Privacy-Preserving Federated Learning and Analytics with FLAME

Marius Herr[1,*,†], Hammam Abu Attieh[2,†], Peter Placzek[1], Mehmed Halilovic[2], Bruce Schultz[3], David Hieber[1], Alexander Röhl[1], Mehrshad Jaberansary[3], Paul Brassel[4], Max Schaible[4], Toralf Kirsten[4], Fabian Prasser[2] and Oliver Kohlbacher[1]

[1]*Institute for Translational Bioinformatics, Maria-von-Linden-Str. 6, Tübingen, 72076, Germany*

[2]*Berlin Institute of Health at Charité – Universitätsmedizin Berlin, Medical Informatics Group, Charitéplatz 1, 10117, Berlin, Germany*

[3]*Institute for Biomedical Informatics (BI-K), Faculty of Medicine and University Hospital Cologne, Kerpener Straße 62, Cologne, 50937, Germany*

[4]*Institute for Medical Informatics, Statistics and Epidemiology, Leipzig University, Härtelstraße 16-18, 04107 Leipzig*

Abstract

FLAME is a modular hub-and-node platform for privacy-preserving federated learning and analytics in medical research, developed within the German national PrivateAIM project. It executes containerized analysis workflows at locally governed nodes, keeping patient-level data under institutional control. Data access, communication between components, analysis execution, and result handling are controlled through policy and security mechanisms. FLAME combines a central coordination hub, institutionally controlled execution nodes, and a software development kit that supports reusable federated workflows across different data sources. Initial deployments at several German university hospitals demonstrated GDPR-compliant federated analytics on clinical and genomic datasets, with performance comparable to centralized training.

Keywords

federated learning, privacy-preserving ML, healthcare infrastructure

1. Introduction

In Germany, the Medical Informatics Initiative (MII) has established Data Integration Centers (DICs) at university medical centers to harmonize clinical routine data and enable secondary use for research [1]. Yet, the value of this infrastructure depends not only on data availability but also on analytical platforms that can process distributed data under real-world hospital IT, governance, and privacy constraints. The PrivateAIM [2] project addresses this need by developing a federated machine learning and data analytics platform for the MII in which analyses come to the data rather than data coming to the analyses. The project explicitly targets limitations of existing distributed solutions, including restricted support for complex machine learning and data science tasks, limited integration of privacy-enhancing technologies, deployment complexity, and insufficient scalability or maturity.

FLAME, the Federated Learning and Analytics Methods Platform, is the core technical outcome of this effort. It is designed as a deployable, privacy-preserving, service-oriented platform for enabling federated analytics in medicine. Rather than centralizing patient-level data, it distributes containerized analysis workflows to locally governed nodes, where computations are executed under institutional control. Only controlled intermediate or final results are exchanged.

FLAME is conceptually embedded in a broader landscape of privacy-preserving medical data analysis. Existing platforms such as DataSHIELD [3], the Personal Health Train implementations [4], Vantage6

Joint Proceedings of the AIME 2026 Workshops: 1st International Workshop on Multicentric and Privacy-preserving Learning in Healthcare, Foundation Models for Public Health and Epidemiology: From Promise to Practice, and First International Workshop on Knowledge Graphs for Health, July 10, 2026, Ottawa, Canada

*Corresponding author.

†These authors contributed equally.

marius.herr@uni-tuebingen.de (M. Herr); hammam.abu-attieh@bih-charite.de (H. Abu Attieh)

0000-0002-6018-9111 (M. Herr); 0009-0006-6149-1503 (H. Abu Attieh)

[5], Flower [6], and Fed-BioMed [7] each address parts of the problem. DataSHIELD, for example, is well established for non-disclosive statistical analysis based on predefined R functions and statistical disclosure control. Personal Health Train and container-based platforms provide greater flexibility by moving analysis containers to data providers. However, many existing approaches face challenges in clinical production environments, including limited analytical flexibility, complex deployment, or insufficient integration into local governance processes. In this context, FLAME represents a modular hub-and-node platform with strong healthcare integration, generic containerized workflows, and embedded governance and policy-enforcement mechanisms. In this paper, we demonstrate how analytical workflows can be executed within the FLAME infrastructure while preserving privacy and maintaining compliance with distributed data governance requirements in federated analysis environments.

2. FLAME Architecture and Analysis

FLAME enables privacy-preserving federated analysis execution through a distributed hub-and-node workflow. Using the FLAME Software Development Kit (SDK), analysts can implement federated workflows without handling low-level communication and orchestration details directly. Users create analysis through the FLAME Hub by defining execution parameters, selecting participating institutions, and uploading analysis scripts and configuration files.

