Deep Science Publishing
https://doi.org/10.70593/978-81-988918-5-3_5

DeepScience
Open Access Books

Chapter 5: Establishing best practices for securing sensitive healthcare information in the cloud

5.1. Introduction

Using Cloud Services to Store and Process Electronic Health Information Healthcare organizations increasingly are turning to encrypted cloud services that offer the convenience effect associated with the tradeoff between storing sensitive healthcare information on their own servers and utilizing a commercially operated cloud service. Concerns about strict, pervasive security protections in HIPAA and other healthcare regulations are being addressed by creative solutions and technical innovation. New cloud-based healthcare information technology services are helping organizations, physicians, and patients manage the massive volumes of electronic health information being generated and consumed every day (Dilsizian & Siegel, 2014; Lee & Yoon, 2017; Davenport & Kalakota, 2019). Cloud services address the often challenging issues of system accessibility and downtime by allowing electronic health information to be stored and retrieved from remote servers accessible through the Internet. Virtualization technology, coupled with a variety of increasingly affordable storage options, allows cloud service providers to store significant amounts of data in a cost efficient manner, keeping subscription fees lower and making affordability a non-issue for many organizations. Yet, providing security protections, user access controls, monitoring, and identification of data breaches, while normal responsibilities of any organization holding sensitive data, take on added complexity when the data resides on a third-party's server. Cloud service customers are responsible for protecting the security and integrity of their sensitive information and must take extra precautions. Both government and private sector cybersecurity agencies and organizations have issued recommendations on steps to take to ensure the security of sensitive healthcare information stored in the cloud, which are discussed in more detail later. What-if scenarios regarding vulnerability to breaches of sensitive information, the associated potential liability, and the value of

reputation for excellent cybersecurity must weigh heavily into the decision whether to utilize cloud services to store and process electronic health information.That is why establishing best practices for managing sensitive healthcare information in the cloud can help organizations to develop effective security risk management processes. Understanding the paths by which security breaches may occur and the harm they could cause is a fundamental component of any risk assessment and decision-making, yet medical organizations may struggle even to identify their most sensitive data, let alone create specific methodologies or best practices for its protection. The available cloud security guidelines and standards may be well known or even widely adopted. However, little guidance is available to assist organizations through the process of protecting sensitive healthcare information stored in the cloud, especially when the information retains its sensitive status only while stored in the cloud. Without best practices, organizations may struggle even to identify their most sensitive data, let alone create specific procedures for its protection (Zhang et al., 2018; Reddy et al., 2019).

