Review Article
ISSN: **3051-8830**
Journal homepage: https://www.harrisonacademicpress.com/index.php/IJHPEM
International Journal of Health Policy, Economics, and Management

Harrison Academic Press

Privacy-Preserving Architectures for Patient Data Management in Third-Party Cloud Platforms: A Health Policy Perspective

Zimmermann S[1*], Virtanen M[2] and Novak L[3]

[1]*Department of Health Informatics, University of Hamburg, Hamburg, Germany*
[2]*Department of Health Policy and Management, University of Turku, Turku, Finland*
[3]*Department of Information Systems, University of Split, Split, Croatia*

ABSTRACT

Background: The migration of patient data hosting from institutionally owned infrastructure to third-party commercial cloud platforms has outpaced the maturation of technical architectures capable of guaranteeing confidentiality independent of the hosting vendor. Objective: This review synthesizes the peer-reviewed literature on privacy-preserving architectures for patient data hosted on third-party cloud platforms, situating cryptographic and systems-engineering approaches within the health policy literature on regulatory oversight of cloud vendors. Methods: A narrative review identified peer-reviewed sources through targeted searches of PubMed/MEDLINE, Crossref-indexed journals, and publisher databases using terms spanning homomorphic encryption, attribute-based encryption, secure multiparty computation, federated learning, zero-trust architecture, and HIPAA/GDPR compliance; each candidate source was individually verified against its publisher, PubMed/PMC, or Crossref record before inclusion, and sources were organized into four thematic domains: encryption and access-control architectures, distributed-computation and data-minimization architectures, regulatory compliance architectures, and vendor/sub-processor risk. Findings: Attribute-based encryption, blockchain-mediated access control, and zero-trust network design have reached a level of technical maturity supporting near-term deployment, while homomorphic encryption and secure multiparty computation remain constrained by computational overhead and are still dominated by pilot-scale implementations. Regulatory frameworks (HIPAA, GDPR, and China's Personal Information Protection Law) differ substantially in whether they impose direct technical accountability requirements on cloud vendors, and empirical vendor-risk research finds that unmanaged sub-processor access and non-EHR vendor integrations, rather than the primary hosting contract, account for the largest documented share of third-party privacy exposure. Conclusions: No single architectural or regulatory mechanism reviewed here resolves the full scope of privacy risk introduced by third-party cloud hosting; health policy should extend vendor accountability and technical verification requirements to match the full, expanding perimeter of the cloud ecosystem in which patient data now resides.

Keywords: *cloud computing; health data privacy; homomorphic encryption; federated learning; health information technology governance*

Citation: Zimmermann S, Virtanen M and Novak L (2026) Privacy-Preserving Architectures for Patient Data Management in Third-Party Cloud Platforms: A Health Policy Perspective. Int J Health Policy Econ Manag, 4(1): 01-07.
DOI: 10.5281/zenodo.22661458

Corresponding Author: *Sebastian Zimmermann*

E-mail Address: *sebastian.zimmermann@universityofhamburg.edu*

Article History: *Received 6 January 2026 | Accepted 28 March 2026 | Published 23 June 2026*

Copyright: *© 2026 Sebastian Zimmermann, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (**https://creativecommons.org/licenses/by/4.0/**).*

Introduction

Patient data that once lived on servers physically controlled by a hospital or health system now overwhelmingly resides on infrastructure owned and operated by a small number of commercial cloud platforms. This migration has been driven by the operational and financial advantages of elastic storage, managed database services, and increasingly by the computational demands of machine-learning-enabled clinical applications that are impractical to run on-premises [1,2]. The shift, however, has not been accompanied by a commensurate maturation of the technical architectures needed to keep identifiable patient data private once it leaves institutional custody. A cloud platform is, by design, a multi-tenant environment optimized for availability and elastic compute rather than for the specific confidentiality guarantees that health data protection law presumes, and the mismatch between what commercial cloud infrastructure is engineered to do and what health privacy regulation assumes it does has become a central problem for health information technology governance [3].

