Google Cloud

Google Cloud Whitepaper
September 2024

Generative AI, Privacy, and Google Cloud

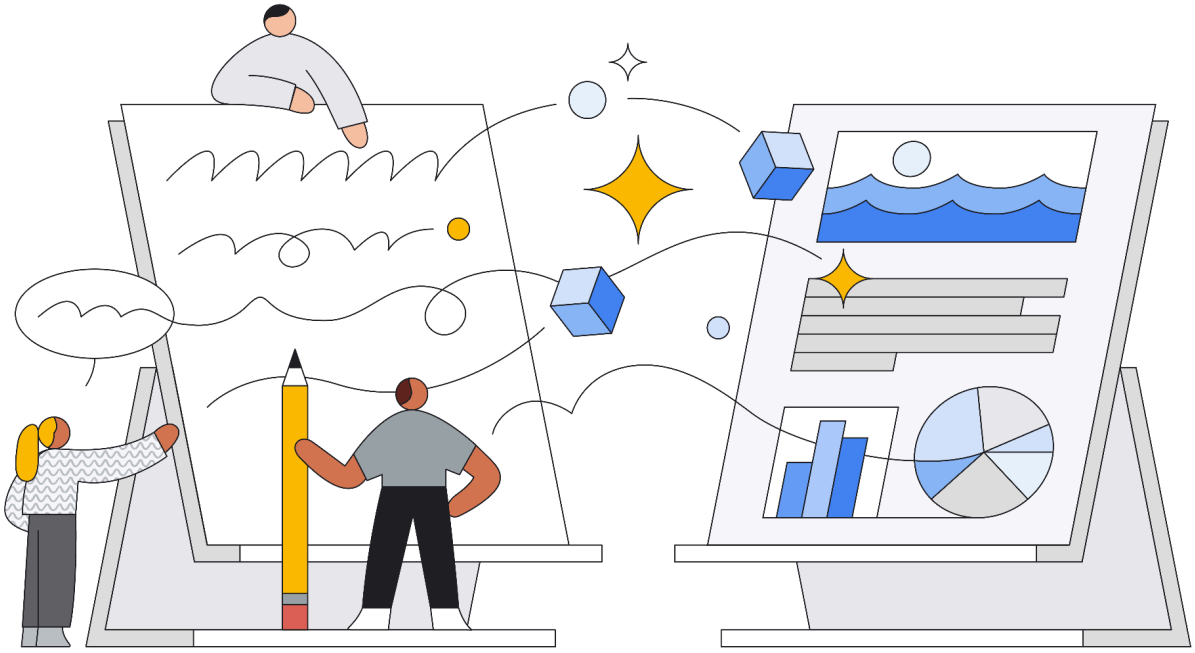

Disclaimer

This whitepaper applies to Google Cloud products described at cloud.google.com. The content contained herein is correct as of September 2024 and represents the status quo as of the time it was written. Google's security policies and systems may change going forward, as we continually improve protection for our customers.

Introduction

With our Cloud AI offerings, we are committed to preserving our customers' privacy and supporting their compliance journey. Google Cloud has a long-standing commitment to GDPR compliance, and AI is no different in how we incorporate privacy-by-design and default from the beginning. We engage regularly with customers, regulators, policymakers, and other stakeholders as we evolve our offering to get their feedback for Google Cloud (formerly known as Google Cloud Platform), Google Workspace, or Google Workspace for Education (together, the "Cloud Services") AI offerings which process personal data.

As part of our privacy reviews, we have a special focus on the intersection of Cloud AI and privacy when developing our offerings. When we bring them to market, Google Cloud AI's approach includes incorporating privacy design principles, designing architectures with privacy safeguards, and providing appropriate transparency and control over the use of data.

Addressing global privacy and data protection requirements enables you to apply AI to accelerate your business with confidence. At Google Cloud, we provide you with comprehensive control and visibility over your data. Transparency creates trust, and trust in AI is necessary for any business to succeed in this arena.

AI Privacy commitments for Google Cloud

Across Google Cloud, we've long shared robust privacy commitments that outline how we protect user data and prioritize privacy. Generative AI doesn't change these commitments — it reaffirms their importance. To support our customers interested in further understanding our privacy perspective, we've outlined our core AI privacy commitments for Google Cloud and Google Workspace.

Essential Commitments

- **Your data is your data.** The data or content generated by a Generative AI Service prompted by Customer Data ("Generated Output") is considered Customer Data[1] that Google only processes according to customer's instructions.[2] We continue to maintain that customers control their data and we process it according to the agreement(s) we have with each customer.
- **Your data does not train our models.** We recognize that customers want their data to be private and not be shared with the broader Google or Large Language Model training corpus. We do not use data that you provide us to train our own models without your permission.
- **We provide enterprise-grade privacy and security.** We provide Cloud AI offerings such as Vertex AI and foundational models with enterprise-grade safety, security, and privacy baked in from the beginning.

