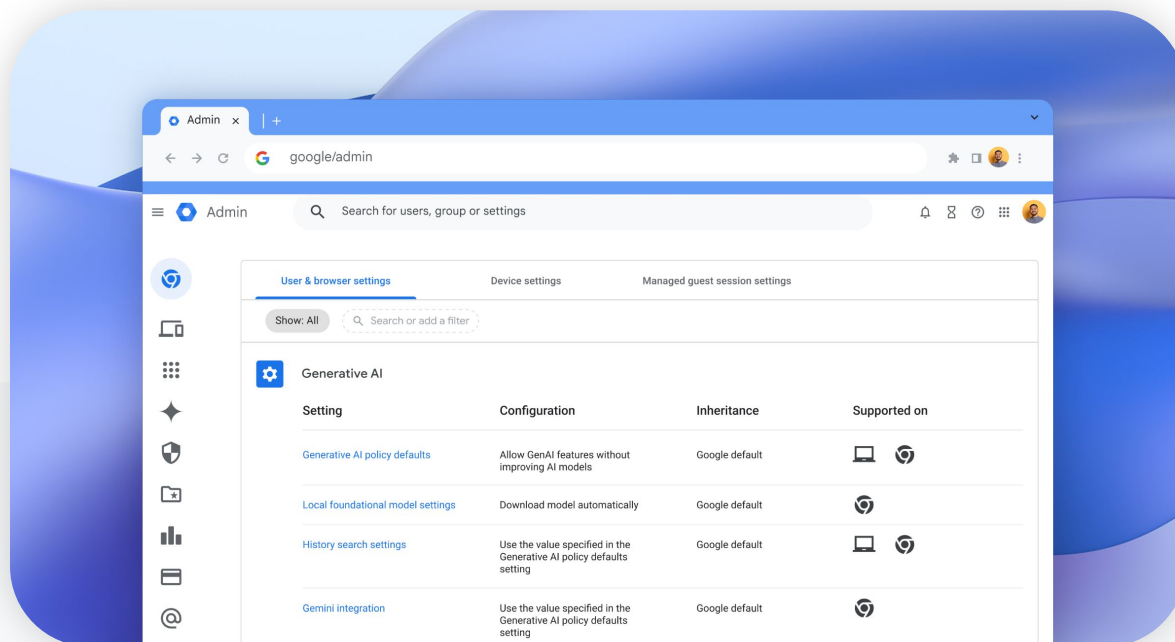




Managing ✨ generative AI in the browser

with Chrome Enterprise



The rapid rise of generative AI (GenAI) is poised to revolutionize the enterprise landscape, offering unprecedented opportunities for innovation and efficiency. However, this transformative technology also introduces new and complex security challenges that demand careful consideration. This paper delves into the implications of GenAI for enterprise security and how Chrome Enterprise gives IT and security teams a variety of ways to manage this powerful new technology based on the needs of their own organization.

Chrome Enterprise's approach to AI

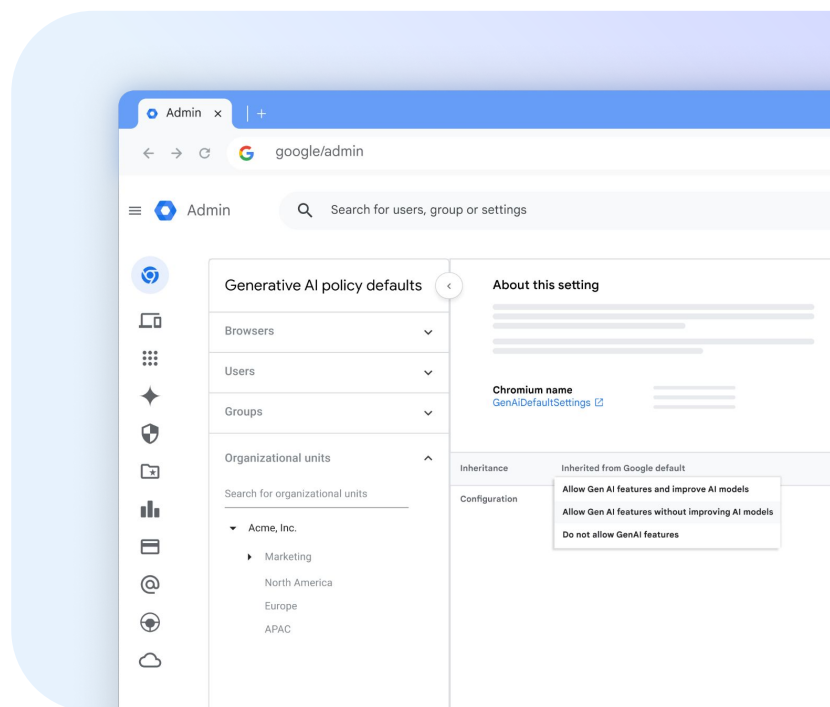
The web browser is now at the center of modern work and workflows, making it the place where GenAI can have the most impact – particularly when it is embedded into the browser itself. Chrome Enterprise enables IT and security teams to deploy browser-based GenAI capabilities to give their employees the time-saving advantages they're looking for while still allowing them to secure and manage how their workforce utilizes these third-party tools and services.

