HIPAA Compliance for the Elastic Stack

Search Guard

Content

Notice: *This document is provided for informational purposes only. Search Guard users are responsible for making their own assessment of the information in this document.*

Introduction

The Health Insurance Portability and Accountability Act (HIPAA) of 1996 was created to modernize the US healthcare information, abandon paper records, launching healthcare into the digital era. Patient data and medical records contain very sensitive information and therefore the Department of Health and Human Services (HHS) has set up rules and regulations on how this data has to be processed.

HIPAA gives patients rights that state how health information can and can't be used. It also defines safeguards to protect the personal information of individuals. For IT providers, HIPAA puts additional restrictions, which makes it impossible to use many Open Source technologies. The same applies to Elasticsearch, as its openly published Open Source version does not implement all the safeguards required.

This whitepaper explains how Search Guard helps meeting HIPAA requirements for medical and personal data stored in Elasticsearch.

Safeguard Rules

HIPAA rules apply to all business associates and covered entities. An example of a covered entity may be: doctors, dentists, pharmacists, nursing homes, or health plans. A business associate may be any person or entity that uses protected healthcare information on behalf of a covered entity.

HIPAA contains the following mandates:

- **HIPAA Privacy Rule** - which states the limits of using patients' information. Protected Health Information (PHI) refers to any personal identifiers. When data is stored in a digital form, the PHI acronym is called ePHI. Personal identifiers may consist of names, email addresses, phone

numbers, but also geographical identifiers, fingerprints, or vehicle license plates. It can be anything that allows identifying a person.

- **HIPAA Security Rule** - which establishes standards needed to secure PHI and ePHI.
- **Breach Notification Rule** - describes what has to be done when a data breach is discovered.
- **Omnibus Rule** - additional rules published in 2013, that allow patients to have electronic access to their medical records, prohibit using PHI for marketing, and strengthen the definition of data.
- **Enforcement Rule** - when a breach occurs, the rule states how an investigation must be conducted. It also states possible fines for negligent parties. It distinguishes if the violation resulted from ignorance or willful neglect. The fine also depends on actions taken, meaning if the violation was corrected within the required time.

The second rule, "HIPAA Security Rule", is the most important for IT providers and technologies used. We will focus on it in the next sections.

Meeting HIPAA Standards

The Department of Health & Human Services has published a document "Security Standards: Technical Safeguards" (available for download), which provides useful guidance for service providers.

– Standard § 164.312(a)(1) - Access Control

"Access controls provide users with rights and/or privileges to access and perform functions using information systems, applications, programs, or files. Access controls should enable authorized users to access the minimum necessary information needed to perform job functions. Rights and/or privileges should be granted to authorized users based on a set of access rules…"

Role-Based Access Control (RBAC)

Search Guard provides role-based access control to data stored in your Elasticsearch cluster. This blocks any unauthorized access to information inside Elasticsearch. It also enables you to control exactly which users can access specific patient PHI, ePHI, and PII.. The roles can be defined on a granular level, making it possible to implement a least-privileges access strategy, so users only have access to *the minimum necessary information needed to perform job functions* Search Guard

integrates with industry-standard Identity Provider systems like Active Directory, LDAP, Kerberos, OpenID and SAML. Third-party IdPs like Okta or Auth0 can be integrated as well.

Cluster- and Index-Level Access Control

Search Guard roles grant access to an Elasticsearch cluster as well as the indices within that cluster by defining a set of access rules. These access control rules specify what users are allowed to do with any PHI and ePHI data that is made available to them. For example, a user may have read and write privileges when it comes to a patients master data, read-only access to some PII data, but no access at all to a patient's medical records.

Document- and Field-Level Access Control

Search Guard also provides access to individual documents and fields stored in Elasticsearch. Restrictions on documents and fields are based on roles as well. For example, you can set up a role with access to medical records from a specific department only. You can also specify that this role only have visibility into fields in a patient's medical record so that all that is exposed to them is simply what they require to successfully perform their jobs.

Field Anonymization

If a user does not have access to cleartext ePHI and PII data, but must still be able to analyze and gather statistical information, Search Guard can anonymize any data stored in Elasticsearch. Anonymization is performed at runtime. You can control who has access to cleartext and anonymized data based on the assigned role.

- Implementation § 164.312(a)(2)(i) - Unique User Identification

"Assign a unique name and/or number for identifying and tracking user identity."

Search Guard assigns each Elasticsearch access request to a user which can be uniquely identified. Anonymous requests are are blocked. Additionally, the action of accessing any Elasticsearch data can be recorded by using the Search Guard audit and compliance logging features.

- Implementation § 164.312(a)(2)(ii) - Emergency Access Procedure

"Establish (and implement as needed) procedures for obtaining necessary electronic protected health information during an emergency."

Search Guard can make use of a so-called *administrative TLS certificate* that allows access to all data in case of an emergency. The usage of this certificate is can be recorded by using the Search Guard audit and compliance logging features as well. This ensures all access is clearly tracked and recorded, even in case of an emergency.

- Implementation § 164.312(a)(2)(iii) - Automatic Logoff

"Implement electronic procedures that terminate an electronic session after a predetermined time of inactivity."

Search Guard uses a stateless architecture for securing Elasticsearch and does not maintain any session state. Each request to Elasticsearch must contain verified credentials. If no credentials are provided, access is denied.

