

Supporting Your Data Protection Impact Assessment (DPIA) for Google Workspace with Gemini

August 2025

Google Workspace with Gemini

Table of contents

01 Executive summary

02 Google's Foundational Commitment to Security and Privacy

03 Protecting Your Customer Data with Robust Controls

04 Flexible Gemini Deployment with Granular Controls

05 Mitigating AI-Specific Risks: A Layered Defense Strategy

06 Monitoring Gemini Usage

07 Other DPIA Considerations



Executive summary

This support guide is designed to assist Google Workspace customers, as controllers, in preparing their Data Protection Impact Assessments ("**DPIAs**") concerning the use of Gemini in Workspace apps and the Gemini app (referred together as "**Gemini**" for simplicity) within their organizations. It supplements, and can be used together with, our DPIA Resource Center.

Similar to our Resource Center, in this guide we use capitalized terms (such as "Customer Data" and "Customer Personal Data") which have the meaning given to them in your contract for Google Workspace. Additionally, the terms "personal data", "data subject", "processing", "controller" and "processor" - as used in this document - have the meanings given in the GDPR (although these terms may have similar definitions under other privacy laws).

As the data controller, your organization is responsible for determining whether a DPIA of Gemini is necessary and for preparing the DPIA. Google cannot conduct DPIAs on behalf of our customers, but we provide assistance by making available the information and other resources set out in this document that can help you complete the DPIA for your use of Gemini.

Note

This support guide applies to Google Workspace products described at workspace.google.com. The content contained therein is current as of August 2025 and represents the status quo as of the time it was written. References to forthcoming features are annotated as such and do not constitute a commitment to a specific release schedule. Google's security policies and systems may change going forward, as we continually improve protection for our customers. The availability of the product features and capabilities described in this paper are subject to license availability of various Google Workspace editions and product offerings.

Please note that this document does not constitute legal advice, and is not a replacement for reviewing relevant regulatory guidance on DPIAs or seeking independent legal advice.



Google's Foundational Commitment to Security and Privacy

At Google, security is paramount, and we have built our cloud services, including Google Workspace with Gemini, to offer a higher level of protection than many traditional on-premises solutions. This commitment is deeply embedded in Google's organizational structure, training, hiring processes, data center operations, technology, and disaster planning.

Key aspects of Google's foundational security include:

Inherited Security Settings:

By default, Gemini integrates with **your existing Google Workspace security settings and data handling practices**. This includes Drive trust rules, IRM controls, and client-side encrypted data. This ensures your Customer Data (including prompts and Generated Output) are appropriately protected and remain within your tenant.

Layered Defense Strategy:

Google implements a **layered defense strategy** against emerging attack vectors, including indirect prompt injections, to help protect Gemini and Customer Data. This is in addition to built-in defenses in Gmail that automatically block nearly 15 billion unwanted emails every day, preventing more than 99.9% of spam, phishing attempts, and malware.

Industry Standards and Regulatory Frameworks:

Google's commitment to security is underscored by our attainment of numerous internationally recognized certifications and standards for Gemini, including the following:

- SOC 1, SOC 2 and SOC 3
- ISO 9001, ISO/IEC 27001, 27701, 27017, 27018
- BSI C5
- ISO/IEC 42001: The world's first international standard for Artificial Intelligence Management Systems (AIMS), certifying responsible development, deployment, and maintenance of Gemini with appropriate ethical considerations, data governance, and transparency.
- Gemini also supports compliance with various regulatory frameworks such as HIPAA and FedRAMP High.

Enterprise-Grade Commitments:

Google adheres to enterprise privacy commitments for all customers, such as not using Customer Data in Google Workspace services, including Gemini, for advertising purposes.





Protecting Your Customer Data with Robust Controls

Protecting your Customer Data is a top priority, and Gemini in Workspace apps and the Gemini app integrate with existing Workspace data protection and security measures to help prevent unauthorized access and exfiltration.

These measures include:

Trust Rules for Drive Sharing:

Administrators can **create and manage trust rules for Drive sharing** to restrict Gemini's access by controlling data sharing between internal and external users. Since Gemini can only retrieve Customer Data that the End User has access to, these rules effectively limit what Gemini can retrieve to generate a response.

Drive Inventory Reporting:

Administrators can use Drive inventory reporting to gain a holistic view of data classification, access, and usage within your organization, enhancing transparency over Gemini’s data access

Information Rights Management (IRM) Controls:

Data Loss Prevention (DLP) policies can apply IRM controls to sensitive files. When IRM is applied (e.g., preventing download, printing, or copying), Gemini will **not retrieve those protected files** to generate a response, thus preventing access to sensitive content.