Despite recent developments and an acceleration in adoption, AI in the browser isn't new. Chrome Enterprise has been infusing AI-powered tools into the browser for more than a decade and continues to develop new ways to empower employees by augmenting, automating, and transforming workflows. The latest GenAI capabilities built into Chrome help both end users and IT admins complete their work more efficiently. For example, now that [Gemini Nano](#) (an on-device model) is packaged with Chrome, developers can produce high-quality AI experiences without having to build them from scratch.

Privacy and security are important considerations, and Chrome's capabilities are developed with [Google's AI principles](#) and foundational privacy commitments as guidelines. [Protected customer data](#) is never shared with third parties or used to train AI models. Policies can be applied to further secure browser-based capabilities, and Chrome strives to minimize the data we collect by utilizing our on-device models when possible.

Chrome Enterprise enables IT and security teams to govern how employees access and use any GenAI tool or service within the browser, via the configuration and application of granular policies and controls. These policies are centrally managed in the Google Admin console and can be easily applied to specific users, teams, or across the entire organization.

GenAI is a fast-growing technology, and as it continues to evolve, Chrome Enterprise will too – providing customers with a productivity platform that enhances the way people work, elevates security, and unlocks opportunities for their business and workforce both now and in the future.



Accessing GenAI capabilities in Chrome Enterprise

To enhance productivity, help employees complete manual tasks faster, quickly create content, and streamline customer engagement, Chrome Enterprise offers many built-in GenAI capabilities.

Check out the [full list](#) of Chrome's GenAI features, which includes:








Feature	Description
→ Tab Compare	Compares information across tabs.
→ History Search Settings	Searches browsing history with AI-powered answers.
→ DevTools Gen AI	Enables AI-powered debugging in DevTools, requiring select debugging data to be sent to Google.

In addition to Chrome-built GenAI features, businesses can also benefit from other Google GenAI features like Google Lens and Google Agentspace, which are accessible in Chrome. Please note that these features are subject to their own terms of service.

Feature	Description
→ Google Lens	Empowers users to search visually with AI – enabling camera, voice, and video-based queries; real-time product insights; image verification; and desktop integration – all with enterprise-grade speed, context, and accuracy. See the Google Lens Terms of Service .
→ Google Agentspace	Unifies enterprise data, Gemini's AI, and secure tools like NotebookLM Plus into a single, branded AI hub – enabling advanced search, insight generation, and task automation across your organization with strong privacy and customization controls. See the Agentspace Service Specific Terms , Google Cloud Platform Terms , and Cloud DPA .
→ Gemini	<p>Gemini in Chrome is designed to help you get more done while you're browsing. This goes beyond having easy access to the existing Gemini experience; Gemini in Chrome can use the context of the webpage you're currently viewing to provide relevant responses.</p> <p>Supported by the GenAiDefaultSetings policy. See the Google Workspace Terms and the Workspace Specific Service Terms.</p>

The security, privacy, and compliance of GenAI features in Chrome

Before your enterprise introduces any new GenAI feature, it is important to first understand the security, privacy, and compliance of the capability you're interested in implementing. You'll want to consider:

- | | |
|--|---|
|  Data privacy |  Retention |
|  What data is collected |  Encryption standards |
|  Why that data is collected |  Data processing and storage |
|  Processing limitations
(Terms of Service) | |

When it comes to these considerations, it will be helpful to know that Chrome adheres to the following policies and guidelines:

Data privacy

Chrome's GenAI tools follow strict rules for [data minimization](#), only collecting what is necessary. The data is also collected in a manner by which Google cannot attribute it to any specific user. Our tools use on-device models when possible but fall back to the server models when necessary to provide higher-quality results.

