


HIPAA Compliance for Healthcare Workloads on IBM Spectrum Scale

Sandeep R. Patil

Sandeep Zende



 **Security**

Storage

In partnership with
IBM Academy of Technology



HIPAA background

When technology workloads process healthcare data, it is important to understand Health Insurance Portability and Accountability Act (HIPAA) compliance and what it means for the technology infrastructure in general and storage in particular. HIPAA is US legislation that was signed into law in 1996.

HIPAA was enacted to protect health insurance coverage, but was later extended to ensure protection and privacy of electronic health records and transactions. In simple terms, it was instituted to modernize the exchange of healthcare information and how the Personally Identifiable Information (PII) that is maintained by the healthcare and healthcare-related industries are safeguarded.

From a technology perspective, one of the core requirements of HIPAA is the protection of Electronic Protected Health Information (ePHI¹) through physical, technical, and administrative defenses. From a non-compliance perspective, the Health Information Technology for Economic and Clinical Health Act (HITECH) added protections to HIPAA and increased penalties \$100 USD - \$50,000 USD per violation. Today, HIPAA-compliant solutions are a norm in the healthcare industry worldwide.

This IBM® Redpaper publication describes HIPAA compliance requirements for storage and how security enhanced software-defined storage is designed to help meet those requirements. We correlate how Software Defined IBM Spectrum® Scale security features address the safeguards that are specified by the HIPAA Security Rule.

¹ Per HIPAA, compliance is “any electronic information that is created or received by a healthcare provider that relates to the past, present, or future physical or mental health of an individual and that identifies the individual”.

HIPAA compliance and storage

HIPAA compliance features overarching requirements that span the entire healthcare solution stack that deals with health-related information. One of the important building blocks of this solution stack is the underlying storage where this sensitive data (that is, ePHI) is stored. Hence, it is important to understand the HIPAA requirements for storage solutions that are used to host ePHI-related data.

Moreover, many institutes are deploying artificial intelligent (AI) workloads that are mining the health data to get meaningful insights for enhancement in healthcare and precision medicine. These AI workloads drive the need for Secure AI solutions that starts with addressing the HIPAA requirements at the storage level and then building up the Secure AI solution stack from there.

What are these requirements? Basically, the Privacy and Security Rule of HIPAA mandates that healthcare-based technology solutions ensure privacy for health information and must also comply with different types of safeguards that are defined by HIPAA. These rules are the primary set that you need to understand from a storage perspective.

Security enhanced software-defined storage

According to the Storage Network Industry Association (SNIA), data security in the context of storage systems is responsible for safeguarding the data against theft, prevention of unauthorized disclosure of data, prevention of data tampering, and accidental corruption. This process ensures accountability, authenticity, business continuity, and regulatory compliance.

Security for storage systems can be classified as shown in the following examples:

- ▶ Data storage (data at rest, which includes data durability and immutability)
- ▶ Access to data
- ▶ Movement of data (data in flight)
- ▶ Management of data

IBM Spectrum Scale

IBM Spectrum Scale is a software-defined storage system for high-performance, large-scale workloads that are on-premises or in the cloud. It was developed and enhanced by IBM for 30 years to address data management at massive scale for large, global organizations. It simplifies management, helps to lower capital and operational costs, is easy to grow, and includes important, enterprise class features.

Its most common use areas are in AI and deep learning, big data analytics, content repository, private cloud, and compute clusters. It is also commonly used for data optimization and resiliency for archive, high-speed back up, and disaster recovery, and Information lifecycle management.

IBM Spectrum Scale addresses HIPAA requirements by including the following features:

- ▶ Securing data at rest by protecting data at rest with snapshots, backups, and immutability features.

- Securing access to data by providing secure management of data, and securing access to data by using authentication and authorization across multiple supported access protocols. Supported protocols include POSIX, NFS, SMB, Hadoop, and Object (REST).
- Securing data in-flight by providing encryption of management and user data over wire on supported protocols.
- By automating data management, it is equipped with powerful information Lifecycle Management (ILM) tools that can help administer unstructured data by providing the correct security for the correct data.