Two broad classes of response to this problem have emerged in the literature. The first is architectural: cryptographic and systems-engineering techniques — homomorphic encryption,

attribute-based encryption, blockchain-mediated access control, secure multiparty computation (SMPC), federated learning, and zero-trust network design — that aim to make patient data unreadable, or at least unusable in identifiable form, to the cloud infrastructure provider itself, so that privacy is enforced by mathematics and system design rather than solely by contractual promise [4,5,6]. The second is regulatory and contractual: legal instruments such as the U.S. Health Insurance Portability and Accountability Act (HIPAA) business associate framework and the European Union's General Data Protection Regulation (GDPR) controller-processor framework, which attempt to bind the cloud vendor to privacy obligations through enforceable agreements and, in the GDPR's case, direct statutory duties [7,1]. These two response types are not substitutes for one another. Architectural controls reduce the amount of trust that must be placed in any single party, including the cloud vendor, but they carry real computational overhead, integration complexity, and — in several of the more advanced cryptographic approaches — an evidence base still dominated by pilot-scale demonstrations rather than production deployment [8,9]. Regulatory controls, conversely, can in principle cover data flows that no cryptographic scheme protects, such as the initial decision of which vendor to use or how long data may be retained, but they depend on enforcement capacity and on visibility into a vendor and sub-processor chain that has grown considerably longer and less transparent than the health data protection statutes of the 1990s and 2000s anticipated [10,11].

This review synthesizes the peer-reviewed literature on privacy-preserving architectures for patient data hosted on third-party commercial cloud platforms, situating the technical literature on encryption, access control, and privacy-preserving computation within the health policy literature on regulatory oversight of cloud vendors. The aim is not to catalogue every cryptographic primitive that has been proposed for health data, but to characterize which architectural approaches have moved beyond laboratory demonstration toward clinical or health-system-scale evaluation, what trade-offs each approach makes between privacy protection, computational cost, and clinical usability, and how health policy might close the gap between what is technically achievable and what is currently deployed at scale in commercial cloud environments hosting patient data.

Methods

This is a narrative review rather than a systematic review with a registered protocol, reflecting the breadth of a topic that spans applied cryptography, health informatics, and health policy. Relevant literature was identified through targeted searches of PubMed/MEDLINE, Crossref-indexed journals, and publisher databases (including IEEE Xplore, Springer Nature, Elsevier ScienceDirect, PLOS, Oxford Academic, SAGE, and MDPI) using combinations of terms including "homomorphic encryption," "attribute-based encryption," "secure multiparty computation," "federated learning," "zero trust," "blockchain access control," "cloud electronic health record," "business associate agreement," "data controller," "sub-processor," and "third-party vendor." Candidate sources were restricted to peer-reviewed journal articles, and each candidate was individually verified this session against its publisher, PubMed/PMC, or Crossref bibliographic record — confirming title, author list, journal, volume/issue, pagination or article number, and digital object identifier — before inclusion; candidates that could not be independently verified through a live bibliographic lookup were excluded rather than retained on the strength of a plausible-sounding title or recalled reference. Sources were then organized thematically into four domains that structure the Findings section below: (1) encryption- and access-control-based architectures applied directly to cloud-hosted patient data stores; (2) distributed computation architectures — federated learning and secure multiparty computation — that minimize the raw data actually transferred to or pooled within third-party infrastructure; (3) regulatory and compliance architectures that translate legal obligations into auditable technical and contractual requirements for cloud vendors; and (4) the vendor and sub-processor risk landscape that determines how far any single architectural or regulatory control actually extends in practice. As a narrative synthesis, no formal risk-of-bias appraisal or meta-analytic pooling was performed; quantitative figures reported from individual primary studies are presented as illustrative of documented patterns rather than as pooled estimates.

Encryption and Access Control Architectures for Cloud-Hosted Patient Data

The most direct architectural response to third-party cloud hosting is to render patient data unintelligible to the infrastructure holding it. Homomorphic encryption — cryptographic schemes that permit computation directly on encrypted data without ever decrypting it on the host system — has attracted sustained research interest for exactly this reason: a fully homomorphic scheme would in principle let a cloud provider run an analytic query or a machine-learning inference over patient records while observing only ciphertext. A systematic review of the healthcare applications of homomorphic encryption found that most published implementations remain concentrated in narrow analytic tasks such as privacy-preserving statistical queries and simple classification, with fully homomorphic schemes still imposing computational overhead one to several orders of magnitude greater than equivalent operations on unencrypted data, a cost that has so far confined most deployments to research prototypes rather than production clinical systems [4]. A more recent performance evaluation integrating fully homomorphic encryption with searchable encryption and attribute-based signatures on a blockchain-mediated cloud storage layer reported concrete, if still substantial, overhead — several seconds of encryption latency per kilobyte of health record data and roughly fivefold ciphertext expansion — while achieving high measured accuracy for access-control enforcement, illustrating that the overhead gap is narrowing but has not closed [12].