Further Commitments

- **Your fine-tuning data is your data.** In our GenAI implementation for enterprise customers, the data of the organization remains in their own instance, whereas our LLM is frozen. The learning and fine-tuning of the model with customer's data is stored in the adaptive layer in the customer's instance.
- **Access transparency.** We will not and cannot look at customer data without a legitimate need to support your use of the service. That's why we led the way in providing meaningful transparency into provider access to customer data and now we're extending that transparency to our AI and ML work.
- **Customers have control over where their data is stored.** Customers maintain control over where their data is stored and how or if it is used, letting them safely pursue data-rich use cases while complying with various regulations. Google does not persistently store, read, or use customer data outside your cloud tenant.

[1] See the Generative AI Service terms as part of the Google Cloud Service Specific Terms
[2] See Cloud Data Processing Addendum

- **Working together to build solutions.** If a customer would like to work together to develop a solution using our AI/ML products, by default our teams will work only with data that the customer has provided and that has identifying information removed. We work with customer raw data only with consent and where the model development process requires it.

AI privacy commitments for all Google Workspace users

GenAI does not change our foundational privacy protections for giving users choice and control over their data. To that end, here are our key commitments for all global Google Workspace users.

- **Your data is your data.** The content that you put into Google Workspace services (e.g., emails, documents) is yours. We never sell your data, and you can delete your content or export it.
- **Your data stays in Google Workspace.** We do not use your Google Workspace data to train or improve the underlying GenAI and large language models that power Gemini, Search, and other systems outside of Google Workspace without permission.
- **Your content is not used for ads targeting.** Google does not collect, scan, or use your content in Google Workspace services for advertising purposes.
- **Your privacy is protected.** Interactions with intelligent Google Workspace features, such as accepting or rejecting spelling suggestions, or reporting spam, are anonymized and/or aggregated and may be used to improve or develop helpful Google Workspace features like spam protection, spell check, and autocomplete. This extends to new features we are currently developing like improved prompt suggestions that help Workspace users get the best results from Gemini for Google Workspace features. These features are developed with strict privacy protections that keep users in control.

  See below for more detail on additional privacy, security, and compliance commitments we make for business customers.

Commitments for business, education, and public sector customers

When Google Workspace commercial customers adopt Gemini for Google Workspace they get the same robust data protection and security standards that come with all Google Workspace services, with specific protections for businesses, education, and public sector customers.

- **Your interactions with Gemini for Google Workspace stay within your organization.** Gemini for Google Workspace stores any prompts or generated content alongside your Workspace content and does not share them outside your organization.
- **Your existing Google Workspace protections are automatically applied.** Gemini for Google Workspace brings the same enterprise-grade security as the rest of Google Workspace, automatically applying your organization's existing controls and data handling practices, such as data-regions policies and Data Loss Prevention.
- **Your content is not shared with or used by any other customers**. None of your content is used for model training outside of your domain without permission.

Workspace's security and privacy commitments for customers can be found on our Workspace Security and Privacy page. Learn more about Gemini for Google Workspace by visiting our website.

AI privacy and data protection frequently asked questions (FAQs)

What is Google Cloud's approach to privacy and data protection when developing generative AI services?

Google Cloud is committed to preserving customers' privacy with AI offerings and to supporting their privacy compliance journey. We incorporate important privacy principles such as privacy-by-design and by-default from the beginning of the development of generative AI services. We engage regularly with customers, regulators, policymakers, and other stakeholders as we evolve our offering to build trust, promote information sharing, contribute to regulation development, and demonstrate compliance.

Does Google Cloud use customer data to train models?

Google Cloud does not use customer data for training models without customer's prior permission or instruction. This commitment is outlined in the 'Training Restriction' sections of both the Service Specific Terms for Google Cloud Platform and the Google Workspace Service Specific Terms.

Where is customer data processed when using Google Cloud generative AI services?

We've seen tremendous adoption of our generative AI capabilities so far. But, we know that around the world, some industries, customers, and use cases have greater requirements around data residency in their jurisdictions before they can leverage these capabilities to their fullest extent. Google Cloud services with data residency lists the services, including Vertex AI Platform and other AI services that can be configured for data location. Customers can control what regions customer data is stored at-rest when using Generative AI and Vertex AI Search & Conversation by using the corresponding regional or multi-regional APIs. To learn more about Google Cloud's regions and multi-regions, please visit our global cloud locations. For more information on data regions for Gemini for Google Workspace please visit our Data Regions page.

What is Google Cloud's approach to ensuring a high level of data quality in the AI model lifecycle?

We recognize that customers are interested in learning how Google Cloud has trained our foundational models and how we ensure the highest levels of quality. Our AI models are trained on large datasets, which allows it to understand and generate human-quality output in response to a wide range of prompts and questions.

Before sharing the model with users, we run dedicated rounds of adversarial testing to find flaws in the model. We enlisted product experts to intentionally stress test the system with an adversarial mindset. Additionally, our policies prohibit users from knowingly generating content

that is sexually explicit; hateful or offensive; violent, dangerous, or illegal; or divulges personal information. We have systems that help us detect and filter such content.

