factor authentication is required for access to the enclave. A proxy service enables internet access via HTTP inside the enclave for system updates and maintenance.

To reduce the compliance load as well as the potential for disaster should anything go wrong, we opted to de-identify data prior to storage inside the enclave. De-identified data will be copied into the enclave on demand to enable scientific computing on said data. In addition to data from patient care, anonymous research data is also abundant in our hospital environment and we will integrate equipment across campus into the research enclave by establishing dedicated Virtual Local Area Networks (VLANs) connected to our file service to enable the flow of research data into the enclave.

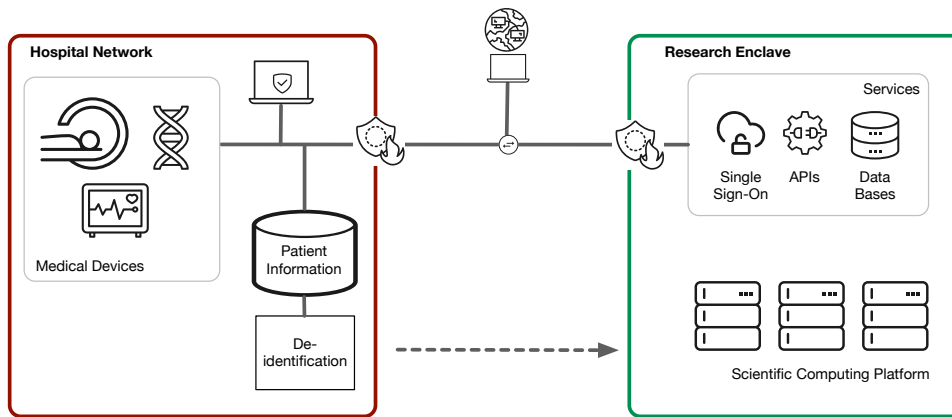


Fig. 1. The research computing enclave (right) is adjacent to the hospital network (left) but separated from it via a firewall. Personal information is de-identified before transfer to the enclave.

While container-based approaches for research are gaining momentum, *traditional* cluster computing approaches with shared file systems and centralized identity management are needed to assist at least some researchers. Therefore, we decided to support both by providing shell access to a managed cluster, including rootless containers (to prevent privileged escalation) and providing access to a bare metal provisioning service.

Finally, since we anticipate moving from our on-premise environment to a third-party IaaS platform, we use an Infrastructure as Code (IaC) stance, coding and documenting the environment in a single source code repository using Ansible [16]. The use of a popular tool like Ansible enables us to share the workload of establishing the enclave between several developers with and without operational roles in running said enclave.

### B. Architecture Details

In the following, we describe and motivate our implementation choices as well as putting them in context.

*Physical / Network Level Separation:* Network level separation is almost a requirement to avoid the significant restrictions imposed by the German legal and compliance framework. A well-defined logical location outside the hospital proper enables more flexibility and reduces the required compliance workload.

*System Security, Firewall, Two-Factor Authentication (2FA):* While data in the enclave is assumed to be free of PHI we nevertheless acknowledge that even in de-identified form, the data still represents a significant value as well as a potential threat to confidentiality. As a consequence, we chose to set up the system in a separate network location, protect it via a dedicated firewall system and require 2FA for access to the entire system. Hardware-based 2FA is currently considered best practice [17] for securing valuable data and computing assets.

*On-Premise and Bare Metal Provisioning Service:* While renting access to computing equipment and storage is certainly de rigueur currently (especially with smaller

scale/experimental test beds), for a production system both the scale of computing resources required along with the vast amount of data that needs to be stored/accessed quickly made the decision to go with a self-hosted, on-premise solution an easy one. The fact that our hospital had a fully equipped but under-utilized data center removed some of the expenses traditionally incurred for a large on-premise hosted solution. We chose Metal as a Service (MAAS) [18] as a bare metal provisioning system with an end-user interface as the basis of our enclave, this provides an environment similar to current popular *cloud* services, giving users an immediate familiarity. MAAS further improves performance against most cloud systems as it does not have the performance limitations that virtualization layers typically impose via abstraction layers, which can be a burden to scientific computing performance.

