

Next '25 Recap slides

Recap of announcements,
claims, offers, and other news



Proprietary



Contents



01 [Executive AI Overview](#)

02 [Overview & Keynote Recap](#)

03 [Google Announcements](#)

[Infrastructure](#)

[Developer](#)

[Databases & Data Analytics](#)

[AI Announcements](#)

[Security](#)

[Google Workspace](#)


04 [Offers](#)

01.

Executive AI overview

At Google Cloud we see AI as the most important way we can help advance your mission — the opportunity with AI is as big as it gets. That's why we are investing in the full stack of AI innovation.

Infrastructure



Key announcements

AI Hypercomputer

- Ironwood, our 7th generation TPU, provides unparalleled performance for demanding AI workloads
- vLLM for TPUs for inference across Compute Engine, GKE, Vertex AI, and Dataflow. Including simultaneous TPU & GPU use with GKE custom compute classes
- Pathways, Google's internal ML runtime for large scale training and inference workloads now available with Jax


Cloud WAN

Globally connect all infra, apps, and employees, with improved network performance and reduces total cost of ownership.

Gemini on Google Distributed Cloud

Run Gemini on your own on-premises infrastructure. Gemini models can run on top of NVIDIA DGX and HGX platforms.

AI platform



Key announcements

Gemini 2.5 Flash & Pro on Vertex AI

Gemini 2.5, our next generation of "thinking models", now with Flash and Pro available on Vertex AI.

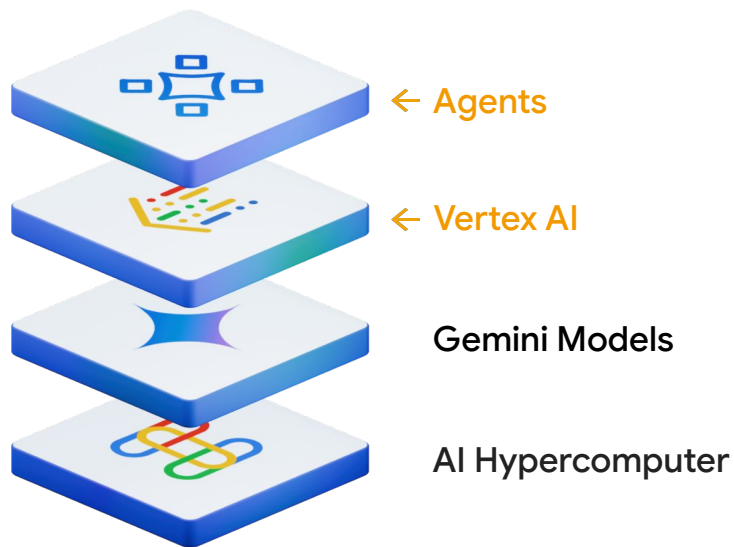
Generative AI Media Models (Veo 2.0, Lyria, Chirp 3)

Expanding our suite of gen AI media models with public availability of Veo 2.0, announcement of Lyria, and release of Chirp 3 with new audio generation and understanding capabilities.

Grounding in Google Maps

New ability to ensure that agent responses relying on location context are factual and fresh.

Agent platform



Key announcements

Google Agentspace

Now offering Google expert agents like Idea Generation agent and Deep Research agent; integration with NotebookLM and Chrome Enterprise; and Agents gallery to find right the right agent for the job.


Open Source Agent Development Kit and Agent Engine

Introducing our new open source Agent Development Kit and runtime (Vertex AI Agent Engine) to help developers build sophisticated multi-agent applications. Launched in Python with support for Model Context Protocol, more languages coming in the future.

Agent Interoperability Protocol

Unite your fragmented multi-agent ecosystems with our new open Agent2Agent protocol. Future proof your agent investments with a universal communication standard that works across frameworks.

Packaged agents



Key announcements

Data Agents

Key launches include: Data engineering agent, data science agent, and conversational analytics agent (with API option).

Security Agents

New Google SecOps agent for triaging alerts and cases. New Google Threat Intelligence agent for malware analysis.