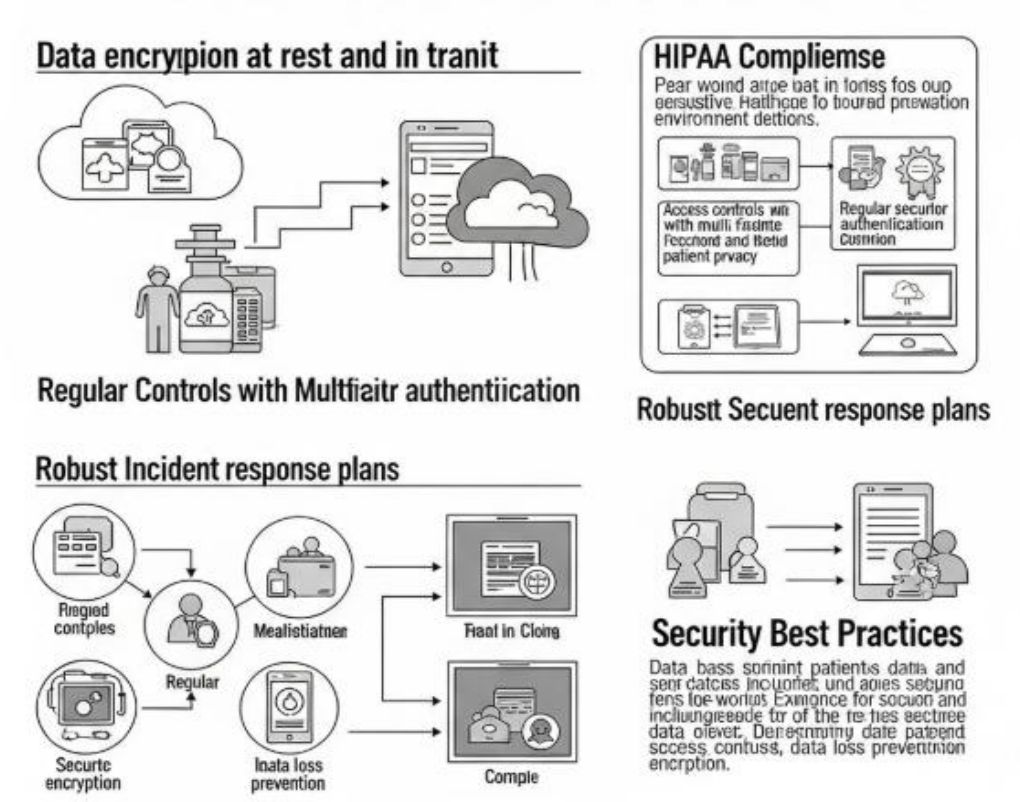

Fig 5.1: Establishing best practices for securing sensitive healthcare

5.1.1. Background and Significance

Securing sensitive healthcare information in the cloud goes beyond addressing cybersecurity concerns or complying with standards and regulations; it is about protecting patients and building trust. The Security Rule mandates that covered entities ensure the confidentiality, integrity and availability of electronic protected health information. To meet this mandate, organizations must conduct security risk analysis to identify and assess potential risks and vulnerabilities to the confidentiality, integrity and availability of electronic protected health information, assess their current security measures, and document those assessments. However, completing a risk assessment is often a confusing process. Many organizations fail to implement appropriate solutions or their processes may fall short because of lack of time or human resources; in addition, they may lack formal risk assessment methodology. Security risk assessments are often conducted solely to satisfy the compliance requirement. As a result, many organizations may possess little knowledge of the risks they face, the vulnerabilities they carry, or the safeguards that will ensure necessary protection.

5.2. Understanding Cloud Computing in Healthcare

Cloud computing can provide many advantages across a variety of industries, allowing healthcare organizations to increase efficiency and population health while simultaneously reducing costs. Cloud solutions allow organizations to respond quickly to changing demands by drawing on pooled resources. For example, cloud computing has the ability to store and share vast amounts of data, which is an increasing feature for many healthcare organizations as they gather and use patient information. Moreover, organizations can access data on demand anywhere, via mobile devices. Whereas traditional technologies require major capital funding for equipment and expertise, cloud computing lowers both costs and barriers of entry.

With cloud computing, private cloud solutions offer single-organization proprietary security, while community cloud solutions allow groups of related organizations to share the costs of services in a secure manner. Hybrid solutions take advantage of public cloud solutions for an organization’s non-sensitive data storage and transfer, using private solutions for sensitive information. Cloud services offer fast expansion of data storage needs but only for non-sensitive data, avoiding major investment in non-committed data needs. Outsourced cloud computing features certified vendors and contracts governing both the vendor's and the insuring organization's security responsibilities. The attractiveness of using the cloud for healthcare-specific processes is tempered by the inability of organization leadership to control all the factors associated with sensitive data storage across operations. Security issues combine with the requirement for Business Associate Agreements to limit cloud-based services for data.

5.2.1. Research design

The knowledge structure of cloud computing has been classified into a hierarchy encompassing the following categories: cloud computing basics, different service and deployment models, technology architecture, significant enablers and platforms, technical security, privacy, and legal issues. The taxonomy has further been expanded and clarified to include other major areas, such as advantages and drivers, challenges and concerns, issues demanding future research, and organizations that can play a role in the promotion and development of cloud computing. The cloud computing adoption landscape has also been reviewed with a focus on the motivations and barriers, including issues of trust, as well as several decision-making models that can be employed to assist organizations in the selection of suitable service providers.

Built upon the above knowledge structure, our research employs an inductive approach using qualitative interviews with healthcare cloud experts, namely experts with experience in using cloud services to store or process sensitive health information. Our objective is to identify and address the most pressing areas of concern around security and privacy of sensitive health information in the cloud, as articulated by these experts. We interviewed four cloud computing experts from different organizations located in the United States but working for both local and international organizations. They have diverse backgrounds in medical, legal, technical, and governmental aspects of cloud computing in healthcare, which allowed us to proscribe a wide span of general and specific questions around the use of cloud services to store and process sensitive healthcare data.