*Infrastructure as Code:* While we recognize that currently the financial and legal situation clearly favors on-premise solutions, we nevertheless acknowledge that in the future our solution might shift and we therefore decided to use an IaC approach allowing us to migrate the entire setup to future platforms. Furthermore this approach increases resilience of the overall computing infrastructure in case of hardware failure as well as facilitates upscaling as required. We decided to implement our infrastructure as a series of Ansible roles, available widely across the institute in a GitHub repository.<sup>1</sup> We invite pull requests by all, utilizing all available expertise rather than frustrating end users with restrictions. Importantly, the infrastructure is documented and pull requests for changes are enabled at the same time. Due to using version control, every change to the infrastructure is transparent and can be rolled back on demand. Using MAAS as the bare metal provisioning service and an IaC approach, we were able to quickly establish basic functionality and work together on an environment conducive to high research productivity.

*File services (NFS, SMB), object storage and single sign-on:* Easily accessible file services covering both the *traditional* (NFS, SMB) and object storage flavors enable a number of computational abstractions and provide ease of use for data transfer and use. We decided to install a server to share storage

for ongoing research via NFSv4, Samba [19] and another server acting as both S3-compatible object storage in/egress service and archive using MinIO [20]. Local disk caching compensates for NFS performance issues. To fully utilize the data storage services and enhance the overall utility of the system we chose a FreeIPA [21] server as our Lightweight Directory Access Protocol (LDAP) single sign-on system to centrally manage users and access credentials.

*A Platform for Container Orchestration – Kubernetes:* Provisioning of services such as databases, central applications, monitoring and messaging buses etc. in a cost effective and resilient way is currently almost synonymous with using Kubernetes [22]. The Helm [23] package manager provides pre-configured production-grade deployments for many data, communication or ML/DL services. The Kubernetes system itself is implemented via the open source version of Rancher [24] and provides local, redundant on-system storage to efficiently implement large and small-scale databases that would otherwise incur significant latency and bandwidth limitation if they were implemented via a network-attached storage. Furthermore, Rancher enables us to easily deploy multiple separate Kubernetes clusters for different use cases, such as infrastructure services on one side and scientific applications on the other side. This enables us to limit access to the underlying infrastructure these services are deployed on, thus increasing the protection of critical central services such as monitoring.

*Bringing Together Cluster and Container-Based Computing:* An analysis of end-user expectations and requirements made it clear that both, container-based approaches and *traditional* cluster computing approaches with shared file systems and centralized identity management are needed.

Therefore we devised two modes of access to the system: i) a Linux cluster accessible via a login server providing shell access including rootless containers and ii) access to the bare metal provisioning service enabling users to provision their own nodes with full root access. Importantly, users can choose to self-administer their own nodes, using the Ansible recipes provided for cluster nodes, losing only the access to the NFS file service. Both in the case of i) and in the case of ii) limitations are necessary to prevent privilege escalation. For mode i) it must be ensured, that users inside a container are not able to alter or write files to the host with arbitrary UIDs<sup>3</sup>. Conversely, users in mode ii) must not be able to access the NFS file service, as they may assume the role of any user through local root privileges.<sup>4</sup> This way it would be possible to bypass access control and inflict arbitrary damage on the overall system.

### C. Planning for Data Ingress and Egress

Access to patient and research data is essential for scientific computing. Hence, building data conduits that work for the

<sup>3</sup>docs.docker.com/engine/security/users-remap/

<sup>4</sup>The authors are aware that using Kerberized NFS would solve this problem but also increase complexity. The available S3-compatible object storage is the primary storage for users with independent platforms.

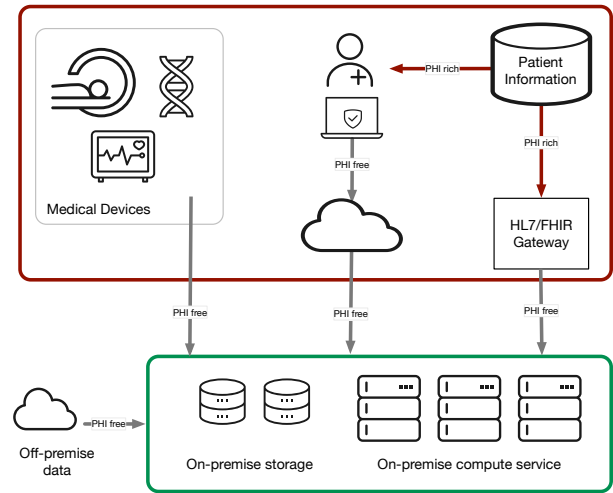


Fig. 2. The data flow into the research enclave.