What data is collected

Most of Chrome's GenAI capabilities collect usage and browser metrics, user content (such as submitted inputs, prompts, files, page context), and generated outputs. For more details on the data collection for each feature, please see our [Help Center articles](#).

Why that data is collected

User content, usage data, and browser metrics are primarily collected to:

1. Deliver the service
2. Prevent abuse and malicious behavior
3. Improve the service (user content will not be used for this when [Chrome's generative AI enterprise data guarantees](#) are enabled)

Processing limitations (Terms of Service)

Use of these services and features is governed by [Google's Terms of Service](#) and [Google Privacy Policy](#), unless you have an agreement that states otherwise.

Retention

Data is only retained as long as necessary to ensure the functionality, quality, and safety of the service. When model training is disabled, any collected user content, prompts, and model outputs are generally stored for 28 days or less for security and abuse prevention purposes. You can see each feature for more specific details. Other pseudonymized metrics data is typically stored for up to 18 months. To better understand Chrome's general data retention practices, see our [data protection center](#) for more details.

Encryption standards

Chrome's [standard encryption practices](#) are applied. All stored and in-transit data is encrypted. Data in transit is protected by TLS/SSL certificates, while data at rest is [encrypted at various layers](#). For more details on how data is strongly protected, you can read [Chrome's SOC 3 Report](#).

Data processing and storage

Chrome optimizes performance and efficiency, often processing and storing data in the closest available region to our users. More details about Chrome's storage and processing locations can be found in [Chrome's SOC 3 Report](#).

Managing GenAI capabilities via policy

Each Chrome-built GenAI capability has an associated policy that IT and security teams can use to manage how the capability is utilized by their organization. With Chrome Enterprise Core, all Chrome policies, across all desktop and mobile platforms, can be centrally managed from the cloud-based Google Admin console.

You can access a full list of policies for Chrome's GenAI capabilities [here](#). Some examples include:

GenAI capability	Description	Policy
Tab Compare	Compares information across tabs.	TabCompareSettings
History Search Settings	Searches browsing history with AI-powered answers.	HistorySearchSettings
DevTools Gen AI	Enables AI-powered debugging in DevTools, requiring select debugging data to be sent to Google.	DevToolsGenAiSettings

Policy States Each AI feature generally has the following states for you to choose from:

State	Description	Default
0	Allow feature and improve AI models.	Third-party managed or unmanaged browsers
1	Allow feature without improving AI models. Note: User content (e.g., inputs, prompts, and page content) may be sent to Google to deliver the service and prevent abuse.	Browsers managed with Chrome Enterprise Core or Chrome Enterprise Premium
2	Do not allow feature. Note: May limit core browser features that reduce the productivity and creativity benefits for your users.	N/A

Global AI governance policies For global AI governance that does not require IT or security teams to set each policy at the individual feature level, Chrome Enterprise provides these additional policies:

Policy	Description
GenAiDefaultSettings (Only available to Chrome Enterprise Core and Chrome Enterprise Premium customers)	<ul style="list-style-type: none">Define preferred default settings for new and existing generative AI featuresDoes not impact any manually set policy values
GenAILocalFoundationalModelSettings	<ul style="list-style-type: none">Ensures the browser is equipped with on-device models (if device specs support it) that can provide faster results, minimize data collection, and empower developers to leverage pre-existing AI capabilitiesReduces reliance on external servers, resulting in faster response times and improved system reliabilityOptimizes bandwidth usage and alleviates network congestion by decreasing the volume of data exchanged between devices and cloud services

Creating a strong AI governance strategy

We hope this framework helps your IT and security teams build scalable, policy-driven oversight that works for your organization.

Step 1

Determine the data categories for your organization

Every organization works with lots of different types of data. A common approach is to group that data based on how sensitive it is. Here are a few general categories to consider:

- **Public data:** generally available to the public.
- **Company data:** exclusively accessible by people within the organization.
- **Proprietary/sensitive data:** confidential trade secrets, designs, plans, financial data, or protected information.
- **Strictly need-to-know data:** information that is only allowed to be accessed by small authorized groups within your organization, due to its sensitive nature.
- **Data not allowed on third-party servers:** data that should never leave the organization's boundaries due to the highest level of sensitivity. In some cases, this may apply to extremely sensitive government data if the third-party platform does not meet the compliance requirements for the government organization.