Figure 1 shows the features of IBM Spectrum Scale security.

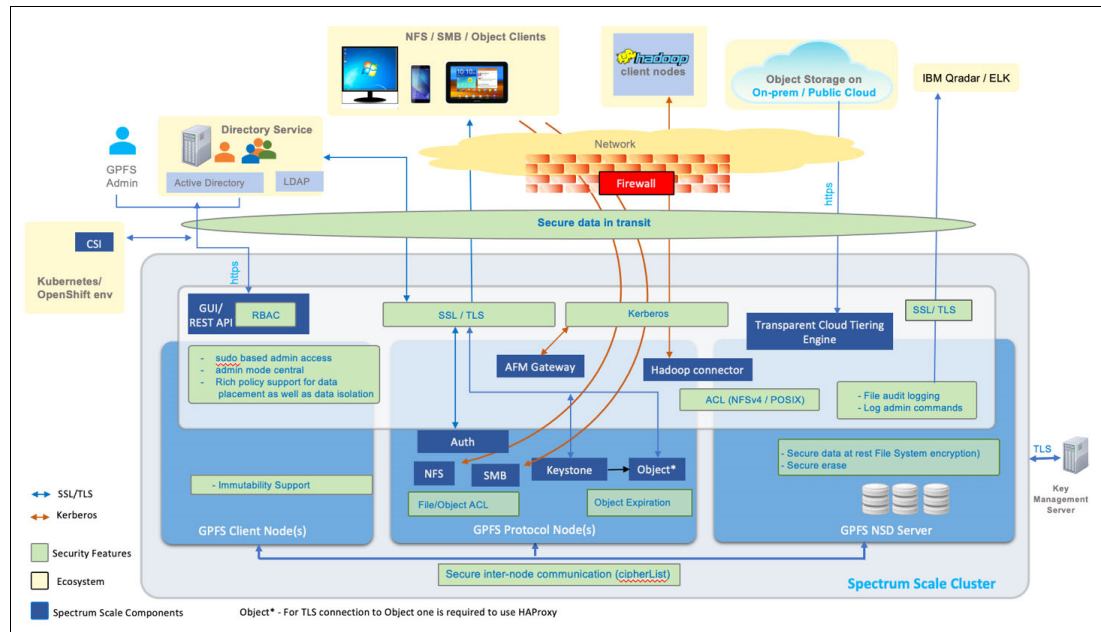


Figure 1 IBM Spectrum Scale security features

This IBM Redpaper publication explains the different safeguards that are specified by the HIPAA Security Rule. It also describes which IBM Spectrum Scale security features can be mapped to them. The Redpaper correlates IBM Spectrum Scale security features to the needs of the healthcare sector and discusses why it is the most suitable file and Object Storage for healthcare workloads.

HIPAA security rules and IBM Spectrum Scale security

The HIPAA Security Rule specifies three types of controls for compliance:

- Administrative Safeguards (Section - 164.308): To document formal policies and practices for data protection, including the organization's security management process and implementation specification.
- Physical Safeguards (Section - 164.310): To protect data from the hazards of fire, weather, the environment, or intrusion.
- Technical Safeguards (Section - 164.312): To protect data and control access to information by individuals and guard against unauthorized access through a network.

Overall, these safeguards require the healthcare industry to measure and keep personal health information secure and to decrease the means of tampering with, destruction of, or inappropriate access to ePHI. We filter and shortlist the specific safeguards that apply for storage systems while the adherence of the overall safeguards is responsibility of the overall solution and the practitioners owning the solution.

These safeguards are applicable to storage in the following ways:

- ▶ Administrative Safeguards (164.308): Semi-applicable to storage components
- ▶ Physical Safeguards (164.310): Minimal applicability to storage components
- ▶ Technical Safeguards (164.312): High applicability to storage components

Table 1 lists the filtered safeguards based on their applicability to file storage and the corresponding IBM Spectrum Scale security features. These features can be used by healthcare solution architects to address the safeguards from the storage perspective.