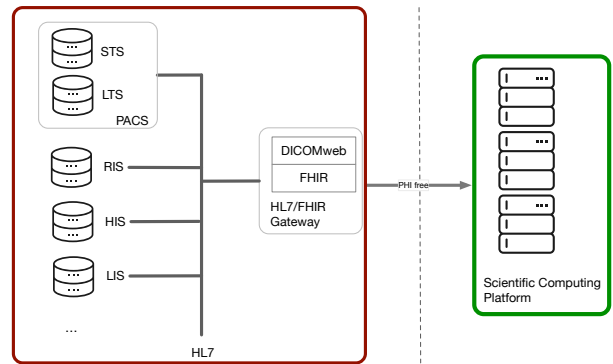


Fig. 3. A dedicated FHIR gateway is used to de-identify health care data from various data subsystems and make it available for research purposes. Abbreviations: STS - short-term storage, LTS - long-term storage, RIS - radiology information system, HIS - hospital information system, LIS - lab information system.

intended audience and yet guarantee privacy is critical.

Research data enters the EMCP on several pathways, as shown in figure 2. As indicated, it is important that solely de-identified data is present within the platform. To ensure this, we i) de-identify data automatically via a Fast Healthcare Interoperability Resources (FHIR) gateway and in cases a fully automatic de-identification is not possible, we ii) manually perform the de-identification.

- **Structured health care applications:** Access to de-identified patient data happens via an authenticated pre-existing local FHIR gateway that includes a Digital Imaging and Communications in Medicine Web Services (DICOMweb) interface. Figure 3 illustrates the concept.
- **De-identified research data:** Access to research data happens via either dedicated solutions like a dedicated research Picture Archiving and Communication System (PACS) implemented with Orthanc [25] or enclave-hosted SMB file services via dedicated VLANs for laboratory devices such as DNA sequencers. We note that human

genome data is a special class of data that we do not consider in this discussion.

- **One-offs special cases:** For one-time data dumps and other incoming transfers from the hospital ecosystem into the research enclave we provide a client for the hospital’s on-premise object store.
- **Global incoming data:** All other data transfer cases are taken care of by the use of an S3-compatible object storage service, implemented via MinIO [20]. The service is available inside and outside the enclave but does require strong multi-factor authentication.

## VII. FROM A USER’S PERSPECTIVE

The platform consumers can choose between working with the infrastructure directly (IaaS) or using one or more pre-installed, extensible platforms (Platform as a Service (PaaS)) and their software as a service (Software as a Service (SaaS)). Figure 4 illustrates this architecture from a user’s perspective.

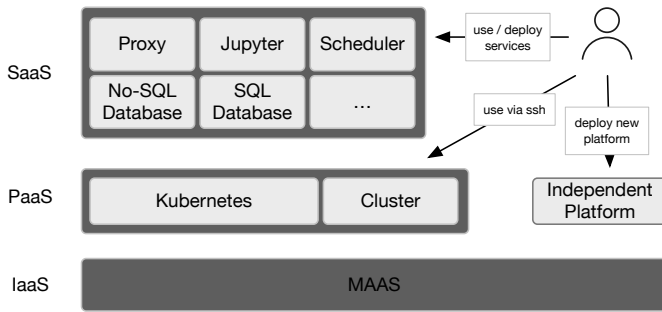


Fig. 4. To the researcher the system presents as a cluster with optional services and the option to resort to the infrastructure level to deploy an independent platform. Managed components are shown with a dark gray background.

We note that while most users will choose the PaaS approach, it is critical to enable advanced users to seek alternative approaches on their own. In addition, by providing the IaC code via a source code repository, we enable pull requests from advanced users. The EMCP at its core provides access to bare metal computing resources that at first glance would appear as a traditional research cluster. The layer just above this bare metal orchestration provided by MAAS demonstrates its power via the ability to be quickly reconfigured for various scientific workloads in a programmatic, repeatable fashion through the use of scripted infrastructure on top of MAAS. Additional flexibility comes from rootless Linux container capabilities that allow the use of pre-defined binary environments for scientific computing, creating instantly repeatable environments to aid in the scientific process.