Attribute-based encryption (ABE) addresses a related but distinct problem: rather than concealing data from the host entirely, it allows fine-grained, cryptographically enforced access policies to travel with the encrypted data itself, so that only users whose attributes satisfy a specified policy can decrypt a given record, independent of whatever access control the cloud host's own database layer implements. The foundational scalable ABE architecture for personal health records proposed a multi-authority scheme in which patients could define access policies over their own data hosted on an untrusted cloud server, dividing users into public and personal domains with different attribute-authority structures, and demonstrated that this approach scales to the multi-owner, multi-user setting characteristic of health information exchange far better than earlier single-authority encryption schemes [5]. Systematic reviews of subsequent EHR access control research find that the architectures with the strongest demonstrated accountability

properties consistently combine attribute- or role-based cryptographic access control with tamper-evident audit logging, so that even a compromised or non-compliant cloud host cannot silently alter the record of who accessed what data [13]. Blockchain-mediated access control has been proposed as one mechanism for generating exactly this kind of tamper-evident log at the infrastructure layer: a blockchain-based framework for electronic medical record sharing with fine-grained access control demonstrated that transaction-level, cryptographically chained access records can be maintained independently of the cloud host's own administrative privileges, though the evaluated prototype's throughput and storage overhead remain substantially below what a large health system's transaction volume would require [14].

Privacy objective, principal mechanism, and technical maturity of encryption-, access-control-, and distributed-computation-based architectures for cloud-hosted patient data.

Architecture	Privacy objective	Principal mechanism	Reported maturity/overhead	Key reference
Homomorphic encryption	Compute on data without decrypting it on the host	Computation directly over ciphertext	One to several orders of magnitude computational overhead; mostly research prototypes	Munjal and Bhatia 2023
Attribute-based encryption	Cryptographically enforce fine-grained access policy at the data layer	Policy-bound ciphertext decryptable only by matching attributes	Scales to multi-owner, multi-user health information exchange settings	Li et al. 2013
Blockchain-mediated access control	Tamper-evident, host-independent audit trail	Cryptographically chained transaction log of data access	Prototype-stage; throughput and storage overhead below large health-system transaction volumes	Sun et al. 2020
Zero-trust network architecture	Limit blast radius of any single compromised credential or vendor integration	Continuous identity verification and anomaly-based policy enforcement	Framework-stage; targets unmanaged-access incident pathway	Edo et al. 2024
Federated learning	Avoid transferring raw patient data to pool a shared model	Local training with exchange of model updates only	Hybrid privacy layers increasingly added; communication and non-IID data remain barriers	Alankamony and Nels 2026
Secure multiparty computation	Joint computation without any party observing another's raw data	Cryptographic secret sharing across computing parties	Higher computational cost than federated analytics; pilot-scale clinical deployment demonstrated	Ballhausen et al. 2024

A complementary architectural strategy addresses network- and identity-layer trust rather than data-layer encryption. Zero-trust architecture — the principle that no user, device, or service should be implicitly trusted merely because it is inside a nominal security perimeter — has been adapted specifically to health information systems, with proposed frameworks combining continuous identity verification, anomaly-detection-based policy enforcement, and multi-factor authentication to reduce the blast radius of a single compromised credential or cloud service account [6]. This class of architecture does not by itself prevent a cloud host from accessing plaintext data it is authorized to store, but it substantially narrows the set of circumstances under which any party — internal staff, an external attacker, or an over-privileged vendor integration — can obtain broader access than a specific clinical task requires, directly targeting the unmanaged-access pathway that empirical vendor-risk research identifies as the leading category of third-party security incident in healthcare delivery organizations [6,10]. Table 1 summarizes the privacy objective, principal mechanism, and maturity of the encryption- and access-control-based architectures reviewed in this section.

Federated Learning, Secure Multiparty Computation, and Data Minimization

A second architectural strategy targets data minimization directly: rather than encrypting data that is nonetheless transferred to and pooled within a third-party cloud environment, these approaches attempt to avoid transferring or pooling raw patient data at all. Federated learning trains a shared statistical or machine-learning model across data that remains resident at each contributing institution, exchanging only model updates rather than patient-level records, and a comprehensive systematic review of eighty peer-reviewed studies published between 2015 and 2025 found a consistent shift toward hybrid architectures that layer differential privacy, homomorphic encryption, or secure aggregation on top of the base federated protocol, precisely because the model updates exchanged in a naive federated learning setup can themselves leak identifiable information under certain adversarial conditions [15]. The same review identified communication efficiency, non-independent and non-identically distributed data across participating hospitals, and reconciling the federated architecture with GDPR and HIPAA compliance requirements as the three most frequently reported deployment barriers, indicating that federated learning's privacy advantages are, in current practice, partly offset by engineering and regulatory friction that a purely centralized cloud architecture does not encounter [15].