Important: Depending on the overall architecture of the healthcare solution, implementation and responsibility of the safeguards can be addressed at different levels of the solution stack. Therefore, many of the safeguards might not be directly applicable for underlying storage, even if the storage is hosting ePHI.

Although Table 1 provides a high-level mapping of the applicable safeguards to the security features of IBM Spectrum Scale, the actual applicability depends on the overall architecture of the healthcare solution and policies set by the practitioners.

Table 1 Filtered HIPAA safeguards and corresponding IBM Spectrum Scale security features

Parts and subparts of Section 164	Section heading and subheading	IBM Spectrum Scale security features
HIPAA 164.308(a)(4)(ii)(B,C)	Information access management (Access authorization and access establishment and modification)	IBM Spectrum Scale provides a rich set of POSIX and NFSV4 ACL for data authorization that can be used by solutions to address this safeguard.
HIPAA 164.308(a)(5)(ii)(B)	Protection from malicious software	IBM Spectrum Scale supports antivirus and integration with threat detection software, such as IBM QRadar, which can be used by solutions to address this safeguard.
HIPAA 164.308(a)(5)(ii)(D) and 164.312(a)(2)(i)	Password management and unique user identification (Procedures for centralized password management and centralized, unique ID for tracking)	IBM Spectrum Scale supports Active Directory and LDAP-based and Kerberos-based centralized authentication and ID management system for comprehensive password management and unique identification, which can be used by solutions to address these safeguards.

Parts and subparts of Section 164	Section heading and subheading	IBM Spectrum Scale security features
HIPAA 164.308(a)(1)(ii)(D), 164.308(a)(5)(ii)(C), and 164.312(b)	Information system activity review, logs, and audit control	IBM Spectrum Scale supports auditing management activity by using a CLI and GUI. It also supports file audit logging, which can be integrated with SIEM, such as IBM QRadar as and Elasticsearch. Solutions can use these features to address these safeguards.
HIPAA 164.312(d)	Person or entity authentication	IBM Spectrum Scale supports authentication by using corporate directory servers, such as AD/LDAP. This feature allows that the data being accessed is authenticated. Also supports authorization by using POSIX and NFSV4 ACL. These features can be used by the solution to address these safeguards.
HIPAA 164.312 (a)(2)(iii)	Automatic logoff (Session termination mechanisms)	IBM Spectrum Scale is Software Defined storage. Users must configure the auto logoff on the operating systems on which the IBM Spectrum Scale server and client are installed. IBM Spectrum Scale GUI supports auto logoff. IBM Elastic Storage® Server (ESS) also supports auto logoff for its CLI and GUI.
HIPAA 164.312(a)(2)(iv)	Encryption and decryption (Mechanism for encryption of stored PHI) Secure Data at Rest	IBM Spectrum Scale supports file system-level encryption (encrypting data on disk) that is powered with policies that allow administrators to select what data should be encrypted and what data should not to be encrypted when it is stored on the backend disks. This feature can be used by the solution to address this safeguard.