Code Assist Agents

New agents to generate code, migrate code, code review, testing, documentation, and Java modernization.

Purpose-Built Agents

Including food ordering AI agent, automotive AI agent, and Customer Experience Suite prebuilt agents for flight booking, ticketing, shopping assistance, and appointment booking.

AI-powered products



Agents

Vertex AI

Gemini Models

AI Hypercomputer

Key announcements

Google Unified Security

Brings together unmatched threat visibility, faster threat detection, continuous virtual red-teaming, the most trusted browser, and Mandiant Expertise solution—in one converged security solution, powered by Google AI, running on a planet-scale data fabric.

Customer Engagement Suite

- Unified console for Conversational Agents combining generative AI and rules-based controls
- AI Coach and AI Trainer to guide human agents and test AI-powered agents
- New features: human-like voices; streaming video support; custom no-code agent builder; API tool calls; and integration with data sources, CRM systems, and 3P platforms

Google Workspace



Key announcements

Data regions support for Gemini

Google Workspace with Gemini now adheres to data regionalization requirements. Admins have flexibility to set up sovereign controls for EU processing, US processing, or a mix of both.

FedRAMP High and BSI C5 certifications

Google Workspace with Gemini is the only AI productivity assistant to achieve the FedRAMP High certification as of today.

Google Workspace with Gemini

- Google Workspace Flows automates work with AI agents in the loop
- 3P data connector framework
- Audio generation in Google Docs
- Help me analyze in Google Sheets

Gemini customers around the world



and many more...

All Next '25 Announcement highlights

Welcome


What's new in the world of Google Cloud? Explore the latest themes, announcements, customer stories, and more from our global exhibition of inspiration, innovation, and education.




Infrastructure



Developer




Data



Security



Agents &
Applications



Google
Workspace

AI Models & Platform

Global network

Infrastructure Highlights

Google Cloud announcements

Cloud WAN: Starting today, this network, which moves at "Google speed" — near-zero latency — for billions of users worldwide, is now available to enterprises everywhere. This backbone-as-a-service offering, takes advantage of Google's 25 years of networking expertise to create a fully managed, high-performance global WAN. Cloud WAN provides high-performance connectivity between geographically dispersed data centers, or connects branch and campus environments over our Premium Tier network. ([learn more](#))

Google's AI Hypercomputer now features **Ironwood TPU** ([learn more](#)), our most performant and scalable TPU to date. Train, tune, and serve AI models with unprecedented efficiency. **Forget speeds and feeds;** this is about **"intelligence per dollar,"** which for customers translates to more efficient AI agents, reduced operational bottlenecks, and ultimately, a better ROI on AI investments. ([learn more](#))

Gemini on Google Distributed Cloud brings Google's hardware and software — designed for sovereignty, regulatory, and low-latency requirements — to customer environments (data center, on-premises) so they can bring AI capabilities closer to where data is generated. It ensures optimal performance for demanding and highly sensitive AI workloads. ([learn more](#))

Customer & Partner Announcements



Toyota deployed ML models for factory workers, saving 10,000 employee hours per year.



Leveraging Google Cloud's highly reliable and resilient infrastructure to ensure business continuity and data protection for mission-critical workloads.



Deploying and managing Cloud GPUs for advanced drug discovery.



Developed a compliant infrastructure platform for AI development that advances aircraft performance, safety, and reliability.

Key Stats

- Cloud WAN delivers up to **40% improved network performance**, while **reducing total cost of ownership by up to 40%**. [Source](#)
- Gemini Flash 2.0, powered by AI Hypercomputer, achieves **24x higher intelligence per dollar** compared to GPT-4o, and 5x higher than DeepSeek-R1. [Source](#)
- Ironwood TPUs deliver a staggering **42.5 exaFLOPS of compute per pod**. [Source](#)

AI Model Innovation Highlights

Google Cloud Announcements

Gemini 2.5, our most intelligent AI model yet, brings enhanced reasoning to enterprise use cases. ([learn more](#))

- **Gemini 2.5 Pro Preview:** SOTA model for complex reasoning, coding, math now available on Vertex AI.
- **Gemini 2.5 Flash** (coming soon): Fast, efficient thinking model optimized for speed/cost.
- **Vertex AI Model Optimizer:** A user establishes a serving budget, calls a single endpoint and gets the best possible answer using all models, context, and tools at our disposal within that budget.
- **Live API:** Production-ready API for real-time, interactive voice and vision agents powered by Gemini 2.0 Flash.