Our participants were recruited using a snowball sampling approach. We began with initial contacts at local academic institutions but gradually expanded our search, and these initial contacts were able to point us towards some key informants, both from the local area and around the United States. In total, we conducted 14 semi-structured interviews that lasted around an hour each. The diversity of our expert participants allows us to collect a large volume of valuable data on this important area, as reflected in some of their background information presented.

5.3. Regulatory Frameworks and Compliance

The healthcare field is subject to the most stringent of regulatory policies, a fact of which organizations using cloud service partners for sensitive health data should be aware. A crucial first step is to understand the technical and legal details of major health risk policies. Organizations should be aware of how these regulations are actually enforced, and at what level of granularity. Clearly privacy violations and data leaks are to be avoided, but the severity of the penalties involved if an organization is discovered to be

in noncompliance varies greatly according to the specific regulation, affected parties, and other factors, such as whether a ruling allows for harm to be ongoing or whether it requires remediation for past actions. Initial evaluations should be at a high level, which is essential for determining if major issues exist that would require a deeper investigation. These evaluations can then use the results of these initial inquiries to create a prioritized action and remediation plan.

The key parts of the core principles about sensitive health data processing are summarized in the Security Rule, Privacy Rule, Breach Notification Rule, and Enforcement Rule of the law. The applicable jurisdiction must be clearly understood as it applies to covered entities and their business associates who handle protected health information. The cloud service partner may itself fall under the category of a business associate, or may instead be used as a subcontractor by the business associate, called a downstream business associate. In either case, organizations must ensure that a Business Associate Agreement is in place for all such partnerships.