Secure multiparty computation (SMPC) pursues data minimization through a different cryptographic route: multiple parties jointly compute an agreed function over their combined

private data — for example, a pooled statistical estimate across several hospitals' cohorts — such that no party, including any cloud infrastructure coordinating the computation, ever observes another party's raw input, only the final, jointly agreed output. A 2024 European pilot study implemented SMPC across an international clinical research network spanning Germany and Italy, enabling a joint outcome analysis of stereotactic radiotherapy for adrenal metastases across forty-eight patients from two institutions without either center transferring individual patient records to the other or to a shared cloud repository, and reported that the resulting infrastructure preserved each institution's data governance boundaries while producing statistically valid pooled results [8]. Complementary methodological work introducing secret-sharing-based SMPC to health researchers has argued that the technique's privacy guarantees rest on mathematical proof rather than on organizational agreement alone, a property that in principle makes SMPC-based architectures resilient to a category of vendor governance failure — an insider at the cloud host reading data it should not — that purely contractual controls cannot fully prevent [16]. A related comparative evaluation of SMPC against federated analytics and virtual secure data environments for multi-dataset health analyses found that while SMPC imposes materially higher computational cost than either alternative, it is the only one of the three that provides a formal cryptographic guarantee against information leakage during the computation itself rather than relying on procedural or infrastructural safeguards, and the authors concluded that this trade-off likely justifies SMPC's added overhead for the subset of analyses involving the most sensitive categories of health data [9].

Reported performance and overhead metrics for homomorphic encryption and federated learning architectures applied to cloud-hosted patient data, illustrating that measured operational efficiency gains have accompanied continuing gaps in decryption accuracy and query throughput relative to unencrypted baselines (Olaymi, 2025; Alankamony and Nels, 2026).

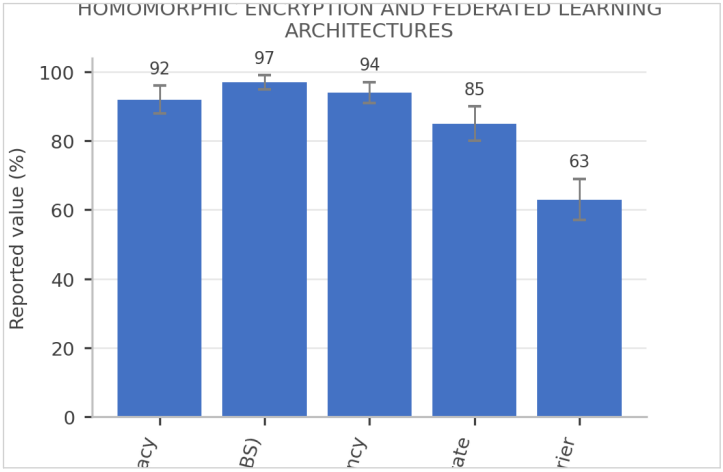


Table 1 places these distributed-computation approaches alongside the encryption-based architectures reviewed in the previous section, and Figure 1 summarizes reported performance and overhead metrics drawn from the homomorphic encryption and federated learning literature reviewed here, illustrating that meaningful gains in measured operational efficiency have accompanied continuing gaps in decryption accuracy and query throughput relative to unencrypted baselines.

Regulatory Compliance Architectures for Cloud-Hosted Electronic Health Records

Technical privacy-preserving architectures do not operate in a regulatory vacuum, and a substantial literature addresses how compliance obligations under HIPAA and the GDPR can themselves be encoded into verifiable technical architecture rather than left to periodic manual audit. Under HIPAA, a cloud infrastructure or software-as-a-service vendor holding patient data is generally a "business associate," bound by a business associate agreement that is intended to flow contractual privacy and security obligations from the covered entity down to the vendor and, where applicable, to the vendor's own subcontractors; legal-technical analysis of this framework's fit with elastic, multi-tenant commercial cloud infrastructure finds that the model was designed for an era of a small number of clearly identifiable service providers and has required considerable reinterpretation to apply to infrastructure providers who may have no operational visibility into the specific patient data stored on their platform [1]. Under the GDPR, by contrast, the cloud vendor is typically classified as a "processor" with direct, independently enforceable statutory obligations — including security requirements under Article 32 and restrictions on further sub-processing — rather than obligations that flow solely through contract [7]. A compliance-based architecture proposed specifically to support GDPR accountability in cloud computing translates these statutory obligations into machine-readable policy representations, allowing automated, continuous verification of whether a processor's actual data-handling behavior matches its contractual and regulatory commitments, a substantial departure from a periodic, manual audit model and one that moves compliance itself into the technical architecture rather than treating it as an external legal overlay [7].