Client-Side Encryption (CSE):

For the highest level of protection, CSE can be enabled. When CSE is on, the protected data is **indecipherable to any unauthorized third-party, including Google or Gemini.**

Limited Data Access and Use:

- Gemini only accesses Customer Data that the **End User has permission to access**.
- Inputs and Generated Outputs are processed only during the active End User session and **are not stored or retained by Google** for Gemini in Workspace apps. For the Gemini app, conversation history retention can be managed by administrators.
- Google does **not use Customer Data**, including prompts and Generated Outputs, **to train or fine-tune any of its generative AI models without your permission or instruction**. We contractually back up this commitment in the "Training Restriction" section of the Workspace Service Specific Terms. To learn more about how Gemini processes prompts, see The life of a prompt: Demystifying Gemini.

Data Residency Options:

For Gemini in Workspace apps, administrators can choose to configure the storage and processing of prompts and responses (i.e. Generated Output) within the United States, Europe, or both, enhancing customer control over data location. Once Generated Output is inserted into a Google Workspace service covered by Data Regions (for example, Google Docs), that content will then be processed and stored at rest in the selected region for that service.

Access Approvals and Management:

Customers can leverage Access Approvals to ensure that Google support staff must obtain explicit consent before accessing covered data, providing an additional layer of control. Additionally, customers can use Access Management to limit the Google staff who can take support actions related to covered data.



Flexible Gemini Deployment with Granular Controls

Google Workspace offers granular controls to help administrators manage End User access and securely deploy AI capabilities across your organization:

Feature Management across Workspace services

Administrators with a supported edition can **enable or disable Gemini features across various Workspace services** like Gmail, Google Drive, and Google Docs, allowing for precise access control based on users, groups, or organizational units (OUs).

Standalone Gemini App Access

Administrators can **grant or restrict End User access to the standalone Gemini app** based on specific users, groups, or OUs.

Configure Gemini App Conversation History

Administrators can configure whether Gemini app should retain conversation history, and for long that that conversation history should be retained.

Workspace Apps & Services Control

Administrators can control whether End Users can utilize Workspace services within the Gemini app, which enables the Gemini app to **leverage content in Workspace** services the End User has access to.

Context-Aware Access (CAA)

Organizations can implement **device-specific policies** by assigning CAA policies to limit access to the Gemini app on covered devices. This control can be based on factors such as End User identity, device security status, IP address, and geographical location.

Endpoint DLP Policies (Chrome Enterprise Premium)

These policies extend robust data loss prevention to the Chrome browser, including the Gemini app. This allows for **granular control over actions such as copy and paste, printing, and uploads/downloads of sensitive content**, along with masking of personal information and screenshot protection. All relevant activities are logged for administrator's review.

Age-Based Access Control

Google Workspace allows for the control of access to its services by age. For Google Workspace for Education, for example, End Users must be at least 13 years old (or the applicable age in their country) to access Gemini if additional data protection is enabled. Specific features like image generation may, at the customer's choice, be restricted for End Users under 18.



Mitigating AI-Specific Risks: A Layered Defense Strategy

In addition to adhering to our AI Principles in the development of our generative AI models, we employ a comprehensive, multi-faceted approach to mitigate risks inherent in generative AI, such as prompt injection attacks, hallucinations, and potential biases:

↪ Prompt Injection Mitigation:

Google has developed a **layered defense strategy** to help protect against prompt injection attacks, an emerging attack vector against AI systems:

- **Prompt Injection Content Classifiers** detect suspicious inputs that may contain malicious instructions, analyzing their structure, keywords, and patterns before they can impact the AI model.
- **Security Thought Reinforcement:** This technique adds targeted security instructions around the prompt content to guide the large language model (LLM) to focus on the End User's intended task and **ignore any adversarial or malicious instructions** present in the content.
- **Markdown Sanitization:** Gemini's markdown sanitizer identifies and **does not render external image URLs**.
- **Suspicious URL Redaction:** Additionally, Google Safe Browsing detects and **redacts suspicious URLs** in responses, helping prevent URL-based attacks and data exfiltration.

- **User Confirmation Framework (Human-In-The-Loop):** Gemini incorporates a contextual user confirmation system that **requires explicit End User confirmation for certain operations** (e.g., deleting a calendar event), preventing undetected or immediate execution of sensitive actions.
- **End User Security Mitigation Notifications:** When security issues are mitigated by Google's built-in defenses, **End Users are provided with contextual information and "Learn more" links** to understand the prevented attack.