How does Google Cloud help customers ensure they have a high level of data quality?

Customers have control and ownership of their datasets used to tune Google Cloud's foundational models. We outline our Intellectual Property Terms for AI/ML Services in detail via our Service Specific Terms. We also provide resources to help customers begin training and tuning foundational models using their data, including supervised tuning and reinforcement learning from human feedback (RLHF) tuning, which uses preferences specified by humans.

How long is data retained and for what purpose?

Google maintains policies and procedures on data classification, protection, and handling throughout its lifecycle according to legal and regulatory requirements. The maximum lifespan for a model trained using Vertex AI depends on the type of model.

Will Google Cloud use any subcontractors or third parties to process personal information?

Google companies directly conduct the majority of data processing activities required to provide Google Cloud services. However, we do engage with some carefully selected third party vendors to perform limited activities in connection with Google Cloud services.

We recognize the importance of transparency about the third parties we engage with who may process your data. We share information about our vendors on our Google Cloud Platform Subprocessor page to provide our customers with visibility. This includes who they are, where they are located, and the specific services they support. We provide this information for our AI / ML services.

Google expects our Subprocessors to meet the same high standards that we do. Before onboarding Subprocessors, Google assesses their security and privacy practices. We do this to ensure that Subprocessors provide a level of security and privacy appropriate to their data access and the scope of the activities they are engaged to perform.

Once Google has assessed the risks, the Subprocessor is required to enter into appropriate security, confidentiality, and privacy contract terms. In particular, Google requires our Subprocessors to only access and use your data to the extent required to perform the obligations subcontracted to them and to do so in accordance with our contract with you. Google will remain fully liable for all obligations subcontracted to our Subprocessors.

To enable customers to retain oversight of our Subprocessors, we will notify customers when we engage a new Subprocessor so that you know in advance before any new Subprocessor starts processing your data.

How does Google Cloud address data minimization when building AI models or products?

Google is committed to complying with applicable privacy and data protection laws. We have a long-standing commitment to global privacy compliance, and generative AI services are no different. These commitments are backed by the strong contractual privacy commitments we make available to our customers in the Cloud Data Processing Addendum for Google Cloud and Google Workspace. We also closely track and monitor industry standards such as the NIST AI Risk Management Framework and the ISO/IEC 42001 AI Management System Standard, and assess our services against them, to ensure we continue to develop and deliver services that serve our customers' needs.

How does Google Cloud ensure security of generative AI services?

Google Cloud's generative AI services benefit from our globally distributed and redundant infrastructure and inherit the platform's foundational controls. Layered security controls protect AI models running on Google Cloud as we do not rely on any single technology to secure our infrastructure. Instead, our technology stack builds security through progressive layers that deliver defense in depth.

Our AI products are built on top of a scalable technical infrastructure designed for maximized availability and reliability while providing security through the entire information processing lifecycle. Google Cloud's core principles include defense in depth, at scale, and by default. Data and systems are protected through multiple layered defenses using policies and controls configured across Identity and access management (IAM), encryption, networking, detection, logging, and monitoring. The platform is further underpinned by a secure-by-design foundation supported by operational controls consisting of in-depth security reviews, vulnerability scanning, ongoing threat monitoring, and intrusion detection mechanisms that enable secure service deployment and safeguard customer data. The security controls specific to GenAI on Vertex AI can be found here.

Does Google Cloud support compliance with applicable regulatory requirements in all relevant jurisdictions?

Our teams closely monitor and analyze new and updated regulations, and we regularly engage regulators through roundtables and other forums. We are committed to preserving our customers' privacy with our Cloud AI offerings and supporting their compliance journey.

We have a long history of doing this with our GDPR privacy efforts, and AI is no different. We are engaging with customers and regulators as we evolve our offering to get their feedback. We are committed to GDPR compliance in our Cloud AI offerings and to supporting our customers' compliance journey. We also closely track and monitor industry standards such as the NIST AI Risk Management Framework and the ISO/IEC 42001 AI Management System Standard, to ensure we continue to develop and deliver tools that serve our customers' needs.

Further resources and information

These are the privacy commitments that Google Cloud and Workspace extends to all of its users. But they aren't just words. To ensure we continually meet these high standards, independent auditors validate our practices against international standards and best practices. Our products have achieved a number of security and privacy certifications from independent auditors who assessed our services' compliance practices in those domains. We apply those practices also to our Generative AI offerings.

Topics related to trust and transparency can often be interconnected. The following resources may be of interest to customers looking to further understand our trust positions.

- **Google's Privacy Policy Recommendations.** In our Generative AI and Privacy Policy Recommendations Working Paper, we offer insights into how organizations and policymakers can apply long-standing privacy principles to protect personal data.
- **Google Cloud's** Trust in Artificial Intelligence white paper

For more information about how enterprise customers can benefit from comprehensive privacy offerings of Google Cloud, please visit cloud.google.com/privacy.