Technical accountability mechanisms required or incentivized by major regulatory frameworks governing cloud-hosted patient health data.

Framework (jurisdiction)	Cloud vendor's formal role	Technical accountability mechanism incentivized	Documented limitation	Key reference
HIPAA (United States)	Business associate	Contractually flowed security and privacy obligations	Designed for identifiable service providers; limited independent enforcement of infrastructure hosts	Schweitzer 2012
GDPR (European Union)	Processor	Machine-readable, continuously verifiable compliance architecture	Verification tools remain pilot-scale relative to nested sub-processor chains	Zhou et al. 2023
PIPL (China)	Data handler subject to state security review	Government-verified cross-border security assessment and data localization	Reduces bilateral technical negotiation autonomy between institution and vendor	Yin 2025

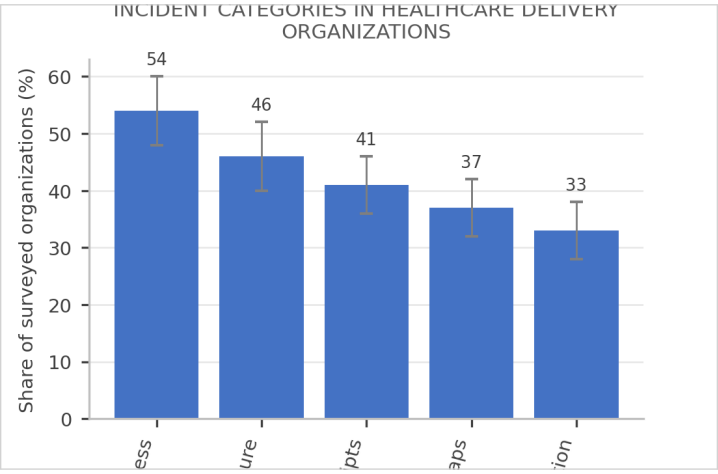
Framework (jurisdiction)	Cloud vendor's formal role	Technical accountability mechanism incentivized	Documented limitation	Key reference
Cloud stewardship maturity frameworks (cross-jurisdictional)	Cloud data repository operator	Outcome-based assessment of access transparency, provenance, and data-subject responsiveness	Most current repositories satisfy only a subset of assessed outcome domains	Rahimzadeh et al. 2026

Systematic legal reviews of health cloud governance frameworks find that jurisdictions vary considerably in whether they impose direct statutory security and privacy obligations on the cloud host or rely on the client institution to enforce compliance contractually, and that this variation is a major determinant of how effectively a given regulatory model can compel a cloud vendor to adopt privacy-preserving architecture in the first place, since a vendor facing only indirect contractual liability has a weaker independent incentive to invest in costly cryptographic or access-control engineering than one facing direct regulatory exposure [3]. China's Personal Information Protection Law illustrates a third regulatory model, treating health data as both sensitive personal information requiring heightened consent and potentially "important data" subject to government security review and data-localization requirements before any cross-border transfer — a framework that substantially increases the state's own role in verifying that a cloud host's technical architecture meets security requirements, at the cost of reducing the operational autonomy that a hosting institution and its cloud vendor would otherwise have to negotiate technical and contractual terms bilaterally [17]. Genomic and other large-scale biomedical cloud data stewardship research extends this compliance-architecture logic into a maturity-framework approach, proposing that responsible cloud data stewardship be evaluated against defined outcome domains — including access transparency, provenance tracking, and demonstrable technical responsiveness to data-subject requests — rather than treated as a binary compliant or non-compliant status; applying this framework, the authors found that most current cloud-hosted biomedical data repositories satisfy only a subset of these outcome domains in practice, indicating that regulatory compliance on paper does not reliably predict the technical maturity of the underlying privacy architecture [18]. A verifiable governance architecture proposed for encrypted cloud storage more broadly extends the same logic to general-purpose commercial cloud infrastructure, using cryptographic logging and attestation so that a data controller, an external auditor, or in principle a patient can independently verify which party accessed a given piece of encrypted data and under what authorization, without needing to trust the cloud host's own self-reported audit trail [19]. Table 2 compares how the major regulatory frameworks reviewed here differ in the technical accountability mechanisms they require or incentivize cloud vendors to adopt.