Parts and subparts of Section 164	Section heading and subheading	IBM Spectrum Scale security features
HIPAA 164.312 (c) (i, ii)	<ul style="list-style-type: none"> ► Integrity of electronic protected health information ► Mechanism to authenticate and maintain systems integrity 	<p>IBM Spectrum Scale supports synchronous mirroring, which can be deployed to protect or recover from health data that is altered or destroyed without human intervention, such as by electronic media errors or failures.</p> <p>IBM Spectrum Scale authentication and authorization ensures the required mechanism to authenticate access to electronic protected information.</p> <p>Also, IBM Spectrum Scale supports file audit logging, which can be used to corroborate the access details to ePHI. File audit logging can be coupled with integration with IBM QRadar, which monitors the access and alteration to the ePHI data and notifies the user if the access was against the defined policy.</p>
HIPAA 164.312(e)(2)(ii), 164.312(e)(2)(ii)	<p>Transmission security (Control, integrity, encryption)</p> <p>Secure data in flight</p>	<p>IBM Spectrum Scale supports POSIX, NFS, SMB, and S3/SWIFT interfaces. All of these interfaces can be configured to perform an integrity check of data that is being transferred and encrypt that data. NFS and SMB use Kerberos, and e POSIX and S3/SWIFT can be considered with TLS for the same.</p>
HIPAA 164.316(b)(1) and (2)	Data Retention	<p>IBM Spectrum Scale supports the immutability feature that allows data to be marked immutable for a specified time. After it is marked as immutable, the data cannot be changed or deleted. Solutions can use this feature to address these safeguards.</p>
HIPAA 164.310(d)(2)(i-iv).	Secure Data Deletion: HIPAA requires physical and electronic forms of PHI to be disposed of securely to prevent impermissible disclosures of PHI.	<p>IBM Spectrum Scale supports secure deletion as a part of its file system encryption feature. Solution scan use this feature to address the stated safeguard.</p>

Table 1 on page 4 also lists how IBM Spectrum Scale can be used by solution architects to create a healthcare solution that addresses safeguards that are mandated by HIPAA and why it should be a preferred choice for file storage.

Also, a strong need exists for a Cyber Resilience solution to safeguard the ePHI and related data from cyber attacks, such as ransomware. Along with IBM Spectrum Protect, IBM Spectrum Scale provides a complete cyber resilience solution that is required for such workloads. For more information, see *Cyber Resiliency Solution for IBM Spectrum Scale*, [REDP-5559](#).

Continuous threat detection and monitoring of malicious access of ePHI data is another concern that administrators must consider. IBM Spectrum Scale integration with IBM QRadar provides enhanced cybersecurity and ensures that any access to health data that violates the compliance or business policies can trigger real-time alerts and actions. For more information, see *Enhanced Cyber Security with IBM Spectrum Scale and IBM QRadar*, [REDP-5560](#).

Regular snapshot-based backup that is supported by IBM Spectrum Scale is another feature that must be considered by the healthcare solution to ensure backup and recovery of the ePHI and related information.

IBM Spectrum Scale software-defined storage is available as on-premises deployment, as an IBM Elastic Storage Server hardware-integrated solution, and for cloud deployment. It excels in hybrid cloud deployment by using IBM Spectrum Scale features, such as Active File management (AFM). It also supports containerized workloads with Red Hat OpenShift and Kubernetes. These HIPAA safeguards and analysis is applicable for these deployment models, which gives healthcare solution various choices that are based on their requirements.

The healthcare solution that is based on IBM Spectrum Scale can also be integrated with IBM Spectrum Discover. This integration allows pluggable deep inspection to enable automated cataloging of healthcare and classifying sensitive data falling under regulatory compliance. It also supports tiering data to IBM Cloud™ Object Storage and workload scheduling software, such as IBM Spectrum Computing that is required in healthcare workloads, such as genomics.

Note: Deploying the IBM Spectrum Scale and using its security features does not mean that your solution automatically complies with HIPAA. However, it can help the solution meet that compliance.

Also, most of the safeguards that are included in the 164-Section of the Security Rule focus on security of ePHI. The stated features in this paper are generic IBM Spectrum Scale features and not specific to ePHI (although it can be applied to it). Therefore, it is advised to consider the stated mapping only as reference information. The checklist that is shown in Table 1 on page 4 is not considered a complete list and might not necessarily be comprehensive from compliance perspective.

Conclusion

This IBM Redpaper publication presented relevant portions of the HIPAA Security Rule and explained how healthcare solutions can use IBM Spectrum Scale and IBM Elastic Storage Server security features to address the security rule safeguards from a storage perspective. Known for its superlative performance, IBM Spectrum Scale can help healthcare workloads run faster and securely and help achieve HIPAA compliance.