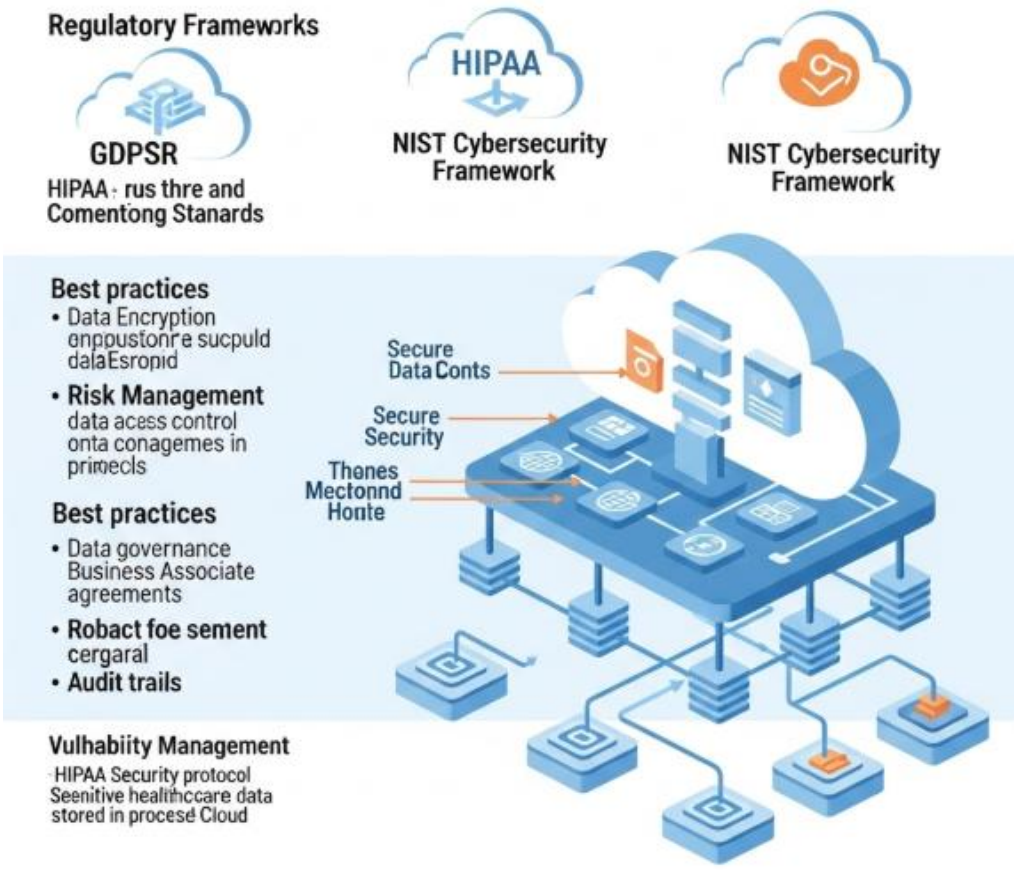

Fig 5.2: Regulatory Frameworks and Compliance

5.3.1. HIPAA Overview

The Health Insurance Portability and Accountability Act of 1996 is a set of regulations intended to protect sensitive healthcare information from being disclosed without the patient's consent or knowledge. The Privacy Rule establishes national standards to protect individuals' medical records and other personal health information provided to health plans, doctors, hospitals, and other healthcare providers. The Privacy Rule also gives patients rights over their own health information, including the rights to examine and obtain copies of their health records and to request corrections. Also, the Security Rule sets national standards for the security of electronic protected health information. The Security Rule requires appropriate administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and security of electronic protected health information.

HIPAA only applies in certain situations. To be subject to its rules, a company must be a “covered entity,” as defined in the statute. Covered entities include health plans; healthcare clearinghouses; and healthcare providers who conduct certain transactions electronically and who transmit those transactions in standardized form. In addition, some companies that do business with covered entities are "business associates" and must also comply with some HIPAA rules; for example, a cloud service provider that has access to electronic protected health information stored on its systems would be a business associate. Business associates are required to sign a Business Associate Agreement in order to be contractually bound to protect electronic protected health information in accordance with HIPAA.

5.3.2. GDPR Implications

The General Data Protection Regulation (GDPR) is a European Union regulation that establishes a comprehensive framework for the collection and processing of personal information. The GDPR was adopted in April 2016 and has been in effect since May 2018. The regulation was established to address the fact that people have lost control over their personal information in a world that’s increasingly digitized. GDPR applies to any organization that does business in the EU or collects files on an EU citizen, regardless of the country in which the organization is located.

The GDPR applies to any organization that does business in the EU or collects files on an EU citizen, regardless of the country in which the organization is located. The healthcare sector has its own specific regulatory standards and requirements, but the GDPR may apply either as a standalone regulation or may be incorporated into the entire compliance program that covers all laws and regulations affecting sensitive data, including PHI, PII, and PCI. Patient information will be part of both GDPR EU and

outside of GDPR for EU citizens' business. If any of the sensitive data being shared is governed by both regulations, it would then have to follow what will be two sets of rules - GDPR for specific sensory data elements, and HIPAA for everything else. Organizations that are collecting or being responsible for the PII or the PHI would then need to create controls to maintain compliance with both sets of requirements.

The GDPR has stricter requirements governing access to and the processing of personal data and also includes fines for violations that are not only enormous but can also be imposed at a staggering rate. Fines for breaching GDPR's security requirement can amount to €20 million or 4% of a company's annual global revenue. The GDPR's security requirement also imposes strict regulations on reporting data breaches. Organizations that use the cloud for their business operations generally delegate some data security responsibilities to their cloud service providers.

5.4. Risks Associated with Cloud Storage

While many healthcare organizations are indeed moving to the cloud for the storage and processing of sensitive data, their understanding of such a transition remains inadequate. Risk factors that are inherent with any transition away from dedicated on-premises IT infrastructure still remain. This includes a lack of understanding of regulatory requirements that govern cloud storage of controlled unclassified data, such as ePHI. While the cloud service models relieve organizations of some typical burdens, not all responsibilities are transferred to the cloud service provider. As a result, non-compliance issues, resulting either from a lack of knowledge as to what tasks the cloud service provider performs and what tasks the organizations remain responsible for or a lack of communication between these entities, may result in heavy fines and liabilities for both partners. Further, many organizations view cloud services as a panacea toward all security-related threats, either underestimating the strength of inherent risks, such as cyber attack and data breach, or completely ignoring them.