After submission and approval of institutions, the FLAME Hub validates the request, applies access and policy controls, and triggers the creation of a containerized analysis environment. Analysis containers are built from pre-approved master images containing validated software dependencies and runtime environments. The worker service extends these master images with user-provided analysis code and configuration files, then distributes the resulting Docker images through the container registry to project-specific repositories accessible by the participating nodes.

At each participating hospital, the FLAME Node pulls the analysis image and deploys it within a Kubernetes-based execution environment. Execution occurs locally at the institution hosting the data, ensuring that patient-level data never leaves the site. Access to clinical and research data is mediated through secure gateways and policy-controlled interfaces supporting standardized data sources such as HL7 FHIR and S3-compatible object storage. A local Policy Decision Point (PDP) enforces permissions and controls communication and data access during runtime.

The execution environments are isolated through containerization, reverse proxies, and network policies that restrict communication to approved interfaces and services. Intermediate updates, metadata, or encrypted partial results can be exchanged through the central message broker. Depending on the workflow, a dedicated aggregator node combines distributed outputs into final analysis results.

After execution, local and aggregated results are transferred back to the FLAME Hub storage and made available to the analyst through the user interface. The FLAME Hub coordinates result collection, workflow completion, and retrieval of generated outputs and metadata for further analysis and interpretation.

3. Security, Privacy, and Governance

Security mechanisms in FLAME are implemented across communication, data management, permission control, analysis execution, and organizational governance. The initial threat model assumes semi-honest, or honest-but-curious, participants: the Hub and Nodes follow the protocol, but may attempt to infer additional information from visible data or metadata.

Communication between nodes via the FLAME Hub is secured using end-to-end encryption. Each node generates an elliptic-curve key pair, while public keys are maintained in a central trust store and private keys remain local. Derived symmetric keys are used for secure data exchange and message distribution. The platform also stores hashes of Docker-based analyses, enabling integrity and provenance checks. Access control is based on a standardized permission and policy model covering users, roles, clients, analyses, data access, execution rights, and result retrieval.

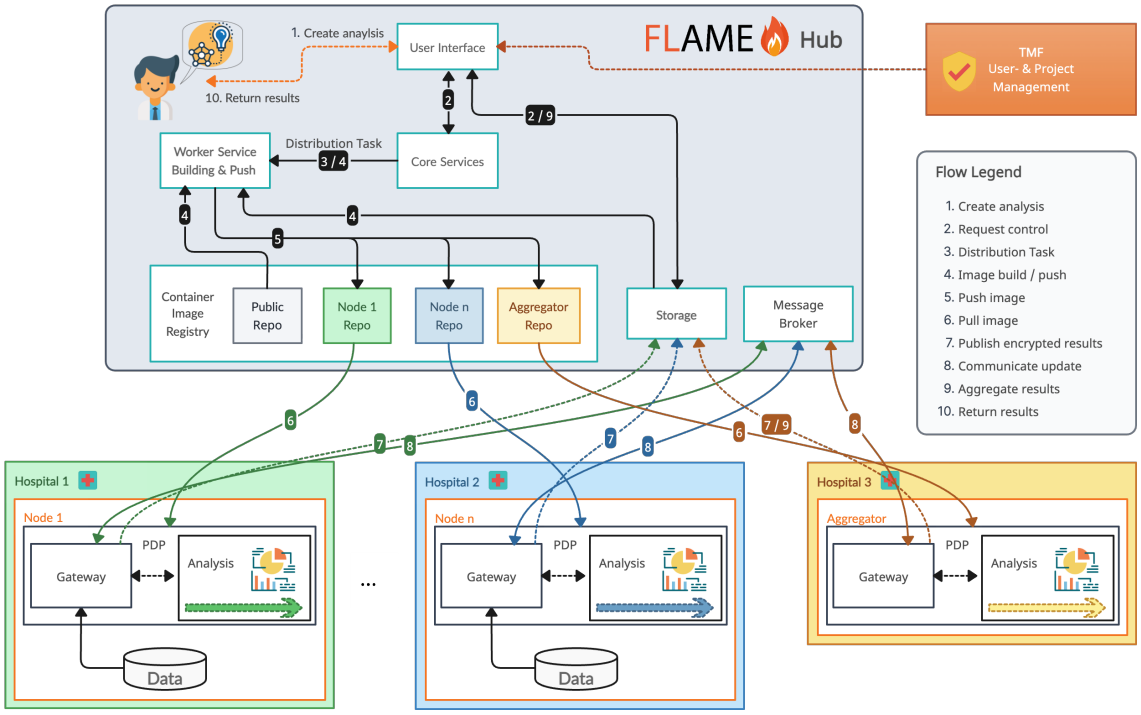

Figure 1: FLAME workflow from analysis creation and container distribution to local node execution, encrypted result exchange, aggregation, and result retrieval. The numbered interactions correspond to the workflow steps described in the legend.