Figure 5 outlines the choices of a researcher with a scientific computing need. A pure-Python computing need, requiring only common libraries allows a zero-configuration quickstart via the managed PaaS and Conda base environment. Similarly, computing needs that can be catered by packages available via Conda can be fulfilled with minimal additional configuration as Conda environment management is available in the

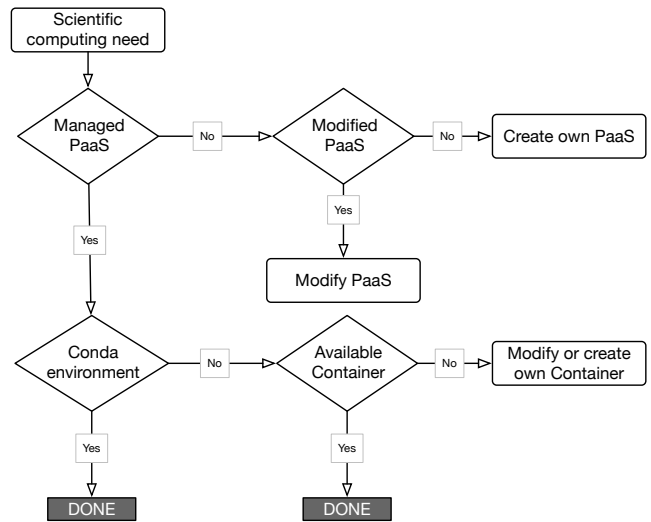


Fig. 5. A wide variety of scientific computing requirements are served with the setup we provide. The graphic shows the decision process from the users’ point of view. Decision nodes determine whether the scientific computing need can be fulfilled by the indicated element.

managed PaaS. We plan to simplify access even further by a JupyterHub [26] instance, making Jupyter [11] notebooks internally available without requiring command-line skills. In particular students are familiar with such environments (e.g., Google Colab [9]), as they rarely own ML hardware nor have access to their institutions’ facilities.

If additional software beyond a Conda environment is required, it can be deployed by using readily available containers or by setting up new ones. If containers cannot cater the researcher’s requirements, the researcher can choose the IaaS to instantiate their own nodes, deploy a modified version of the platform or start from scratch.

## VIII. CONCLUSION AND DISCUSSION

We believe we have successfully established our computational research enclave and evidence suggests that research productivity is growing.

A critical component is separating research computing from patient care IT. It is this separation of concerns that enables the freedom required to conduct research, e.g., giving researchers the freedom to install and use tools without being limited by legitimate hospital IT security concerns.

Possibly the most critical point of all is that all parties trust the data privacy rule set, the operational model and agree with the implemented restrictions.

This infrastructure is a work in progress, and by design it allows newly emerging computational paradigms and/or software to be easily accommodated. More importantly, the process of change is de-centralized (or rather crowd sourced) by sharing the infrastructure as code of the platforms.

We think that lowering the barrier to access and flattening the learning curve is critical to help, e.g., students be productive with the system. Future user modalities are likely to include readily accessible web-based workspaces via

Gitpod [27]. We will continue to explore options for use of our system and implement them as needed.

#### LIST OF ACRONYMS

**2FA** Two-Factor Authentication  
**ASIC** Application-Specific Integrated Circuit  
**CUDA** Compute Unified Device Architecture  
**DL** Deep Learning  
**DICOMweb** Digital Imaging and Communications in Medicine Web Services  
**EHR** Electronic Health Record  
**EMCP** Essen Medical Computing Platform  
**FHIR** Fast Healthcare Interoperability Resources  
**FPGA** Field Programmable Gate Array  
**GDPR** General Data Protection Regulation  
**GE** Gigabit / second Ethernet  
**GPU** Graphics Processing Unit  
**IaaS** Infrastructure as a Service  
**IaC** Infrastructure as Code  
**LDAP** Lightweight Directory Access Protocol  
**MAAS** Metal as a Service  
**ML** Machine Learning  
**NFS** Network File System  
**PaaS** Platform as a Service  
**PACS** Picture Archiving and Communication System  
**PHI** Protected/Personal Health Information  
**SaaS** Software as a Service  
**SMB** Server Message Block  
**VLAN** Virtual Local Area Network