When ePHI is stored and processed within the cloud and specifically a multi-tenant cloud system, the paradigms for normal threat assessment change dramatically. Data stored in the cloud becomes subject to a revised set of different processes to achieve authorization and authentication. The reliability of the cloud service provider becomes an integral part of the healthcare organization’s information security processes and procedures. The concepts of data ownership, access, and protection of sensitive data change significantly when applied in the cloud model. The storage, processing, and transmission of large amounts of sensitive data involve a variety of different risks and an increased attack surface. In particular, financial data, personal data, locations, and sensitive dates serve as great targets for cyber attackers. Acquiring ePHI is often combined with medical identity theft or health insurance fraud.Thus, the temperature monitored over time is key

to determining whether that tunnel actually has become an unauthorized supplementary path between the user and an outside network domain — especially if that domain happens to be close to certain sensitive target thresholds. These activities generally exist behind the perimeter defenses; the actual packet flows may not leave the local network and communications such as queries may not trigger anomaly alarms on routers are part of the security posture. However, selective logging to enable vulnerability scanning of high-profile accounts can be performed to alert security at any site.

5.4.1. Data Breaches

There are more than 150 well-documented data breach incidents involving patient records in just the last five years, reportedly involving loss of over 1.5 million records. While many of these incidents had nothing to do with cloud computing, the increasing maturity of cloud services means that healthcare organizations are moving more sensitive patient data into the cloud, and it’s not simply replicated staging data. A single breach incident can expose thousands of sensitive patient records. Security experts have recently expressed increased awareness over cloud storage and the risk of data breaches, now leading to illegal patient record enumeration and sharing. Even with the sophistication of today’s client-side libraries, securely accessing cloud-hosted services requires careful review so that corporate credentials do not end up in the cloud storage account.

Treat patient data residing in the cloud as any other data that needs to be secured. Access should only be given to patient data in cloud storage as needed, least privilege ought to be exercised, and all data access should be logged and audited. If patient data needs to be shared in the cloud, investigate third-party solutions that address these concerns. There are now several services that enable secure sharing of files in cloud storage, and the financial implications of a breach are substantial. A breach of data stored can incur substantial billable amounts, and the incident response is both critical and costly. Any involved cloud service provider has obligations at law to protect a patient’s information, and those obligations extend to preventing, detecting, and responding to a breach in a timely manner.

5.4.2. Insider Threats

Yet fundamentally, this crack is an internal threat. So, while we've defined key employee groups, based on their jobs and responsibilities, watching what they do is a good first step, it isn't enough. This enables management to track the major changes in an employee's behavior, and in the long run, it can help identify those employees who are planning to leave. And, of course, it's a two-way street: disgruntled employees can use

their privileged access to collect proprietary information such as customer databases, business plans, and trade secrets, all of which are invaluable to competitors. So, while these actual incidents involving insiders describe babies that grew up to be bad, it's very important to remember that the key influential condition is simply not changing your signature of trusted behavior. A key component to detecting, as opposed to preventing and denying, bad intent is determining in advance the expected set of behaviors. By monitoring specific items, it becomes possible to detect if a user suddenly decides to shift his or her browsing from, for example, clothes and cars to sensitive topics of a subversive nature. It may even be that user tunneling traffic to a site while internal security folks are blocked from even monitoring the breach is legitimate — it might just mean that the user is doing research. But it's a red (or possibly even yellow) flag nonetheless.

5.5. Key Security Principles

Sensitive healthcare environments contain protected health information (PHI), intellectual property, trade secrets, payment card data, research data, or a combination of these. Access to such information should be limited based on the principle of least privilege to people with a need to know. Sensitive information must be kept from unauthorized parties to avoid consequence and remain secure while preserving its utility, privacy, and authenticity. To achieve this goal, disparate technologies interact to ensure PHI is encrypted, accessible only to approved individuals and entities, and that there is a chain of custody throughout information exchange. Many of these technologies are incorporated into technology safeguards established by the Department of Health and Human Services.

The principle of confidentiality describes how PHI must be accessible only to authorized individuals. This mandate preserves patient trust, safeguards the security of the healthcare system, and prevents a host of consequences for unauthorized disclosure. Confidentiality of PHI and related sensitive healthcare data is further protected by donor wishes for organ transplantation, information provided to faculties of medical emergencies, and other sensitive information systems. Authorized users are granted access based on the systems they administer, the role they fill, and the processes they need to conduct. For example, ruling out prostate cancer requires specific laboratory tests that are conducted by authorized personnel, and the resulting data is accessible to those who have a need to know. The principle of least privilege goes further. Individuals should be given the minimum access necessary to accomplish tasks. For example, upon hiring or terminating an employee, much input may be required from different departments to effectuate their onboarding or offboarding process, but the IT department should assign minimal access based on that person’s role.