Generative AI Media Models available in Vertex AI, now across all modalities. ([learn more](#))

- **Introducing Lyria** (text-to-music) in Private Preview
- **Veo 2** (text-to-video and image-to-video) in GA with Allowlist, and **new Veo 2 editing features** in Preview
- **Chirp 3 instant custom voice and diarization** (audio generation and understanding) in Preview
- **Imagen 3** (text-to-image): improved image generation, editing, and inpainting capabilities

Customer & Partner Announcements



Speeding up campaign creation from 8 weeks to 8 hours using Imagen and Veo on Vertex AI.



Creates unique visuals and videos of travel destinations with Imagen and Veo on Vertex AI.



Bringing advanced Imagen 3 and Veo 2 models to creative professionals, starting with Adobe Express and Project Concept.



Uses Veo 2 as a creative partner to generate diverse, cinematic shots across 20 countries and languages.



Accelerates content creation for consumer-ready visuals at scale, using generative AI, for global brands like Oreo and Cadbury.

Key Stats

- With **Gemini, Imagen, Chirp, Lyria and Veo**, Google is the **only company** that offers **generative media models across all modalities** and all are available today on Vertex AI. ([Source](#))
- We recently launched Gemini 2.5 Pro which is state-of-the-art on a wide range of benchmarks and **debuts at #1 on LMArena** by a significant margin. ([Source](#))
- Since Next 2024, we've seen over **40X growth in Gemini use on Vertex AI**, now with billions of API calls per month. ([Source](#))
- More than **4 million developers** are using Gemini. ([Source](#))

AI Agents & Platforms Highlights

Google Cloud Announcements

Customer Engagement Suite ([learn more](#))

- Conversational Agents console
- Prebuilt agents for common use cases
- Third-party connectors to data sources and CRM
- AI Coach & AI Trainer for customer service reps

[Agents] **Vertex AI Agent Builder:** New ways to build and manage multi-agent systems with Vertex AI ([learn more](#))

- Agent Development Kit (ADK)
- Vertex AI Agent Engine
- Agent2Agent Protocol
- Agent Garden
- Grounding with Google Maps

[Agents] **Google Agentspace** ([learn more](#))

- Agents gallery
- Agent Designer (no-code agent builder)
- Two new pre-built expert agents from Google: Idea generation agent and deep research agent
- Support for partner agents through Google Cloud Marketplace
- Chrome search bar integration

Customer & Partner Announcements



Transforming the restaurant experience by integrating AI into core operations to spot issues, predict maintenance, answer key questions, and much more.

Honeywell

Using Vertex AI, BigQuery, and Gemini to streamline product lifecycle management, aiming to accelerate engineer output by up to 70%.



Developing next-gen AI agents, with Google Agentspace, to dynamically orchestrate procurement, internal audit, accounting, and finance expertise.

verizon

Personalizing the customer service experience for more than 115M connections with Google Cloud's Customer Engagement Suite.



The Home Depot built "Magic Apron" using Gemini and Vertex AI. This agent delivers expert home improvement guidance 24/7 online and in their app.

Key Stats

- **Tens of thousands of enterprises are building** with Vertex AI and Gemini. [Source](#)
- **Google Agentspace** is our **fastest-growing enterprise offering** with a powerful set of features unique to Google.
- As of April 8, **50+ partners** such as Box, Deloitte, Elastic, Salesforce, ServiceNow, UiPath, UKG, Weights & Biases, and many more are committed to working with us on the **Agent2Agent protocol**.
- We are seeing huge momentum for Vertex AI Search with **billions of daily queries** executed by our customers. [Source](#)

