Compliance Edition Beta released

Jochen Kressin · 2018-03-28

We are very proud to announce the arrival of the Search Guard Compliance Edition Beta! We added some neat new features that will help you to stay compliant with regulations like GDPR, HIPAA, SOX or ISO.

The technical preview of the Compliance Edition that we released in January already had important compliance features on board:

Read History

With this feature you are able to track exactly which fields in your documents have been accessed by what user. This especially helps you to meet several aspects of GDPR: Under GDPR, you are required to tell a user or customer exactly which persons in your organization have access to her or his personal data it, and for what purpose it is being processed. With the read history you can do that: Monitor and track access to sensitive fields like first name, last name, email, etc. and store the generated events in Elasticsearch, Kafka, Cassandra or any other system of your choice.

Write History

Similarly, the write history makes it possible to track creation, changes, and deletions of sensitive documents in your cluster. Search Guard stores changes in JSON patch format, so you can see how a document has changed over time. How is that useful? If you store *personally identifiable information* (PII), the customer can

- Request information about what PII data is stored
- Request information about when the data has been created
- Request changes to the PII data
- Request the deletion of the data ("right to be forgotten")

With the write history you can keep an audit trail about all these events, and provide it to your customer.

> The read and write history alone make GDPR compliance much easier, but we do not stop there.

The Compliance Edition Beta adds many new stuff on top, giving you an arsenal of features to meet compliance regulations.

Field anonymization

With the Field-level security feature of Search Guard you can filter out sensitive fields from your documents: Configure a whitelist or blacklist of fields, and Search Guard will filter them accordingly. However, sometimes you don't want to remove fields completely, but instead anonymize them. This means replacing the cleartext value with a consistent and secure hash.

Until now the only way of doing so was to store a field twice: Once in cleartext, and once as a hash value. A standard way of replacing a value with a hash at ingest time is to use the logstash fingerprint plugin for example. With the field anonymization feature of Search Guard, this is no longer necessary: You store the value in cleartext, and let Search Guard replace it with a hash at query time. Field anonymization is configured per role: For example, you can define that an *admin* or *manager* role can see the cleartext value, but all other roles can only see the hash.

Search Guard uses Blake2bDigest to calculate the hash. This algorithm strives a right balance between speed and security and has built-in support for a salt for randomized hashing. The salt can be configured in elasticsearch.yml and has to be at least 32 characters for security reasons.

Apart from being very easy to use and configure, anonymization also affects GDPR: Anonymized data is not PII data. In other words, if a user only sees the hashed equivalent of the cleartext data, several compliance rules do not apply anymore. Consequently, anonymized fields are excluded from the Search Guard read history automatically!

Immutable Indices - Data integrity

To protect the integrity of your data, Search Guard offers two main features. By using TLS you can make sure that the data is not changed while in transit. And by applying role-based access control you can define which users can create, modify or delete data. Immutable indices take data integrity to another level.

> If an index is marked immutable, you can create documents in it, but you can never change them again.

This follows the write-once / read-many approach and is extremely helpful if you are required to make sure that data, once recorded, cannot be changed. All audit and compliance events fall into this category - once written, you don't want anyone to change them afterward. Marking an index as immutable is extremely easy - just list them in elasticsearch.yml:

```
searchguard.compliance.immutable_indices:
  - immutable1
  - immutable2
```

In addition, Search Guard also makes sure that no one changes your index directly. The following operations are forbidden for an immutable index:

- Deleting the index
- Opening and closing the index
- Performing a reindex
- Snapshot / restore

Audit event routing

Audit events can be stored to various endpoints. This includes Elasticsearch, Webhooks, Kafka, Cassandra and many more. With the regular audit log module, you could configure and use exactly one endpoint for all events. The Compliance Edition now offers flexible event routing and multiple endpoints.

For example, you can store events from the read- and write history audit trail in Kafka while still shipping all security-related events like failed logins or missing privileges to a SIEM system. The routing is based on the event category, and you can store events to multiple endpoints at the same time.

This makes the audit and compliance logging extremely flexible with regards to the storage endpoint. The new configuration is completely backward-compatible, so if you used the audit log module before, you can introduce new endpoints step-by-step.

Elasticsearch installation and Search Guard configuration integrity

Last but not least, Search Guard can now also track the integrity of your Elasticsearch installation and the Search Guard configuration.

Elasticsearch installation integrity

On node startup, Search Guard can emit an event that lists:

- Settings in elasticsearch.yml
- Environment variables used in elasticsearch.yml
- Java properties
- Files used by Search Guard like PEM certificates, keystores or Kerberos keytabs

Search Guard also calculates a hash of these settings, so you can detect any changes to your Elasticsearch installation immediately.

Search Guard configuration integrity

Likewise, the Search Guard configuration integrity tracking records a trail about read- and write-access to the Search Guard configuration. You can see who has accessed the configuration and what changes have been made. Do you want to know when a particular role was added? Or what the permissions for a role looked like at any point in time? With our new configuration tracking, you can see that. Store these events in an immutable index and you can always be sure to have the full control over your sensitive security configuration.

What to do next?

- Give our new compliance features a test drive by downloading the Beta
- Head over to the compliance documentation and learn about the new featured in-depth
- Become a Search Guard Vanguard by answering the compliance survey

Image: shutterstock / Inspring

This article was published on the Search Guard blog: https://search-guard.com/blog/elasticsearch-compliance-edition-beta/