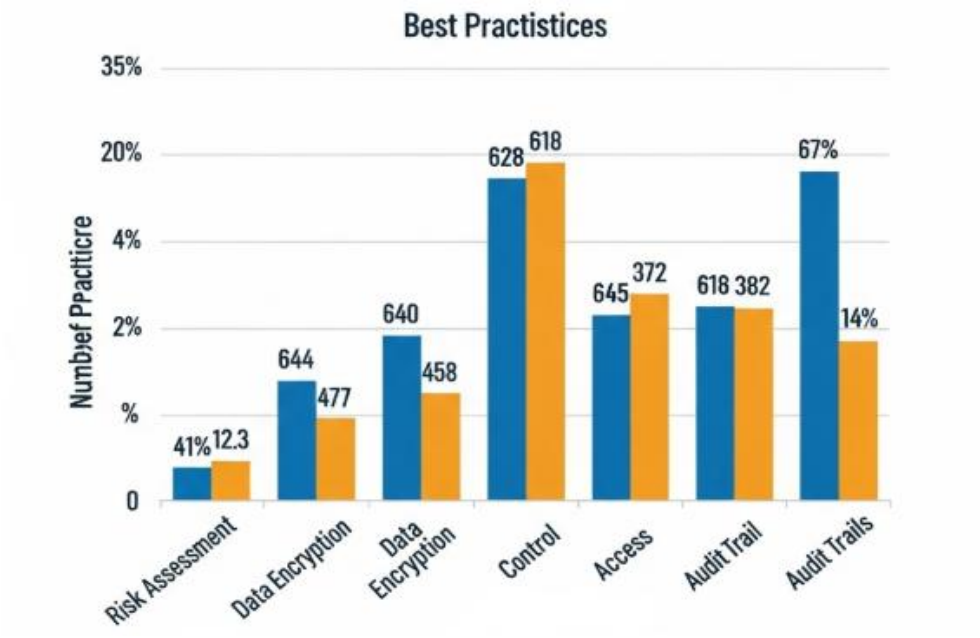

Fig : Establishing Best Practices for Securing Sensitive Healthcare Information in the Cloud

5.5.1. Confidentiality

When discussing sensitive healthcare information, particularly when allowing third-party storage in the cloud for such information, confidentiality is the most important principle. It is not enough that only the people with a right to see the information are able to see it. We must also, by law, prevent even the third-party data holders from being able to see any of the information we have entrusted to them. Non-disclosure agreements, background checks and physical security cannot accomplish that – only encryption can.

Organizations storing sensitive healthcare information in the cloud should use encryption such that only the organization keeping the sensitive healthcare information can ever un-encrypt that information. Furthermore, the encryption should be of high quality, having been recommended or approved by the most trusted cryptographic experts. The encryption should also be of a type that the most trusted cryptography experts are confident will be able to keep secrets hidden from even quantum computers with all their enormous power. Simply relying on cloud providers to encrypt with whatever weak, obsolete or broken crypto scheme they prefer would be allowing other organizations to have access to secrets we are legally and ethically required to keep from them. And even with these weak, obsolete or broken schemes, just relying on the cloud provider doing that encryption is also insufficient, for numerous reasons.

The reasons for this are actually rather complex, but to greatly simplify, the data owner is not the only party that has legal obligations to protect a data subject's privacy. The fact that transactions may involve multiple parties may impose obligations on other parties that make joint action necessary. The parties sharing data in a cloud need to mutually agree on encryption policies and the means for enforcing them. Non-disclosure agreements may only be effective in simple operations. If any party's interests become misaligned, that party might leverage its possibly unique access to all data to breach another party's legal obligation.

5.5.2. Integrity

Integrity is assuring the receiver of information that it originated from a trusted source, has not been modified in transit, and remains whole when received. In many cases, protecting the integrity of data may require additional overhead to be added to a cloud application and to the infrastructure it executes upon. Adding or verifying the integrity of data incurs some extra performance overhead to the application. For the cloud, the integrity of data often means a one-way hash. The sender of data to the cloud hashes the file, stores the hash value somewhere secure, and then sends the data to the cloud and allows the cloud to store it. Later, the sender can request the data from the cloud and hash it again, trusting that the two hash values match. Integrity not only protects against data modification but also ensures that the data received comes from the sender and has not been tampered with in transport. What is most important for integrity in the cloud is the use of cryptographically strong hash functions.

