Commvault

Customer FAQ: AI

1 How do you use AI in your products and services?

For over a decade, we have leveraged artificial intelligence, machine learning, and intelligent automation to enhance our platform's capabilities. This AI improves data protection and recovery operations, and assists customers in managing risk, locating sensitive data, and more.

At the heart of our AI innovations is Metallic AI, our intelligent control layer that powers Commvault Cloud. From automation and machine learning to advanced AI, Metallic AI brings greater performance and efficiency to every aspect of the Commvault Cloud platform.

To learn more about Metallic AI, visit commvault.com/platform/metallic-ai

2 How do you use Generative AI?

We leverage Generative AI to streamline the configuration and use of Commvault Cloud. Specifically, it is used to power Arlie, our AI assistant. Users can interact conversationally with Arlie to learn how to use the product and troubleshoot errors, making it easy to use, even for novices. This approach not only simplifies the user experience but also optimizes feature utilization to its fullest potential. Arlie's capabilities include:

Active Insights: AI troubleshooting to discover operational issues in Commvault Cloud. Get insights, recommendations, and steps for resolving job issues.

API Code Assist: Offers a "no-code" way to build an integration or code an action. Simply type a description of what you want to do in natural language and get step-by-step instructions on the APIs to use and how to use them. Arlie maps commands to REST APIs supported by Commvault, providing a user-friendly experience for developers of all skill levels.

Custom Walkthroughs: Context-sensitive, guided product walk-throughs that make it easy to perform new tasks in Commvault Cloud.

Arlie™ Chatbot: Ask "how-to" questions and get step-by-step guidance on configuring and using Commvault Cloud to maximize cyber resilience.

3 What ethical principles guide your development and use of Generative AI?

We seek to harness the power of Generative AI through responsible AI deployment. The NIST AI Risk Management Framework (RMF) 1.0 best practices are our guiding principles. We implement continuous monitoring systems to detect and mitigate any unintended behaviors or vulnerabilities. Additionally, we prioritize transparency by providing clear documentation and explanations of our AI systems' decision-making processes. For example, in certain Arlie features, including Arlie™ Chatbot and API Code Assist, we provide users with references to the sources from which the output has been extracted. Additionally, our product documentation provides information on Arlie's architecture and the data used to generate output. Read the documentation here.

These measures help our AI deployment to be responsible, safe, and secure so that our customers can confidentially leverage the benefits of AI.

Learn more about our Responsible AI Policy here.

4 What data do you collect to train your Generative AI?

None. We utilize pre-trained LLM from OpenAI and feed relevant data into them as context to extract insights and responses. For example, Arlie utilizes pre-trained LLM along with context from Commvault product and API documentation, user query, error codes, error descriptions, and anonymized customer support resolutions. User queries guide generative AI responses, they are not used for the AI's training purposes, preserving the integrity of our data-driven insights. Human supervision oversees our generative AI output to confirm anonymization.

5 Can I opt out of or disable Generative AI features in your products and services?

Yes. Arlie is optional for installed software and can be activated or deactivated at any time according to customer preference. In Commvault Cloud SaaS, Arlie is enabled by default to enhance user experience, but it does not collect any data in the background, maintaining customer data privacy. SaaS customers may opt out of using Arlie entirely if they prefer.

6 How do you address potential risks associated with Generative AI deployment?

We address potential risks associated with AI deployment through a multifaceted approach that prioritizes reliability, fairness, and security. Firstly, to mitigate the risk of inaccurate or inappropriate responses, we meticulously train and test our AI using diverse datasets that are agnostic to any specific user and do not contain protected characteristics such as age, race, or gender. This helps minimize biases and avoid discriminatory practices.

We also maintain transparency by citing the sources of responses and recommendations provided by Arlie, alongside a clear disclaimer about the potential for inaccuracies in AI-generated responses. Human oversight is integral to our process, where all AI outputs undergo review by human experts who assess their accuracy before any actions are taken.

7 What security measures are in place to protect my data used for AI?

Our AI features are integrated within the Commvault Cloud platform which restricts access to its functionalities from outside the platform. This means Commvault Cloud's industry-leading security features are in place, including role-based access controls (RBAC), which help prevent unauthorized access and keeps systems secure and resilient.

We adhere to CIS Level 1 benchmarks to harden the infrastructure and minimize vulnerabilities. Additionally, we implement secure communications through encryption, loopback mode, and enforcement of client certification authentication.

8 Do you have a human-in-the-loop approach to producing Generative AI insights?

Yes. We firmly believe in maintaining human oversight and control, viewing AI as a supportive assistant rather than an autonomous decision-maker. All AI-generated insights are carefully reviewed by humans to inform their judgment, ensuring that decisions are based on a blend of AI-driven recommendations and human expertise. This approach ensures responsible decision-making and mitigates risks associated with fully automated AI outputs.

9 Can you explain how your Generative AI arrived at a specific response or recommendation?

Yes. We feed controlled data into the generative AI and maintain logs of the data sources from which responses are derived. Such responses are accompanied by links to relevant data sources: Commvault product and API documentation, or anonymized customer support resolutions. Arlie's responses are generated based on these data sources, then tailored to address user queries. It's important to note that while user queries guide Arlie's responses, they are not used for the AI's training purposes, preserving the integrity of our data-driven insights.

10 How do you monitor and audit the performance of your Generative AI?

We ensure thorough monitoring and auditing of our AI through several proactive measures. We start by clearly delineating AI entry points, ensuring users are well-informed when AI intervenes. Additionally, we establish a robust feedback loop where users can easily provide input on AI-generated responses, recommendations, and insights. This feedback mechanism includes a simple like/dislike feature within our product, accompanied by an option for users to provide more detailed comments. This direct user feedback is invaluable for continuously enhancing the performance and accuracy of our AI over time.

11 What future plans do you have for AI development and implementation?

We will continue to focus on using AI to enhance data protection and recovery operations efficiencies in the most intelligent, secure, and automated way. Stay tuned for more updates by visiting commvault.com/platform/whats-new.

12 Who should I contact with further questions about your AI practices?

If you have questions, comments, or feedback about Commvault AI practices, please contact compliance@commvault.com. If you have requests, queries, or complaints about your personal data, please reach out to privacy@commvault.com.