Developer Highlights

Google Cloud Announcements

Firebase: We've introduced Firebase Studio, a cloud-based, agentic development environment powered by Gemini that includes everything developers need to create and publish production-quality AI apps, as well as several more updates across the Firebase platform. ([learn more](#))

Gemini Code Assist: Gemini Code Assist is introducing agents to help developers with tasks like code migration and translation, feature implementation, code review, unit test generation, documentation, and more. Gemini 2.5 is now available in Gemini Code Assist for individuals to use in personal projects. ([learn more](#))

Gemini Cloud Assist: Announcing new features to simplify cloud operations and management with AI, including Application Design Center Integration, Investigations, enhanced chat experience, and more. ([learn more](#))

Cloud Hub: A brand new product! Cloud Hub is a central, unified hub designed to streamline Google Cloud operations and management. ([learn more](#))

Cloud Run: Cloud Run GPU and functions are now generally available, and have even more features in preview, including the ability for developers to deploy apps from Vertex AI Studio to Cloud Run in a click, and ability to run Cloud Run in multiple regions for high availability.

Customer & Partner Announcements



Partners including Atlassian, Sentry, and Snyk launched pre-built connections, called tools, into Gemini Code Assist to bring more context into the IDE. Customers include Major League Baseball where we're announcing the winner of the MLB + Google Cloud Hackathon, Sphere showcasing the enriching of the Wizard of Oz with Google AI, and X-Games highlighting a Google AI-driven fan experience this past Aspen Games.

Key Stats

- CME Group, which runs the Chicago Mercantile Exchange, states that Gemini Code Assist provides most of its developers more than **10.5 hours of productivity gain** per month. [Source](#)
- Developers are estimated to get common development tasks done over **20% faster** with Gemini Code Assist.
- GKE delivers differentiated performance for gen AI applications with the addition of new inference capabilities. These capabilities **reduce serving costs by up to 30%, tail latency by up to 60%, and increases throughput by up to 40%** compared to other managed and open source Kubernetes offerings based on our internal benchmarks. [Source](#)
- **Millions of developers use Firebase** to engage their users, powering over 70 billion instances of apps every day, everywhere — from mobile devices and web browsers, to embedded platforms and agentic experiences. [Source](#)

Databases & Data Analytics Highlights

Google Cloud Announcements

Data Analytics ([learn more](#))

BigQuery's next evolution: The autonomous data to AI platform to provide customers with a flexible, integrated, and intelligent platform with automation and agentic capabilities for the entire data lifecycle. Including Gemini in BigQuery.

Looker updates include a full conversational analytics suite of features that allow customers to interact with their data more easily and intuitively, including Gemini in Looker, Conversational Analytics API, and Looker Reports.

Databases ([learn more](#))

AlloyDB AI is evolving with Google Agentspace integration, natural language support, optimized SQL and vector search, three new gen AI models, and more.

Expanded Oracle and Google Cloud services: Oracle Base Database Service, a flexible, controllable way to run Oracle databases in the cloud.

SQL Server modernization: Database Migration Service supports SQL Server to PostgreSQL migrations for Cloud SQL and AlloyDB.

Customer & Partner Announcements



Synthesizes millions of points of customer feedback to develop insights that deepen customer relationships and drive product innovation.



Built Distributed Trade Transformation on BigQuery and Looker to process 75,000 orders daily in emerging markets.



Enhanced flexibility using Spanner, which has enabled operational continuity while optimizing cost and performance.



Improved the online search experience, enhancing natural language search query accuracy by 20%, by using AlloyDB to combine structured and unstructured data.

Key Stats

- Together BigQuery and Vertex AI are between 8–16x more cost efficient compared to other data warehouse and AI platforms. [Source](#)
- BigQuery has 5x more customers than the two leading cloud companies that exclusively offer data warehouse and data science platforms. [Source](#)
- Compared to the HNSW index in standard PostgreSQL, AlloyDB AI's ScaNN index delivers up to 4x faster vector search queries and up to 10x faster filtered vector search queries. [Source](#)

Security Highlights

Google Cloud Announcements

Google Unified Security (Generally Available) is a converged security solution powered by AI that brings together best-in-class security products for threat intelligence, security operations, cloud security, and secure enterprise browsing, along with Mandiant Expertise solution to provide a unified view and improved threat detection across complex infrastructures.