Vendor and Sub-Processor Risk in the Third-Party Cloud Ecosystem

Even a well-designed privacy-preserving architecture protects only the data flows it is built to cover, and a growing empirical literature documents that a substantial share of real-world privacy exposure in cloud-hosted patient data occurs outside the primary electronic health record hosting contract altogether. A large survey of healthcare delivery organizations found that unmanaged vendor and sub-vendor access is now the single most frequently reported category of third-party cybersecurity incident, ahead of delayed breach notification and analytics-related exposure, and that many organizations lack a current, comprehensive inventory of which vendors and sub-vendors hold access to which categories of patient data — a basic precondition for applying any of the encryption, access-control, or compliance architectures reviewed above to the full scope of an institution's actual data exposure [10]. This finding is corroborated by analyses of hospital website tracking technologies, which found that a substantial majority of hospital websites transmit potentially identifying visitor information to third-party advertising and analytics vendors under contractual arrangements that fall well short of a business associate agreement, exposing the hospital to a category of privacy and legal risk that originates entirely outside its primary cloud EHR vendor relationship and that no patient-record encryption scheme is designed to address [11].

Reported categories of third-party and vendor-related cybersecurity incidents affecting healthcare delivery organizations, illustrating that residual privacy risk concentrates outside the primary cloud electronic health record hosting contract (Gellert et al., 2025).



Federated and multi-institutional health data-sharing networks introduce a related but distinct governance challenge: legal and empirical analysis of anonymization and access-control practices in these networks finds that responsibility for a downstream re-identification or misuse event is frequently difficult to trace to a specific party, even where privacy-preserving architecture is nominally in place, because obligations are distributed across multiple contributing institutions and intermediary infrastructure providers in ways that formal governance documents do not always resolve unambiguously [20]. Quantitative evaluation of layered governance frameworks — combining ISO information security management standards with GDPR and HIPAA contractual requirements on top of blockchain-based technical infrastructure — finds that healthcare organizations applying multiple overlapping frameworks simultaneously report measurably higher data governance effectiveness scores than organizations relying on any single framework, but also report correspondingly higher administrative and compliance-monitoring burden, a direct empirical illustration of the trade-off between governance rigor and operational cost that pervades this literature [21]. Finally, patient-facing consent architecture remains a persistent gap in

the vendor risk landscape: survey research on dynamic, purpose-specific consent finds that patients strongly prefer the ability to revise or withdraw permission for specific downstream uses of their data, but that this preference is difficult to honor reliably once data has passed through a multi-vendor cloud hosting chain in which the party operationally holding the data is often several contractual steps removed from the party that originally obtained consent — a structural mismatch that neither encryption nor regulatory compliance architecture alone currently resolves [22]. Figure 2 summarizes the reported categories of third-party and vendor-related cybersecurity incidents affecting healthcare delivery organizations, illustrating where the residual risk concentrates once primary-vendor encryption and compliance architectures are in place.

Discussion

Three general conclusions follow from this review of privacy-preserving architectures for patient data hosted on third-party commercial cloud platforms. First, the technical maturity of the architectures reviewed varies considerably by category. Attribute-based encryption and blockchain-mediated access control for electronic health records have moved furthest from theoretical proposal toward evaluated implementation, with published systematic reviews now able to compare accountability properties across a substantial number of deployed or near-deployed systems [13,5]. Homomorphic encryption and secure multiparty computation, by contrast, remain constrained chiefly by computational overhead — several seconds of latency per kilobyte for fully homomorphic operations, and materially higher cost than federated analytics for SMPC — such that the evidence base for both is still dominated by pilot and pathfinder implementations rather than production deployment across large health systems [12,9,8]. Second, no single architectural approach reviewed here addresses the full scope of privacy risk that third-party cloud hosting introduces. Encryption and access control protect data at rest and in use against an untrusted or compromised cloud host, but they do not by themselves resolve the sub-processor visibility gap, the consent-propagation gap, or the governance ambiguity that arises once data crosses institutional and vendor boundaries in a federated network [10,20,22]. Third, regulatory compliance architecture and technical privacy-preserving architecture are increasingly convergent rather than parallel tracks: the more recent work reviewed here — machine-readable GDPR compliance verification, cryptographically verifiable access governance, and outcome-based cloud stewardship maturity frameworks — treats regulatory compliance itself as something to be encoded into and verified by technical architecture, rather than checked against it after the fact [7,19,18].