Step 2

Define the types of GenAI tools your organization allows

There are many ways to classify a GenAI tool. For example, you can classify a tool based on:

- What it outputs (e.g., image, video, text)
- What are its inputs (e.g., pre-selected vs open inputs)
- Where the model processes data (e.g., on-device vs server-based)
- How the data is used (e.g., model improvement vs service delivery)

Every organization uses different factors to classify GenAI for users. Take some time to figure out what works best for you. In this case, we'll use:

- Server-based where the feature uses user content to improve
 - A communal model
 - A model that is exclusively available to your organization
 - A model that is exclusive to a specific set of users
- Server-based feature that does not use any user content to improve any models
- On-device feature that does not use any user content to improve any model

Step 3

Create a simple mapping to determine which GenAI tools can be used with different types of data

This helps make it clear what's allowed – and what's not – within your organization, so teams can confidently choose the right tools to boost productivity, work more efficiently, and get creative.

In the example below, you'll see data categories across the top and AI model types in the first column.

User content is used to ...	Public data	Company data	Proprietary/ sensitive data	Strictly need-to-know	Data not permitted to be sent to a third-party server
Improve communal models (server-based)	✓	✗	✗	✗	✗
Improve models exclusively available to your organization (server-based)	✓	✓	✗	✗	✗
Improve user- and team-specific models (server-based)	✓	✓	✓	✓	✗
Deliver service but not to improve ANY models (server-based)	✓	✓	✓	✓	✗
Deliver service (on-device)	✓	✓	✓	✓	✓

Step 4

Share a clear list of approved (and not-approved) AI tools to help everyone make smart choices

Giving users a clear view of which tools are good to go—and which ones aren't—makes it easier for them to choose the right tools while keeping data safe. It helps reduce risk, avoid accidental data leaks, and keeps everyone on the same page. Plus, sharing how these decisions are made builds trust and helps teams stay thoughtful when exploring new tools.

Step 5

Automate your AI policy management

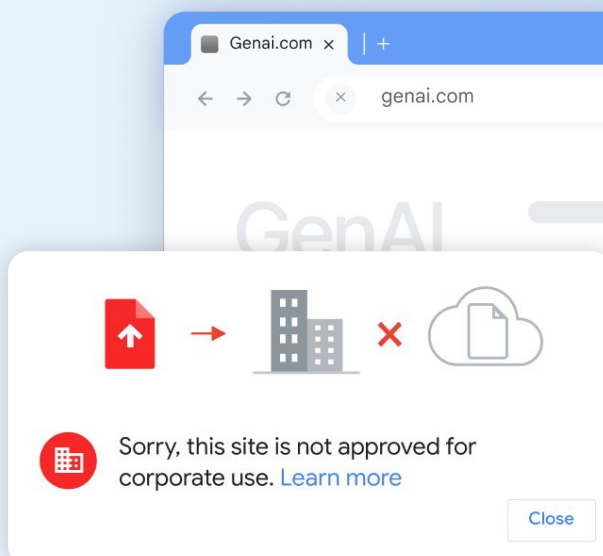
Chrome Enterprise Core and Chrome Enterprise Premium admins can use the [GenAiDefaultSettings](#) policy to define their organization's default settings for all AI capabilities. Admins don't need to manage this at the individual generative AI feature level and can confidently expect new capabilities to launch with the controls they have selected.

Even with default settings in place, Chrome Enterprise customers are recommended to subscribe to the [Chrome Enterprise and Education release notes](#) to hear about upcoming GenAI capabilities.