Security Agents & Gemini is transforming security operations with 1) Alert triage agent, which performs dynamic investigations of alerts and renders a verdict. And 2) Malware analysis agent, which autonomously identifies security issues and recommends whether code is malicious or not.

Google Security Operations can now deliver end-to-end data pipeline management. These new capabilities empower customers to better manage scale, reduce costs, satisfy compliance mandates, and most importantly, drive better security outcomes.

Security Command Center now includes AI Protection to secure AI systems and sensitive AI training data, along with a Compliance Hub to simplify regulatory management. These updates aim to provide comprehensive security and compliance visibility across Google Cloud.

[\(learn more\)](#)

Customer & Partner Announcements



Uses Google SecOps to proactively respond to cyber threats, gaining visibility and reducing investigation and resolution time.



Centralizes the monitoring of AI security threats alongside their other cloud security findings in Google's Security Command Center.



Enhanced security operations ingesting 22 times more log data, detecting three times more events, and closing investigations 50% faster.



Used Vertex AI along with open-source tools and Google Cloud's security foundation to establish an AI security governance layer.

Key Stats

- Mandiant spent 450k+ hours investigating incidents in 2024. Source: Google Internal Data
- With Gemini in Security Operations, we are the first major cloud provider in the market to make generative AI generally available to defenders in a unified SecOps platform. [Source](#)

Google Workspace Highlights

Google Cloud Announcements

Google Workspace is expanding its AI capabilities with Gemini to enhance productivity and collaboration through features:

- **Help me analyze** transforms Google Sheets into a business analyst, identifying insights from your data without the need for prompting that help you make informed decisions. (coming soon)
- **Audio Overviews** in Google Docs, where you can interact with Docs in an entirely new way by creating high-quality audio versions of your content. (coming soon)
- **Google Workspace Flows**, where Gemini's agentic capabilities help you automate time-consuming, repetitive tasks and make decisions with more context. (Preview)

[\(learn more\)](#)

Customer & Partner Announcements



Uses Google Workspace with Gemini to transform building vehicles, driving efficiency, collaboration, and creativity to accelerate building groundbreaking, sustainable vehicles.



Migrating 575,000 employees to Google Workspace to deliver a sovereign cloud-based collaboration solution for regulated industries in Germany and the EU.

FRESHFIELDS

Adopted Google Workspace with Gemini to enhance legal work productivity and collaboration, expedite tasks. And conduct due diligence.



Using Gemini in Sheets to reduce time spent analyzing customer feedback from hours to minutes, enabling them to quickly identify trends and improve customer interactions.

Key Stats

- We're seeing over 2 billion AI assists monthly to business users within Google Workspace. [Source](#)
- We're seeing major boosts in productivity through Gemini for Google Workspace, with customers saving an average of 105 minutes per user, per week, according to our recent study of enterprise customers.¹
- A Forrester Total Economic Impact™ study, commissioned by Google, estimates that Google Workspace's secure and collaborative environment can result in a 336% return on investment and net present value of \$57.3 million over three years.²

¹ Based on an internal Google analysis of survey data from completed Gemini for Google Workspace pilots. Some pilots began in December 2023 and some are still in progress. Data was aggregated in July 2024, across 18 enterprise customers globally, spanning North America, the Middle East, and Asia Pacific. We looked at more than 3,200 responses from employees across a broad range of industries, including manufacturing, healthcare and life sciences, retail, energy, financial services, and software and technology.

² The Total Economic Impact™ of Google Workspace, January 2024, a commissioned study conducted by Forrester Consulting

02. Overview & keynote recap

Details sourced from publically available Next '25 event details, opening keynote, and Thomas Kurian's welcome address.

Cloud Next themes

Theme 1

*Google Cloud is the **most enterprise-ready cloud for the AI era**, and we're winning*

Message: Google Cloud is the only hyperscaler that offers an open platform that is fully multi-cloud, works with all of your current investments, and addresses evolving sovereignty, security, privacy, and regulatory requirements.