These findings carry direct implications for health policy. Regulators overseeing cloud-hosted patient data should consider requiring covered entities and controllers to maintain a current, auditable inventory of every party — including sub-processors and non-EHR vendors such as website analytics providers — with technical access to patient data, since the evidence reviewed here indicates that a substantial share of documented third-party incidents occurs precisely where such inventories are absent or outdated [10,11]. Procurement and certification standards for cloud vendors serving the health sector could reasonably distinguish between vendors that rely solely on contractual and organizational privacy commitments and vendors that have implemented independently verifiable technical architecture — cryptographic access logging, attribute-based encryption, or machine-readable compliance monitoring — since the latter category reduces the degree to which enforcement depends on trusting the vendor's own self-reported compliance [19,7]. Finally, given that federated learning and secure multiparty computation both reduce the amount of raw patient data that must be transferred to or pooled within third-party cloud infrastructure in the first place, health policy that funds or mandates pilot implementations of these architectures for specific high-sensitivity data categories — genomic data, mental health records, and multi-institutional research cohorts among them — could plausibly reduce aggregate privacy exposure more directly than incremental strengthening of contractual terms applied to a fully centralized, unencrypted cloud architecture [15,16,18].

This review has limitations inherent to its narrative rather than systematic design. The literature search, while targeted and individually verified for each included source, was not exhaustive, and it likely under-represents proprietary vendor-published performance data and national-level legal scholarship outside the United States, European Union, and China. The performance figures drawn from individual primary studies — encryption latency, decryption accuracy, governance-effectiveness scores — are heterogeneous in method, hardware, and dataset, and are presented as illustrative of documented directional patterns rather than as a pooled meta-analytic estimate. Finally, because both applied cryptography and cloud health data regulation are advancing rapidly, conclusions about the maturity of specific architectures, particularly homomorphic encryption and federated learning, may understate progress made after the search period underlying this review.

Conclusion

Privacy-preserving architecture for patient data hosted on third-party commercial cloud platforms has advanced along two largely complementary tracks: cryptographic and systems-engineering techniques that reduce how much trust must be placed in the cloud host itself, and regulatory or compliance architectures that translate legal obligation into auditable technical requirement. Attribute-based encryption, blockchain-mediated access control, and zero-trust network design have reached a level of technical maturity that supports near-term deployment at scale, while homomorphic encryption and secure multiparty computation remain constrained by computational overhead that current pilot implementations are only beginning to close. None of these architectures, however, resolves the governance gap that opens once patient data crosses institutional and vendor boundaries — the sub-processors, analytics integrations, and federated network participants that fall outside the scope of the primary hosting contract and that empirical vendor-risk research identifies as the leading source of documented third-party exposure. Realizing the full privacy benefit of the technical architectures reviewed here will require health policy that extends regulatory oversight and vendor accountability requirements to match the actual, and still expanding, perimeter of the cloud ecosystem in which patient data now resides.

References

1. Schweitzer EJ (2012) Reconciliation of the cloud computing model with US federal electronic health record regulations. Journal of the American Medical Informatics Association, 19(2):161-165, DOI:10.1136/amiajnl-2011-000162.

2. Rodrigues JJPC, de la Torre I, Fernandez G, Lopez-Coronado M (2013) Analysis of the security and privacy requirements of cloud-based electronic health records systems. Journal of Medical Internet Research, 15(8):e186, DOI:10.2196/jmir.2494.

3. Zandesh Z, Ghazisaeedi M, Devarakonda MV, Sayad Haghighi M (2019) Legal framework for health cloud: a systematic review. International Journal of Medical Informatics, 132:103953, DOI:10.1016/j.ijmedinf.2019.103953.

4. Munjal K, Bhatia R (2023) A systematic review of homomorphic encryption and its contributions in healthcare industry. Complex & Intelligent Systems, 9(4):3759-3786, DOI:10.1007/s40747-022-00756-z.

5. Li M, Yu S, Zheng Y, Ren K, Lou W (2013) Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption. IEEE Transactions on Parallel and Distributed Systems, 24(1):131-143, DOI:10.1109/TPDS.2012.97.