For Kibana, the session lifetime is 60 minutes by default but can be changed to any value by the customer. After the session terminates the user is logged off automatically.

- Implementation § 164.312(a)(2)(iv) - Encryption and Decryption

"Implement a mechanism to encrypt and decrypt electronic protected health information."

Search Guard enforces TLS (Transport Layer Security) for all data channels. TLS is required for client-to-server communication and server-to-server communication between each of the Elasticsearch nodes in a cluster. The customer can configure which TLS versions and cipher suites are allowed. For example, to increase the level of security a customer may only allow TLS 1.3 and Elliptic Curve Cryptography (ECC).

Encryption at REST can be achieved at the operating system level with tools like dm-crypt, EcryptFS, or others.

– Standard § 164.312(b) - Audit Controls

"Implement hardware, software, and/or procedural mechanisms that record and examine activity in information systems that contain or use electronic protected health information."

The Search Guard Audit Logging feature tracks any time someone has accessed Elasticsearch down to the second. It can create audit trails for every activity inside the Elasticsearch cluster. This includes security-related events like failed log-in attempts, and also data access. Search Guard can record which users accessed PHI, ePHI, and PII down to the individual field or document. It can also be used to create audit trails capturing the complete lifecycle of PHI, ePHI and PII data. This includes when the data was created, how it changed over time, and when it was deleted.

All audit events contain information like the timestamp (date and time when a request was submitted), the username (who submitted the request), and the remote IP (from where was the request submitted). In addition, Search Guard can store which documents and fields have been accessed, and what changes have been made (if any).

This provides administrators with full oversight over all activities inside their Elasticsearch cluster. The generated audit events can be stored at different storage endpoints by using the included storage connectors.

– Standard § 164.312(c)(1) - Integrity

"Implement policies and procedures to protect electronic protected health information from improper alteration or destruction."

Access to medical data stored in Elasticsearch is governed by Search Guard roles. Search Guard roles can be freely defined and assigned, and implement fine-grained access control to data. If the company has set up access policies in a third-party system like Active Directory, LDAP, Okta, or Auth0, Search Guard can integrate with those systems and use the already implemented security policies.

- Implementation § 164.312(c)(2) - Mechanism to Authenticate Electronic Protected Health Information

"Implement electronic mechanisms to corroborate that electronic protected health information has not been altered or destroyed in an unauthorized manner."

Search Guard provides several features to shield health information from improper alteration or destruction. Using TLS for data in transit certifies that the data is encrypted and cannot be altered while traveling across the network.

With role-based access control (RBAC), you can configure which users have read-only access to data, and which users are allowed to change data. Unauthorized or anonymous access will not be possible.

Audit logging ensures that any alteration to data is recorded. This includes both the identity of who altered the data as well as the changes that took place.

Finally, if sensitive data like medical records must not be altered or deleted once entered into Elasticsearch at all, Search Guard provides the Immutable Index feature which protects any document from being changed or deleted after it has been created.

– Standard § 164.312(d) - Person or Entity Authentication

"Implement procedures to verify that a person or entity seeking access to electronic protected health information is the one claimed."

Search Guard integrates with all industry standards for authentication and authorization. This includes Active Directory, LDAP, Kerberos, JWT, OpenID, SAML, and others. These systems are used to verify the identity of each person accessing any data stored in Elasticsearch.

– Standard § 164.312(e)(1) - Transmission Security

"Implement technical security measures to guard against unauthorized access to electronic protected health information that is being transmitted over an electronic communications network."

Search Guard mandates the use of TLS for all data being transmitted. TLS makes sure the data is encrypted, cannot be sniffed, and cannot be altered while being transmitted through the network. Search Guard users can leverage their own PKI infrastructure to ensure that only company-issued TLS certificates can be used for the purpose of encryption.

– Implementation § 164.312(e)(2)(i) - Integrity Controls

"Implement security measures to ensure that electronically transmitted electronic protected health information is not improperly modified without detection until disposed of."

As with "Transmission Security" Search Guard uses TLS to shield against unauthorized data modification on the network. All data is encrypted. Suppose data was modified while traveling across the network. In that case, it becomes unusable for the receiver, so any modification is immediately detected.

- Implementation § 164.312(e)(2)(ii) - Encryption

"Implement a mechanism to encrypt electronic protected health information whenever deemed appropriate."

Search Guard uses TLS on both the Elasticsearch transport layer and also the REST API layer. This ensures that any data that is being processed on an Elasticsearch cluster is encrypted while traveling across the network. Encryption at REST can be achieved at the operating system layer with tools like dm-crypt, EcryptFS or others.

Summary

Search Guard is an enterprise grade security suite for protecting and controlling access to any PHI, ePHI and PII data stored in Elasticsearch. It's powerful features for encryption, role-based access control tfield and document level security, as well as its extensive audit logging capabilities make Search Guard an ideal solution for organizations of all sizes to meet even the most rigorous HIPAA compliance standards.

Search Guard relies on industry standards like TLS, Active Directory / LDAP, OpenId, Kerberos, SAML amongst others. By only working with well known open standards Search Guard allows for clear and painless security assessments to be possible.

In addition to its core security features, Search Guard also provides versatile anomaly detection and alerting functionalities. These critical features can automatically detect any unauthorized access attempts and can notify administrators (through Email, SMS, PagerDuty, Slack, etc) in the event that a data breach has taken place.