Theme 2

*Google Cloud offers **highly differentiated AI products, built for multi-agent systems***

Message: Only Google Cloud offers an AI-optimized platform with leading price, performance, and precision. It includes advanced infrastructure and databases, world-class models, a robust developer platform in Vertex AI, and the most comprehensive agent portfolio, including the broadest range of enterprise-ready tools to build AI agents to build and manage multi-agent systems.

Theme 3

*Customers are getting **real business value** from Google AI investments*

Message: 500+ customers, developers, and partners around the world are using Google AI to deliver real business transformation and impact.

Theme 4

*Google Cloud has the **best AI partner ecosystem***

Message: Google Cloud is working with leading ISVs and System Integrators to make interoperability between agents easy and to speed up their ability to build on top of its Enterprise AI platforms.



Keynote recap (1/2)

Technology is evolving and it can enable us to reimagine what's possible. The chance to improve lives and reimagine things is why Google has been investing in AI for more than a decade. Getting technology advances into the hands of both consumers and enterprises is core to Google's mission.

With Google Cloud, AI is the most important way to help drive real business growth and transformation. Google Cloud is the most enterprise-ready cloud for the AI era, offering an open, multi-cloud platform, and building for interoperability, speeding up time to value from AI investments.

More details: [Watch Google Cloud Next '25 Opening Keynote](#); [Read Thomas Kurian's welcome address for Next '25](#)

Keynote recap (2/2)

Google is the only company investing in the full stack of AI innovation with the best infrastructure, leading models, tools, and agents, enabling leading price, performance and precision.

- **Infrastructure:** Our AI Hypercomputer includes hardware, software, and consumption models — all optimized to deliver more intelligence at a consistently low price for training, tuning, and serving AI workloads. We continue to make massive improvements in TPU performance and efficiency, and now introducing Ironwood: more than a 10x improvement from our most recent high performance TPU, with over 9,000 chips per pod. [TPU v7 Ironwood, Google Cloud WAN]
- **Research & Models:** Building on the groundbreaking research of Google DeepMind, we're delivering rapid innovation across many AI models, starting with Gemini, our most capable family of AI models. [2.5, Veo, Lyria, Imagen, Chirp]
- **Products & Platforms:** We're continuing to invest in Vertex AI, our comprehensive platform to orchestrate the three pillars of production AI — models, data, and agents. We've seen 20x growth in Vertex AI over last year and have 500+ customer stories showing strong ROI. [Vertex AI innovations - agent interoperability protocol]
- **Agents:** Only Google Cloud offers a robust developer platform in Vertex AI, and the most comprehensive agent portfolio, including the broadest range of enterprise-ready tools to develop agents, and build and manage multi-agent systems. [Google Agentspace, agents for data, security, and code]

More details: [Watch Google Cloud Next '25 Opening Keynote](#); [Read Thomas Kurian's welcome address for Next '25](#)

As seen in the Next '25 Keynote,

Our customers are building the future

[Explore the videos](#)

Deutsche Bank 

Honeywell

INTUIT



verizon

Deutsche Bank uses DB Lumina, an AI research agent powered by Gemini and Vertex AI, to achieve faster, secure data analysis, driving productivity and competitive advantage.

Using Vertex AI, BigQuery, and Gemini to streamline product lifecycle management, aiming to accelerate engineer output by up to 70%.

Intuit simplifies tax filing for millions with Google Cloud Document AI to enhance data comprehension and autofill, creating a "done-for-you" experience that frees up users' time.

Mattel synthesizes millions of points of customer feedback to develop insights that deepen customer relationships and drive product innovation.

McDonald's leverages Google Cloud's AI and edge computing to revolutionize restaurant operations, enhance customer experiences, and empower its workforce.

Mercado Libre is utilizing Vertex AI Search to revolutionize product discovery for its 100M customers across Latin America, significantly boosting revenue with more intuitive searches.

The Nevada Department of Employment, Training and Rehabilitation developed an AI-powered Appeals Assistant on BigQuery and Vertex AI to process case approvals 4x faster.

