



Safeguarding Your Enterprise: 5 Pillars of GenAI Security

White Paper

5 Pillars of Security

The world is transfixed by generative AI (GenAI). It's not just the human-like answers it gives, but it's also the amorphous movement of letters across the screen that creates an allure. Your employees have already adopted it (whether sanctioned or not) and the C-suite is eager to make something happen — before competitors do.

Yet, but the biggest hurdles remain. How do you deploy a GenAI solution without putting your company at risk?

Locking down the risks means adopting these five pillars of enhanced security:

- Secure Content Retrieval
- Results Grounded in Context
- Auditability
- Zero Retention by LLM Provider
- Compliance With Data Security Protocols

1 Secure Content Retrieval

GenAI works by passing content from your system into a large language model (LLM). However, you don't actually want to move your data, and you don't want all your content going into the LLM. And beyond that, you want to limit what people from various departments can see.

You could build a standalone solution that houses content answers. But that would require the time-consuming job of recreating roles and permissions.

Another alternative would be to create separate systems that can only be accessed by e.g., support, or HR or finance. While that requires less roles and permissions building, you now have created duplicate content from your knowledge bases — and a Gordian knot of maintenance efforts to keep permissions and content current.

A better alternative is **secure content retrieval**.

This is both critical and very hard to do. First, you don't want to move your content — you want to index a representation of it. You also want to index all the permissions and access controls around *every piece of content* that is stored in various repositories.

And if these are proprietary systems you will want end-to-end encryption from system to application. Ideally you use a middleware component that allows you to open a port only as the index needs to be updated.

Coveo Advantage

Coveo excels at secure content retrieval, providing native access to +100 cloud and on-premise content sources of both structured and unstructured data. All document-level security permissions are maintained and indexed. This means users see only content that they are authorized to.



2 Results Grounded in Context


In terms of processing cost, it would be astronomical to send your entire corpus to your LLM. Not only that, you would greatly increase the chances of hallucinations. Instead, focus on querying the index using input data, information, and authorized chunks to generate the answer. We call this **grounding context**.

In the field of artificial intelligence, grounding is a technique used to create associations between entities for the machine. These chunks are then sent to the LLM, which in turn generates the answer. This real-time grounding is essential for enterprise applications.

Coveo Advantage

When the query is made, Coveo's index pulls only the chunks of content that the person making the query is allowed to see. These personalized search results are then sent to the LLM to generate the answers. This "grounding context" helps to ensure that the output is relevant, accurate, consistent and secure.

Note that if you have two users with two different sets of permissions asking the same question — you may see different responses.



▶ 3 Auditability


Track and record all queries/prompts and responses generated by the LLM for auditing purposes. This ensures that the model is being used to generate answers in a responsible and ethical manner. And no doubt, it will be required to satisfy regulatory concerns — once the regulators finish drafting their rules.

And while you might think of this as merely an application feature, consider it from a security or risk mitigation point of view. Let's say you have a self-service portal, and a customer complains that an answer caused them to break their product. You must be able to audit LLM answers so you can take appropriate actions.



Coveo Advantage

The LLM is stateless — meaning customer data is never used by the LLM to train, retrain or improve the models. Coveo, however, will automatically log the query for you. This allows organizations to understand how the AI model is being used to generate search results and to ensure that it is being used in a responsible and ethical manner.



► 4 Zero Retention by LLM Provider

In the world of cloud computing, data ownership can get fuzzy. If your enterprise wants to maintain control over your data's privacy, there can't be any fuzziness. You need to retain complete ownership over all your data.

Make certain that your contract with your LLM vendor stipulates they can not retain nor train from either your content nor queries/prompts. Your LLM must be stateless.



Coveo Advantage

Because Coveo uses the grounding context technique, **minimal data is shared with the LLM**, and none of it is kept longer than our retention policy requires. Customer data will remain stored only in Coveo AWS data centers — as stipulated by each individual contract.

► 5 Compliance With Data Security Protocols

A locked-down environment will help ensure necessary governance against unauthorized access, use, disclosure, disruption, modification, or destruction of any enterprise data. And a secure environment must comply with standards or regulations like AICPA SOC 2 Type II, HIPAA, Cloud Security Alliance, and ISO 27001.

This means if your customer requests that you forget all about him (i.e., lose my number and other personally identifiable information or PII), and you remove that record in your CRM, in order to be in compliance, you will need to remember to remove that person's data from any application connected to your CRM.



Coveo Advantage


The Coveo platform is a multi-tenant platform that adheres to all of the above regulatory standards. Your index is secure to you and the people you authorize. So if a customer wants to be forgotten — as the CRM data is updated so too is your index. (For usage analytics, you can make a request to our dedicated API endpoint, and we will proceed with the deletion from our end.) There is no need to log in and manually remove.

Further, Coveo offers encryption in transit and at rest. Communication between client and server is encrypted and secure, preventing potential attacks such as eavesdropping, tampering or data theft. TLS endpoints use cryptographic protocols to provide authentication, confidentiality, integrity and non-repudiation services, allowing secure web communication between Coveo and Microsoft Azure's servers.

Summary

GenAI is the next make-or-break moment for technology leaders, bringing significant leaps in employee proficiency and customer satisfaction. Companies can't, and don't have to, solve for this alone.

The Coveo Relevance Cloud™ is the leading, trusted AI platform that has helped global enterprises become AI-proficient over the last 10 years, paving the way for the next evolution of AI — GenAI.



Coveo can help CIOs harness GenAI for their organizations — while mitigating the risks outlined in this guide.

Deploy Secure AI Today

The Coveo Relevance Cloud™ is the leading, trusted AI platform that has helped global enterprises become AI-proficient over the last 10 years — paving the way for the next evolution of AI — GenAI. Coveo can help CIOs harness GenAI for their organizations — while mitigating the risks outlined in this guide.



Learn more about Coveo

Coveo is the world's leading cloud-based relevance platform. The Coveo Relevance Cloud™ uses applied AI to deliver relevant experiences in all digital interactions, from search to recommendations to personalization.

Contact us



coveo.com