The secure analysis lifecycle consists of three phases: creation, execution, and result retrieval. During creation, analysts define workflows, upload scripts, select nodes, and specify required permissions. Optional privacy-enhancing technologies can be configured, including secure multiparty computation, homomorphic encryption, or differential privacy [8]. During execution, nodes verify the authenticity of received images and run analyses in isolated environments under the control of a local Policy Decision Point. During result retrieval, the aggregator compiles outputs and encrypts final results so that sensitive data is not exposed unencrypted to intermediary services.

FLAME also distinguishes three maturity levels for privacy-enhancing technologies. At the (1) Analysis-Based level, PETs are implemented directly in the analyst's code, requiring substantial trust in the submitted workflow. At the (2) Library-Based level, PETs are implemented using the FLAME SDK or compatible libraries, increasing consistency and auditability. At the (3) System-Based level, PETs are enforced automatically by the platform through APIs, policies, and permissions, reducing reliance on user-defined logic and enabling stronger verifiable guarantees.

Governance remains a central part of FLAME's privacy model. The Charité data protection concept, for example, describes the FLAME Node as a local component for federated analyses and machine learning on pseudonymized data, without transferring the data outside the institution. Scientific use is project-based and requires formal project definition, ethics approval, a data protection concept, a Data Use and Access review, and data use agreements between participating institutions. It further distinguishes roles such as administrators, data managers, and researchers, granting access on a limited, need-to-know basis.

4. Use Case and Evaluation

FLAME was initially test-deployed at four German university hospital sites to evaluate whether the platform can be installed and operated under realistic institutional IT, governance, and security constraints.

These deployments focused on validating the core operational workflow, including containerized analysis distribution, local execution at FLAME Nodes, encrypted communication through the Hub, policy-controlled access to local data interfaces, and aggregation of distributed results.

For an initial functional evaluation, a federated machine learning use case was defined to classify cardiomyopathy subtypes using structured clinical data represented via HL7 FHIR profiles and gene expression data. This use case was selected because cardiomyopathy classification is clinically relevant and requires integration of heterogeneous data modalities. The evaluation resulted in federated models achieving performance equivalent to centralized training while maintaining local control over patient-level data. The federated workflow reproduced the centralized differential expression analysis, including the identification of 5,567 significant genes. The experiments demonstrated that FLAME can support distributed model training and result aggregation across institutional boundaries without requiring central data pooling. Beyond model performance, the evaluation confirmed the feasibility of the platform's technical workflow, including secure container execution, node-local data access, encrypted exchange of results, and coordinated result retrieval through the Hub.

Future work will focus on transferring these findings from test deployments to productive clinical operations. In particular, the first clinical use case using real patient data from four university hospitals, focusing on Atrial Fibrillation, is planned for Q2 2026, followed by a broader deployment across all participating university hospitals by March 2027. This next phase will evaluate not only technical scalability and robustness, but also the practical interaction between deployment, data access approvals, local governance processes, data use agreements, and operational coordination across DICs.

5. Discussion

FLAME addresses a central challenge in data-driven medicine: enabling meaningful multi-center research without undermining patient privacy, institutional data sovereignty, or regulatory compliance. Its main strength lies in combining flexible containerized analytics with a governance-aware, healthcare-oriented infrastructure. In contrast to frameworks focused primarily on federated model training, FLAME supports general-purpose federated analysis workflows. In contrast to more restrictive statistical platforms, it enables complex multi-step workflows, supports multiple programming languages, and integrates with domain-specific libraries.

At the same time, FLAME does not eliminate privacy risks solely through architectural measures. Federated outputs, model parameters, and intermediate values may still leak information if they are not adequately protected. The platform's value, therefore, lies in its layered approach: local execution, encrypted communication, policy enforcement, isolated containers, privacy-enhancing technologies, governance workflows, and output control.