Note: Clients are responsible for ensuring their own compliance with various laws and regulations, including HIPAA. Clients are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that might affect the clients' business and any actions the clients might need to take to comply with such laws and regulations. The products, services, and other capabilities that are described herein are not suitable for all client situations and might have restricted availability. IBM does not provide legal, accounting, or auditing advice or represent or warrant that its services or products ensure that clients are in compliance with any law or regulation.

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this paper.

IBM Redbooks

The following IBM Redbooks® publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only:

- ▶ *IBM Spectrum Scale Security*, REDP-5425:
<http://www.redbooks.ibm.com/abstracts/redp5426.html>
- ▶ *IBM Spectrum Scale Best Practices for Genomics Medicine Workloads*, REDP-5479:
<http://www.redbooks.ibm.com/abstracts/redp5479.html>
- ▶ *Cyber Resiliency Solution for IBM Spectrum Scale*, REDP-5559:
<http://www.redbooks.ibm.com/abstracts/redp5559.html>
- ▶ *Enhanced Cyber Security with IBM Spectrum Scale and IBM QRadar*, REDP-5560:
<http://www.redbooks.ibm.com/abstracts/redp5560.html>
- ▶ *IBM Spectrum Scale Functionality to Support GDPR Requirements*, REDP-5489
<http://www.redbooks.ibm.com/abstracts/redp5489.html>
- ▶ *IBM Spectrum Discover: Metadata Management for Deep Insight of Unstructured Storage*, REDP-5550:
<http://www.redbooks.ibm.com/abstracts/redp5550.html>
- ▶ *IBM Hybrid Solution for Scalable Data Solutions using IBM Spectrum Scale*, REDP-5549:
<http://www.redbooks.ibm.com/abstracts/redp5549.html>

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials, at this website:

<http://www.redbooks.ibm.com/>

Online resources

The following websites are also relevant as further information sources:

- ▶ Health Information Privacy:
<https://www.hhs.gov/hipaa/for-professionals/security/guidance/index.html>
- ▶ IBM QRadar®: Health Insurance Portability and Accountability Act (HIPAA) Reporting Extension:
<https://ibm.biz/BdqrTR>
- ▶ Health Insurance Portability and Accountability Act of 1996:
<https://aspe.hhs.gov/report/health-insurance-portability-and-accountability-act-1996>

- ▶ IBM Spectrum Scale Knowledge Center:
<https://ibm.biz/BdqrTX>
- ▶ IBM Elastic Storage Server Knowledge Center
https://www.ibm.com/support/knowledgecenter/SSYSP8/sts_welcome.htm

Help from IBM

IBM Support and downloads:

<http://www.ibm.com/support>

IBM Global Services:

<http://www.ibm.com/services>

Authors

This paper was produced by a team of specialists from around the world working with the IBM Redbooks, Tucson Center.

Sandeep R. Patil is a Senior Technical Staff Member who works as a Storage Architect with IBM System Labs. He has over 18 years of product architecture and design experience. Sandeep is an IBM Master Inventor, an IBM developerWorks® Master Author, and a member of the IBM Academy of Technology. Sandeep holds a Bachelor of Engineering (Computer Science) degree from the University of Pune, India. He is recognized and listed by Wikipedia in the World Wide Prolific Inventors list.

Sandeep Zende is a storage solutions architect in IBM Systems Group primarily for solutions that are related to healthcare industry. He has designed and delivered storage migrations, disaster recovery solutions, and software-defined solutions for several clients, including healthcare clients.

Thanks to the following people for their contributions to this project:

Larry Coyne
IBM Redbooks, Tucson Center

Dave McDonnell
John Sing
Prashant Sodhiya
Carl Zetie
IBM Systems

David Jenkins
IBM Global Markets

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
developerWorks®
IBM®

IBM Cloud™
IBM Elastic Storage®
IBM Spectrum®

QRadar®
Redbooks®

The following terms are trademarks of other companies:

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.



REDP-5591-00

ISBN 0738458600

Printed in U.S.A.

Get connected