Rigorous Testing and Collaboration:

As new security threats emerge, Google is committed to deploy improved security mitigations through e.g.:

- **Manual and Automated Red Teams:** Continuous, rigorous testing by internal teams.
- **Generative AI Security BugSWAT Events:** Dedicated events to find and fix flaws.
- **Secure AI Framework (SAIF):** Adherence to strong security standards.
- **Partnerships with External Researchers:** Collaboration via the Google AI Vulnerability Reward Program (VRP) and industry peers through the Coalition for Secure AI (CoSAI) to identify and responsibly disclose AI security vulnerabilities and share threat intelligence.

Hallucination and Accuracy Mitigation:

- **Responsible Model Development:** Google's Gemini models are built responsibly, incorporating rigorous safeguards, testing, and human oversight throughout the entire development lifecycle. This approach is designed to mitigate risks, avoid unfair bias, promote privacy and security, and ensure the models generate high-quality outputs.
- **Adversarial Testing:** Before integrating models into Google Workspace services, Google conducts dedicated adversarial testing rounds to uncover flaws.
- **End User Warning:** We strongly encourage End Users to double-check Generated Output. Gemini includes a prominent notice that: "**Gemini can make mistakes, including about people, so double-check it**" to encourage End User to check the accuracy, reliability and relevance of Generated Outputs. Gemini in Workspace apps includes a similar, prominent notice that "**Gemini in Workspace can make mistakes, so double-check responses**".
- **Generated Output Validation:** Generated Outputs include sources from the End User's Google Workspace content, enabling them to verify the accuracy, reliability and relevance of these outputs.
- **"Double Check" and "Retry with Google Search":** To further mitigate against the risk of hallucinations, End Users can choose to use the "Double Check" feature or the "Retry with Google Search" option, which leverages Google Search to ground responses in authoritative information and relevant sources, improving accuracy. The data transmitted to Google Search is processed in compliance with the Cloud Data Processing Addendum.

- **Voluntary End User Feedback:** End Users can, if they wish, provide feedback about Generated Outputs to Google, specifically reporting inaccuracies or unsafe content. Google publicly discloses that it may process End User feedback as a controller in accordance with the Google Cloud Privacy Notice. Google may use feedback to: track overall user satisfaction and identify trends to inform future improvements and prioritization; identify and resolve issues quickly; and improve Gemini safety features that are designed to help prevent the generation of harmful or offensive content. However, it is not used to train any of the generative AI models that support Gemini. Feedback may be retained for up to 18 months. Note that **any feedback provided is entirely voluntary and End Users can use Gemini without providing any**.

If End Users choose to provide feedback while using Gemini, they can provide the following data to Google:

- The End User's rating for Gemini suggestions (thumbs up / thumbs down).
- The issue found by the End User, selected from a set of pre-defined options such as "offensive / unsafe" or "not factually correct" *(Gemini in Workspace only)*.
- Any additional feedback the End User may want to provide via the open text box *(Gemini in Workspace only)*.
- Generated Output, only if the checkbox to share that output is explicitly selected by the End User *(Gemini in Workspace only)*.

Bias Mitigation:

Google actively works on detecting and mitigating biases in the Gemini AI models and uses diverse training data to reduce bias.



Monitoring Gemini Usage

Google Workspace provides administrators with comprehensive tools to monitor Gemini usage and activity in your organization:

Usage and Data Access Reports

Administrators can **review overall Gemini usage and data access** across your organization, including adoption levels in the Gemini app and Gemini in Workspace, and on a per-app basis.

Audit Logs for User Activity, Interactions, and File Access:

Administrators can leverage the Reports API to track End User activity and interactions with Gemini. They can also use the Security Investigation Tool with Drive log events to identify instances when Gemini accessed a Drive file in response to an End User prompt. For eligible customers, these logs can also be exported to BigQuery for further analysis of End User activity when interacting with Gemini.

eDiscovery with Google Vault:

For eDiscovery purposes, administrators can leverage **Google Vault to search and export relevant prompts and responses** from the Gemini app. Gemini in Workspace apps do not retain prompts and responses beyond the End User's session.