Maintaining data integrity also requires access control. Well-documented procedures must exist to assign the proper level of access to all data and systems, as well as an ongoing need to conduct administrative review and an automatic or manual authorization process for additional access periods. Additional non-technical considerations are also for data integrity. One consideration is the human factor; negligent or poorly trained employees are the leading cause of integrity breaches. Another consideration is documentation. The Integrity Documentation Plan must contain detailed plans for specific actions of how data and system integrity will be maintained.

5.6. Conclusion

Healthcare organizations should adopt a well-defined and comprehensive cloud security framework governing at least the Eleven Practices we covered above. The Eleven Practices comprise essential aspects and have made critical contributions toward the adoption of sensitive healthcare information in the cloud. Further, while individual elements of each of the eleven can bolster protection efforts, it is their deployment as a

cohesive architecture that will maximize defense against cloud-based cyberattacks. The framework presented here can benefit the sensitive healthcare data cloud ecosystem not only by enhancing the security posture of the healthcare industry but also by leveraging cloud infrastructure providers' capabilities, which they can scale, optimize, and continually improve.

Cloud computing is fundamentally a model for managing computing resources in a more centrally managed, automated manner to lower overall costs through economies of scale and enabling greater strategic use of technology. Security is and will continue to be a concern for cloud computing users, especially for the Healthcare industry. The emergence of cloud computing has facilitated sharing publicly available health information and may represent a solution to the insufficient participation in health-related electronic social networks. Nevertheless, the sensitivity of data in cloud computing in Healthcare makes some users hesitant to give access to their data to third-party cloud providers. Thus, it is essential the service agreements support the health users in protecting the service from threats and vulnerabilities, helping overcome the existing reluctance to use cloud computing.

Several tools and services are emerging that can help healthcare industry organizations with the challenges we discussed in this work. For example, several organizations have created products to help with identity and access management and logging. In turn, as CSPs continue to manage the hypervisor and control the servers underlying client VMs, they will likely need to further enhance their security posture.

5.6.1. Emerging Trends

Cybersecurity continues to evolve in an arms race that pits creative miscreants against dedicated developers who strive to keep pace, patching systems and programming in ways that make breaches harder to accomplish. The battle is growing more sophisticated and technical, leading to the prioritization of cloud architecture, policies, and frameworks that enshrine security by design. Miscreants are deploying AI as an assistant to carry out common cyberattacks and writing generative AI programs that assist coding and testing of applications. AI exploration is gaining steam and use in a variety of tasks; it augments development and operational tasks – some security-related, such as log analysis, incident detection, and user detection. These are efforts that security teams must always perform. AI enhances the productivity of those working with the tools, enabling analysts and engineers to work on higher-impact projects. Cloud service providers today prioritize security as not just a key feature but part of their architecture and a primary product pillar that differentiates their offerings. This emphasis is based on a number of trends. One is prioritizing technical investment in infrastructure and systems that provide secure foundational cloud and services, again, making security-by-design an enshrined

principle. Another is increased emphasis on assurance that a cloud service surface has been inspected and safeguarded in terms of your organization's unique requirements – that the services are running trusted code and are consistent in their operations, as services spin up and down in a demand-oriented manner. Cloud customers expect this emphasis to be part of their service contracts, complying with security, operational, and broader due diligence frameworks.

References

Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. Future Healthcare Journal, 6(2), 94–98.

Zhang, Y., Milinovich, G. J., Xu, Z., Bambrick, H., Mengersen, K., Tong, S., & Hu, W. (2018). Monitoring pertussis infections using internet search queries. Scientific Reports, 8(1), 1–10.

Lee, C. H., Yoon, H. J. (2017). Medical big data: Promises and challenges. Kidney Research and Clinical Practice, 36(1), 3–11. https://doi.org/10.23876/j.krcp.2017.36.1.3

Dilsizian, S. E., & Siegel, E. L. (2014). Artificial intelligence in medicine and cardiac imaging: Harnessing big data and advanced computing to provide personalized medical diagnosis and treatment. Current Cardiology Reports, 16(1), 441.

Reddy, S., Fox, J., & Purohit, M. P. (2019). Artificial intelligence-enabled healthcare delivery. Journal of the Royal Society of Medicine, 112(1), 22–28.