#### REFERENCES

- [1] A. Jain, D. Way, V. Gupta, Y. Gao, G. de Oliveira Marinho, J. Hartford, R. Sayres, K. Kanada, C. Eng, K. Nagpal, K. B. DeSalvo, G. S. Corrado, L. Peng, D. R. Webster, R. C. Dunn, D. Coz, S. J. Huang, Y. Liu, P. Bui, and Y. Liu, "Development and Assessment of an Artificial Intelligence-Based Tool for Skin Condition Diagnosis by Primary Care Physicians and Nurse Practitioners in Teledermatology Practices," *JAMA Network Open*, vol. 4, no. 4, pp. e217249–e217249, 04 2021. [Online]. Available: <https://doi.org/10.1001/jamanetworkopen.2021.7249>
- [2] S. M. McKinney, M. Sieniek, V. Godbole, J. Godwin, N. Antropova, H. Ashrafian, T. Back, M. Chesus, G. S. Corrado, A. Darzi, M. Etemadi, F. Garcia-Vicente, F. J. Gilbert, M. Halling-Brown, D. Hassabis, S. Jansen, A. Karthikesalingam, C. J. Kelly, D. King, J. R. Ledsam, D. Melnick, H. Mostofi, L. Peng, J. J. Reicher, B. Romera-Paredes, R. Sidebottom, M. Suleyman, D. Tse, K. C. Young, J. De Fauw, and S. Shetty, "International evaluation of an ai system for breast cancer screening," *Nature*, vol. 577, no. 7788, pp. 89–94, Jan 2020. [Online]. Available: <https://doi.org/10.1038/s41586-019-1799-6>
- [3] S. Liu, K. C. See, K. Y. Ngiam, L. A. Celi, X. Sun, and M. Feng, "Reinforcement learning for clinical decision support in critical care: Comprehensive review," *J Med Internet Res*, vol. 22, no. 7, p. e18477, Jul 2020. [Online]. Available: <https://www.jmir.org/2020/7/e18477>
- [4] F. Marazza, F. A. Bukhsh, J. Geerdink, O. Vijlbrief, S. Pathak, M. v. Keulen, and C. Seifert, "Automatic process comparison for subpopulations: Application in cancer care," *International Journal of Environmental Research and Public Health*, vol. 17, no. 16, 2020. [Online]. Available: <https://www.mdpi.com/1660-4601/17/16/5707>
- [5] M. Wiesweg, C. Preuß, J. Roeper, M. Metznermacher, W. Eberhardt, U. Stropiep, K. Wedeken, H. Reis, T. Herold, K. Darwiche, C. Aigner, M. Stuschke, H.-U. Schildhaus, K. W. Schmid, M. Falk, L. Heukamp, M. Tiemann, F. Griesinger, and M. Schuler, "Braf mutations and braf mutation functional class have no negative impact on the clinical outcome of advanced nscl and associate with susceptibility to immunotherapy," *European Journal of Cancer*,

- vol. 149, pp. 211–221, May 2021. [Online]. Available: <https://doi.org/10.1016/j.ejca.2021.02.036>
- [6] T. Thomas, J. Gilbert, and M. F., "Metagenomics - a guide from sampling to data analysis." *Microb Informatics Exp*, vol. 2, 2012. [Online]. Available: <https://link.springer.com/article/10.1186/2042-5783-2-3>
- [7] F. Meyer, S. Bagchi, S. Chaterji, W. Gerlach, A. Grama, T. Harrison, T. Paczian, W. L. Trimble, and A. Wilke, "Mg-rast version 4- lessons learned from a decade of low-budget ultra-high-throughput metagenome analysis," *Brief Bioinform*, 2017. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/29028869>
- [8] B. Grüning, R. Dale, A. Sjödin, B. A. Chapman, J. Rowe, C. H. Tomkins-Tinch, t. B. T. Renan Valieris, and J. Köster, "Bioconda: Sustainable and comprehensive software distribution for the life sciences," *Nature Methods*, 2018.
- [9] T. Carneiro, R. V. Medeiros Da Nóbrega, T. Nepomuceno, G.-B. Bian, V. H. C. De Albuquerque, and P. P. R. Filho, "Performance analysis of google colab for a tool for accelerating deep learning applications," *IEEE Access*, vol. 6, pp. 61 677–61 685, 2018.
- [10] 2018 reform of eu data protection rules. European Commission. [Online]. Available: [https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes\\_en.pdf](https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf)
- [11] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, "Jupyter notebooks – a publishing format for reproducible computational workflows," in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, F. Loizides and B. Schmidt, Eds. IOS Press, 2016, pp. 87 – 90.
- [12] GitLab Inc. (2021) An open source end-to-end software development platform with built-in version control, issue tracking, code review and CI/CD. [Online]. Available: <https://gitlab.com/gitlab-org/gitlab-foss/>
- [13] The Rocket.Chat project. (2021) An open source team chat platform. [Online]. Available: <https://github.com/RocketChat/Rocket.Chat>
- [14] The Mattermost project. (2021) An open-source, self-hostable online chat service. [Online]. Available: <https://github.com/mattermost>
- [15] P. Mell and T. Grance, "The nist definition of cloud computing (technical report)," National Institute of Standards and Technology: U.S. Department of Commerce., Special publication 800-145. A/HRC/27/37, 2011.
- [16] The Ansible project. (2021) Ansible – ansible is a radically simple it automation platform that makes your applications and systems easier to deploy and maintain. [Online]. Available: <https://github.com/ansible/ansible>
- [17] J. Colnago, S. Devlin, M. Oates, C. Swoopes, L. Bauer, L. Cranor, and N. Christin, "It's not actually that horrible: Exploring adoption of two-factor authentication at a university," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Apr. 2018, p. 456.
- [18] The MAAS project. (2021) MAAS – very fast server provisioning for your data centre. [Online]. Available: <https://git.launchpad.net/maas/>
- [19] The Samba project. (2021) Samba – opening windows to a wider world. [Online]. Available: <https://samba.org>
- [20] The MinIO project . (2021) MinIO – multi-cloud object storage. [Online]. Available: <https://github.com/minio/minio>
- [21] The FreeIPA project . (2021) FreeIPA – open source identity management solution. [Online]. Available: <https://freeipa.org>
- [22] The Kubernetes project . (2021) Kubernetes – production-grade container orchestration – automated container deployment, scaling, and management. [Online]. Available: <https://github.com/kubernetes/kubernetes>
- [23] The Helm project . (2021) Helm – the package manager for kubernetes. [Online]. Available: <https://github.com/helm/helm>
- [24] The Rancher project. (2021) Rancher – complete container management platform. [Online]. Available: <https://github.com/rancher/rancher>
- [25] S. Jodogne, "The Orthanc ecosystem for medical imaging," *Journal of Digital Imaging*, vol. 31, no. 3, pp. 341–352, Jun 2018. [Online]. Available: <https://doi.org/10.1007/s10278-018-0082-y>
- [26] The JupyterHub project. (2021) A multi-user version of the notebook designed for companies, classrooms and research labs. [Online]. Available: <https://github.com/jupyterhub/jupyterhub>
- [27] gitpod.io. (2021) Gitpod. [Online]. Available: <https://www.gitpod.io>