Data Export Tool

You retain full control over your Customer Data. Administrators can **export covered data in the Gemini app** for archiving or backup using the Data Export tool. While prompts and responses for Gemini in Workspace apps are not stored or retained – meaning they cannot be exported – covered content stored in other Workspace services that may be accessed by Gemini can be exported using the Data Export tool, as well as Generated Output inserted into such a service (e.g. when an End User inserts a Gemini response into Google Docs).



Other DPIA Considerations

As mentioned above, this page has been designed to supplement and be used together with our DPIA Resource Center. With that in mind, when preparing your DPIA, you should assess whether the following considerations are appropriately taken into account in relation to Gemini specifically:

Data Processing Description:

Clearly identify the types of personal data intended to be processed through inputs and Generated Outputs, the categories of data subjects involved (e.g., employees, customers), and the specific purposes and benefits of using Gemini within your organization.

Lawful Basis for Processing:

Establish and document the lawful basis for processing personal data, which may vary for different categories of data subjects or processing activities (e.g., legitimate interests, consent, performance of a contract, public interest).

Necessity and Proportionality:

Assess whether the processing of personal data is necessary to achieve the intended purposes for using Gemini and if less data, and/or anonymized/pseudonymized data, could achieve the same goals. Consider leveraging Workspace features like DLP to limit personal data processing in Gemini.

Data Quality and Minimization:

Implement measures to ensure the accuracy of inputs (such as content in Google Docs, Sheets, and Slides that may be accessed by Gemini) and raise End User awareness of potential AI "hallucinations." Consider establishing guidelines for End Users on prompt optimization, good practice for reviewing and validating Generated Outputs, and on reporting inaccuracies (within your organization and/or to Google as feedback).

Fairness and Transparency:

To address potential over-reliance on Generated Outputs, inform End Users that they are interacting with an AI system when they use Gemini. Review whether your organization's privacy notices to data subjects are appropriately transparent based on the intended data processing via Gemini.

Retention:

- **Gemini in Workspace apps** does not retain End User prompts or responses after the session, but Generated Outputs inserted into other Workspace services are stored like any other Customer Data (and, therefore, can be deleted by the customer in a manner consistent with the functionality of those services).
- For the **Gemini app**, administrators have full control over conversation history retention periods. When the Gemini app conversation history is off, new chats are saved in End User accounts for up to 72 hours, so Google can provide the service and process any user feedback. This activity doesn't appear in the Gemini app Activity. When conversation history is on, new chats are saved in End User accounts until they are auto-deleted according to the auto delete control. Workspace Administrators can choose to automatically delete conversations after 3, 18, or 36 months. The default is 18 months.

Data Processing Commitments:

Gemini in Workspace apps and the Gemini app are Workspace Core Services. This means they are governed by robust contract commitments, including the Cloud Data Processing Addendum, which applies to prompts, inputs, and Generated Output in Gemini (which we treat as Customer Data). As well as committing to act as a processor in respect of any Customer Personal Data, Google commits to (among other things) implementing and maintaining technical, organizational and physical measures to protect Customer Data against accidental or unlawful destruction, loss, alteration, unauthorized disclosure or access.

International Transfer Safeguards:

Gemini adheres to the same international transfer safeguards as the rest of Google Workspace (as described in more detail in the DPIA Resource Center).

Data Subject Rights Support:

Utilize Google Workspace Admin Console features and Google's Data Subject Requests (DSR) Guide to assist in responding to data subject requests under applicable data protection law.

Organizational Safeguards:

In addition to the safeguards provided by Google, consider implementing appropriate customer-side technical and organizational measures to mitigate identified risks, such as:

- **Clear Data Governance Policies:** Define data classification, prompt/Generated Output management protocols, and retention/deletion policies tailored for Gemini.
- **Training and Guidance:** Provide End Users with training on ethical AI use, data entry best practices, and the importance of critically reviewing Generated Outputs for accuracy, reliability, and relevance.
- **Human Oversight:** Establish workflows that require human review and approval for decisions involving the use of Gemini (including using Generated Output), especially in sensitive applications.
- **Bias Detection and Mitigation:** Implement techniques like Generated Output analysis, adversarial testing, and End User feedback loops to detect and address biases (in addition to Google's efforts).
- **Incident Response Plan:** Develop and implement a specific incident response plan for issues related to Gemini, including data breaches, biased Generated Outputs, and misuse of those outputs by End Users.

By leveraging the comprehensive security and privacy controls embedded in Workspace with Gemini, combined with diligent implementation of your own organizational policies and human oversight, you can confidently deploy Workspace with Gemini while maintaining a strong security posture and ensuring data protection across your organization.