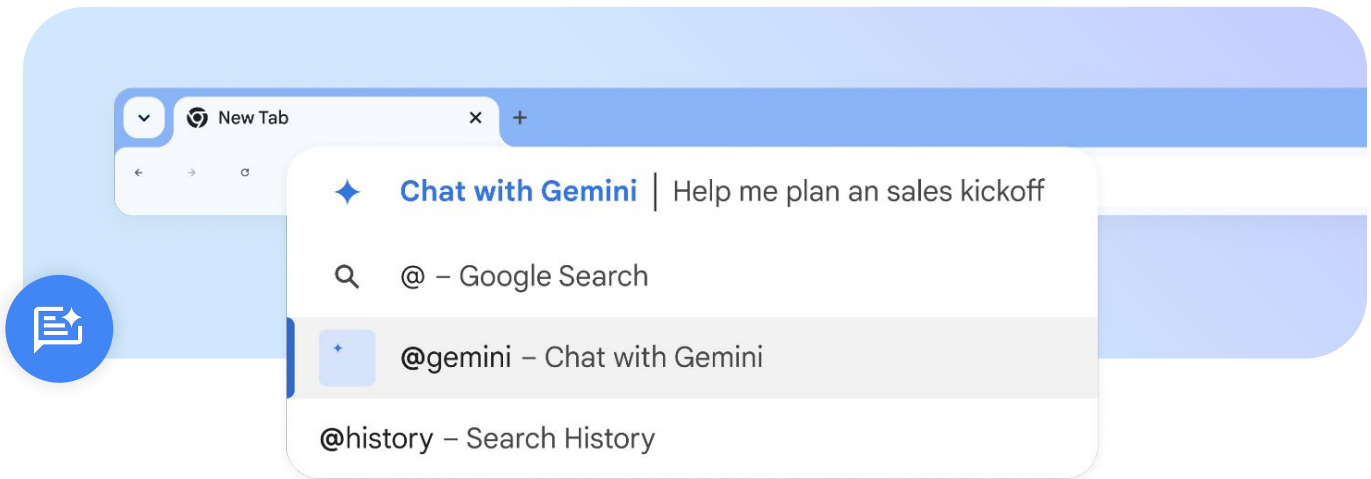
Taking GenAI security a step further with Chrome Enterprise Premium

Once you've got your AI governance strategy in place, you might want to level up your security. [Chrome Enterprise Premium](#) gives IT and security teams powerful tools like context-aware access controls, data loss prevention (DLP), URL filtering, and advanced protections against malware and phishing – all designed to help keep your organization's browsing secure.

When it comes to GenAI, these features can help block access to unapproved tools and protect sensitive data from being shared where it shouldn't be.



Additional considerations for AI best practices in your organization



To help, here are a few questions IT and security teams might want to consider as they build out their approach:

1. When should your enterprise users let others know that a piece of work was AI-generated?
2. Are there specific tasks that you don't want teams using generative AI for?
3. Are there tasks that you want to encourage using generative AI for?
4. Should there be limits on how much of a final work can be AI-generated?
5. Are there any types of personal or sensitive data that should never be shared with AI tools?
6. What's the best way for users to double-check AI-generated work?
7. How can teams spot potential AI "hallucinations" or inaccuracies?

Generative AI is growing fast, and here at Google, we are focused on building secure, trusted tools to help you make the most of it. These tools provide a strong foundation upon which you can build a robust set of best practices for users within your organization.

Appendix

Shaping your organization’s GenAI strategy and governance policy

Next Steps

Use the worksheet below as a starting point for your company’s generative AI strategy – it’s designed to help you create a governance policy that aligns with your organization’s specific security and compliance needs.

Just a reminder, the policy states include:

- 0 (Default): Enabled with data available for AI model improvements
- 1 (Enterprise default): Enabled, but Google cannot use data for AI model improvements
- 2 (Disabled): Feature is blocked

Capability	Policy	Preferred policy state
Create Themes with AI	HelpMeWriteSettings	
Tab Compare	TabOrganizerSettings	
AI History Search	TabCompareSettings	
Help Me Read	HistorySearchSettings	
Generative AI VC Background	HelpMeReadSettings	
Generative AI Wallpaper	GenAIVcBackgroundSettings	
Admin Assistant	GenAIWallpaperSettings	
Enables AI-powered debugging in DevTools, requiring select debugging data to be sent to Google	DevToolsGenAiSettings	

Global AI Governance	
Defines default AI feature settings. Overrides apply only if manually set.	GenAiDefaultSettings
Manages local AI model downloads and inference.	GenAILocalFoundationalModelSettings

User content is used to ...	Public data	Company data	Proprietary/ sensitive data	Strictly need-to-know	Data not permitted to be sent to a third-party server
Improve communal models (server-based)					
Improve models exclusively available to your organization (server-based)					
Improve user- and team-specific models (server-based)					
Deliver service (but not to improve ANY models (server-based)					
Deliver service (on-device)					

Helpful GenAI resources

The following resources can help IT and security teams learn more about Chrome's current and upcoming GenAI capabilities, and how their organization can use Chrome Enterprise to optimize its utilization of this powerful new technology.

- [AI innovations in Chrome](#)
- [Chrome DevTools AI](#)
- [Chrome Enterprise AI Capabilities](#)
- [Chrome Enterprise overview](#)
- [Chrome – Generative AI features and policies](#)
- [Chrome Enterprise blog](#)
- [Chrome Enterprise Generative AI policy list](#)
- [Contact a Chrome Enterprise expert](#)
- [Chrome Enterprise Generative AI help](#)