6. Edo OC, Ang D, Billakota P, Ho JC (2024) A zero trust architecture for health information systems. Health and Technology, 14(1):189-199, DOI:10.1007/s12553-023-00809-4.

7. Zhou C, Barati M, Shafiq O (2023) A compliance-based architecture for supporting GDPR accountability in cloud computing. Future Generation Computer Systems, 145:104-120, DOI:10.1016/j.future.2023.03.021.

8. Ballhausen H, Corradini S, Belka C, Bogdanov D, Boldrini L, Bono F, Goelz C, Landry G, Panza G, Parodi K, Talviste R, Tran HE, Gambacorta MA, Marschner S (2024) Privacy-friendly evaluation of patient data with secure multiparty computation in a European pilot study. npj Digital Medicine, 7(1):280, DOI:10.1038/s41746-024-01293-4.

9. Kerr S, Robertson C, Sudlow C, Sheikh A (2025) Enabling health data analyses across multiple private datasets with no information sharing using secure multiparty computation. BMJ Health & Care Informatics, 32(1):e101384, DOI:10.1136/bmjhci-2024-101384.

10. Gellert GA, Borgasano D, Palermo R, Gellert GL, Kelly SP (2025) Third-party access cybersecurity threats and precautions: a survey of healthcare delivery organizations. Applied Clinical Informatics, 16(5):1518-1530, DOI:10.1055/a-2713-5725.

11. Friedman AB, Merchant RM, Maley A, Farhat K, Smith K, Felkins J, Gonzales RE, Bauer L, McCoy MS (2023) Widespread third-party tracking on hospital websites poses privacy risks for patients and legal liability for hospitals. Health Affairs, 42(4):508-515, DOI:10.1377/hlthaff.2022.01205.

12. Olaymi SEZ (2025) Performance and security analysis of fully homomorphic encryption in cloud-based healthcare blockchain. Health Informatics Journal, 31(4):14604582251394616, DOI:10.1177/14604582251394616.

13. Cobrado UN, Sharief S, Regahal NG, Zepka E, Mamauag MB, Velasco LC (2024) Access control solutions in electronic health record systems: a systematic review. Informatics in Medicine Unlocked, 49:101552, DOI:10.1016/j.imu.2024.101552.

14. Sun J, Ren L, Wang S, Yao X (2020) A blockchain-based framework for electronic medical records sharing with fine-grained access control. PLOS ONE, 15(10):e0239946, DOI:10.1371/journal.pone.0239946.

15. Alankamony A, Nels N (2026) Systematic review of privacy preservation in federated learning for secured healthcare applications. Information, 17(7):647, DOI:10.3390/info17070647.

16. Kerr S, Escudero D (2026) Introduction to secret sharing based secure multiparty computation in health data. BMJ Open, 16(5):e110916, DOI:10.1136/bmjopen-2025-110916.

17. Yin Q (2025) Data privacy rights in health data sharing: a government-centric paradigm in China. International Data Privacy Law, 15(3):238-255, DOI:10.1093/idpl/ipaf007.

18. Rahimzadeh V, Walsh B, Rehm HL, Cho M, McGuire AL (2026) Characterizing core outcomes of responsible stewardship for human genomic data in the cloud. BMC Medical Genomics, 19(1):105, DOI:10.1186/s12920-026-02382-x.

19. Rodriguez-Aragon JF, Zato C, Pinto-Santos F, Sanchez-Pravos L (2026) KATENA: a verifiable governance architecture for encrypted cloud storage systems. Frontiers in Big Data, 9:1829960, DOI:10.3389/fdata.2026.1829960.

20. Kogut-Czarkowska M, Shabani M (2026) Anonymization, accountability, and access: legal dimensions of health data sharing in federated networks. Perspectives from empirical study. Frontiers in Digital Health, 8:1719728, DOI:10.3389/fdgth.2026.1719728.

21. Ahmed A, Shahzad A, Naseem A, Ali S, Ahmad I (2025) Evaluating the effectiveness of data governance frameworks in ensuring security and privacy of healthcare data: a quantitative analysis of ISO standards, GDPR, and HIPAA in blockchain technology. PLOS ONE, 20(5):e0324285, DOI:10.1371/journal.pone.0324285.

22. Muller SHA, van Thiel GJMW, Mostert M, van Delden JJM (2023) Dynamic consent, communication and return of results in large-scale health data reuse: survey of public preferences. DIGITAL HEALTH, 9:20552076231190997, DOI:10.1177/20552076231190997.