The Atrial Fibrillation use case is important because it tests this interaction under realistic conditions. It requires not only technical execution but also data access approvals, contracts, local production deployment, data connection, script adaptation, and operational coordination across DICs. Its success would demonstrate that FLAME is not merely a prototype, but a practical infrastructure for privacy-preserving medical analytics in real clinical environments.

6. Conclusion

FLAME provides a promising architecture for privacy-preserving federated learning and analytics in medicine. Developed within PrivateAIM, it combines a modular hub-and-node architecture, secure local execution environments, containerized workflows, SDK-supported method development, encrypted communication, policy enforcement, and integration with MII governance processes. Its dynamic policy framework bridges legal, technical, and organizational requirements for multi-institutional collaboration while preserving institutional control, data sovereignty, and GDPR compliance. Initial deployments at several German university hospitals demonstrated the feasibility of federated analytics on clinical and

genomic data with performance comparable to centralized approaches, while reducing the need for central data pooling.

Acknowledgments

This research was partly funded by the German Federal Ministry of Research, Technology, and Space (BMFTR) through the PrivateAIM project (grant numbers 01ZZ2316A·B·K·L).

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] S. C. Semler, M. Boeker, R. Eils, D. Krefting, M. Loeffler, J. Bussmann, F. Wissing, H.-U. Prokosch, Die medizininformatik-initiative im überblick – aufbau einer gesundheitsforschungsdateninfrastruktur in deutschland, Bundesgesundheitsblatt - Gesundheitsforschung - Gesundheitsschutz 67 (2024) 616–628. doi:10.1007/s00103-024-03887-5.

[2] PrivateAIM Consortium, Privacy-preserving analytics in medicine (privateaim), https://privateaim.de/, 2026. Accessed: 2026-05-13.

[3] A. Gaye, Y. Marcon, J. Isaeva, P. LaFlamme, A. Turner, E. M. Jones, J. Minion, A. W. Boyd, C. J. Newby, M.-L. Nuotio, R. Wilson, O. Butters, S. E. Wallace, I. Budin-Ljøsne, C. Oliver Schmidt, P. Boffetta, M. Boniol, M. Bota, K. W. Carter, N. deKlerk, C. Dibben, R. W. Francis, T. Hiekkalinna, K. Hveem, K. Kvaløy, S. Millar, I. J. Perry, A. Peters, C. M. Phillips, F. Popham, G. Raab, E. Reischl, N. Sheehan, M. Waldenberger, M. Perola, E. van den Heuvel, J. Macleod, B. M. Knoppers, I. Fortier, J. R. Harris, B. Woffenbuttel, P. R. Burton, DataSHIELD: Taking the analysis to the data, not the data to the analysis, International Journal of Epidemiology 43 (2014) 1929–1944. doi:10.1093/ije/dyu188.

[4] S. Welten, M. de Arruda Botelho Herr, L. Hempel, D. Hieber, P. Placzek, M. Graf, S. Weber, L. Neumann, M. Jugl, L. Tirpitz, K. Kindermann, S. Geisler, L. O. Bonino da Silva Santos, S. Decker, N. Pfeifer, O. Kohlbacher, T. Kirsten, A study on interoperability between two personal health train infrastructures in leukodystrophy data analysis, Scientific Data 11 (2024) 663. doi:10.1038/s41597-024-03450-6.

[5] F. Martin, M. Sieswerda, H. Alradhi, et al., vantage6, 2022. doi:10.5281/zenodo.7221216.

[6] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, J. Fernandez-Marques, Y. Gao, L. Sani, H. L. Kwing, T. Parcollet, P. P. B. de Gusmão, N. D. Lane, Flower: A friendly federated learning research framework, arXiv (2020). URL: https://arxiv.org/abs/2007.14390. arXiv:2007.14390.

[7] F. Cremonesi, M. Vesin, S. Cansiz, Y. Bouillard, I. Balelli, L. Innocenti, others, M. Lorenzi, Fedbiomed: Open, transparent and trusted federated learning for real-world healthcare applications, in: Federated Learning Systems: Towards Privacy-Preserving Distributed AI, Springer Nature Switzerland, Cham, 2025, pp. 19–41.

[8] H. Abu Attieh, M. Halilovic, R. Kerkouche, S. Rahimian, H. P. Wang, R. Rehms, G. Kaissis, A. Ziller, J. Kuntzer, K. Otte, T. Meurers, M. Fritz, D. Rueckert, F. Prasser, An exploratory survey study toward a conceptual framework for structuring the integration of privacy-enhancing technologies in federated learning and analytics, IEEE Access 14 (2026) 41508–41528. doi:10.1109/ACCESS.2026.3672991.