## IX. APPENDIX: DESCRIPTION OF HARDWARE IN THE ENCLAVE

The authors thought it useful to provide an overview of the component bits of the enclave as well as a list of the services established.

Our guiding principle was KISS<sup>5</sup>. We assume that complexity will arise no matter what and by deliberately keeping the components simple we can reduce the potential for complex errors. We chose 1 Gigabit / second Ethernet (GE) and 10 GE network ports and components over potentially faster but more expensive and/or more complex alternatives, dedicated low-cost network hardware over multi-purpose, complex setups supporting multiple VLANs per device.

TABLE I  
THE ENCLAVE HARDWARE SHOPPING LIST

Name	Short description	Function
Network switches	Layer 2 ethernet switches (10 GE or faster)	Top-of-rack Layer 2 switch providing internal enclave network. The switches are cross-connected via 100 GE. We chose traditional ethernet over fibre.
Baseboard Management Controller switches	Layer 2 ethernet switch (1 GE or slower)	Remote management out-of-band using Baseboard Management Controllers, dedicated device for security and to minimize overall complexity.
Infrastructure nodes	Dual CPU Servers with approx. 200GB RAM and local redundant storage	External connectivity via proxy, reverse proxy and bastion host, IaaS, LDAP for auth
File servers	Dual CPU servers with approx. 800GB RAM, 10 GE and local storage	NFS, SMB, S3 services
Compute-Servers	Dual CPU servers with approx. 200GB RAM and dedicated local data cache respectively	CPU-bound activities, prefer many independent IO pipelines in small machines over fewer pipelines in fewer larger machines
GPU Servers	Dual CPU servers with 1TB or more RAM, 6 or more GPUs and large local NVMe data cache	Machine Learning, Deep Learning

TABLE II  
CURRENT LIST OF INTERNAL SERVICES

Name	Purpose	Comment
MAAS	IaaS platform	Ubuntu MAAS allows user driven OS deployments on bare metal
LDAP	Auth and single sign-on	implemented by FreeIPA
Slurm	Batch processing and queueing system	
NFS	industry standard shared file system	performance issues are compensated for with local disk caching
SMB	industry standard protocols for accessing shared file systems	used for connecting to laptops and lab devices
S3	industry standard storage API	implemented via MinIO
Kubernetes	industry standard for container orchestration platform	implemented via Rancher

<sup>5</sup>Abbreviation for the design principle “keep it simple, stupid”.