Reddit Answers harnesses AI, powered by Vertex AI Search, to deliver authentic, human-generated insights and discussions from Reddit's vast conversational database.

Sphere used Google Cloud's AI and infrastructure, including DeepMind, to create immersive, high-impact experiences like The Wizard of Oz, accelerating production and delivery.

Verizon is personalizing the customer service experience for more than 115M connections with Google Cloud's Customer Engagement Suite.

03. Announcements deep-dive

Announcements Overview (1/2)

Infrastructure	AI Model Innovation	AI Agents & Platforms	Developer
AI Hypercomputer : A4 and A4X VMs, and 7th gen TPU Ironwood	Gemini 2.5 : Pro and Flash	Vertex AI Agent Builder: Agent Development Kit (ADK), Agent Engine, Agent Garden, Agent2Agent Protocol, Grounding in Google Maps	Firebase : Firebase Studio & Gemini Code Assist Agents
Google Distributed Cloud : Gemini on-premises, and Google Agentspace search	Veo 2 : Text to Video Generation, Image-to-Video Generation		Gemini Code Assist : Agents and tools, and Gemini 2.5 integration
Networking : Cloud WAN added to our Cross-Cloud Network solution	Imagen 3 : Improved quality and editing	Google Agentspace : Agents gallery, Agent Designer (no-code agent builder), 2 new pre-built agents, support for partner agents, Integration with Chrome	Gemini Cloud Assist : Application Design Center Integration, Investigations
Workload-optimized Storage : Anywhere Cache, Google Cloud Managed Lustre, Storage Intelligence and Rapid Storage	Chirp 3 : Instant custom voice + Diarization		Cloud Hub : New! Streamline Google Cloud operations and management
Axion processors : New database integrations, and ISV adoption	Lyria : our new text-to-music model	Customer Engagement Suite : Conversational Agents console, Prebuilt agents for common use cases	Cloud Run : Cloud Run GPU and Cloud Run functions, and high availability.
Compute : Introducing C4D, H4D, Z3H, M4, and Hyperdisk Instant Snapshots			Google Developer Program : Enhanced benefits and Enterprise tier
GKE : GKE Inference Gateway, Cluster Director, and Autopilot for standard clusters			

Announcements Overview (2/2)

Databases & Data Analytics

[BigQuery](#): Gemini in BigQuery, BigQuery multimodal tables, and metastore.

[Looker](#): Gemini in Looker, Conversational Analytics API, and Studio in Looker

[General Database updates](#): Arm support for Cloud SQL and AlloyDB, SQL Server modernization

[AlloyDB](#): Integration with Agentspace, AlloyDB AI natural language, Optimized SQL and vector search, AlloyDB AI query engine

Security

[Google Unified Security](#): An all new converged security solution.

[Security Agents & Gemini](#): Alert triage agent, Malware analysis agent

[Google Security Operations](#): Data pipeline management, Mandiant Threat Defense

[Security Command Center](#): AI protection including Model Armor, Data Security Posture Management (DSPM), and Compliance experience

[Chrome Enterprise Premium](#): More options to protect employees and data, including Employee phishing protections, Data masking, and Android Enhancements

[Mandiant Consulting](#): Mandiant retainer for on-demand access to experts, and new partnerships with Rubrik and Cohesity

[Trusted Cloud](#): new security controls and capabilities on our cloud platform to help organizations meet evolving policy, compliance, and business objectives.

Workspace

[Google Workspace Flows](#): Create and run workflows that automate work across the Google Workspace suite.

[3P Data Connector framework](#): Easily build integrations and access third party data

[Gemini in Sheets](#): describe what you would like to do using simple language, and Gemini in Sheets will carry out your instruction

[Help me analyze](#): click to analyze data and proactively surface key insights in Sheets.

[Start a Chat from Gmail \(aka Airmail\)](#): Move conversations from Gmail to Chat or add Gemini as a conversation partner.

[Gemini in conversations](#): Generate a detailed list of open actions, provide summaries, find key details.

[Unified Workspace apps](#): Enables developers to deliver a single app that will work across Workspace, wherever add-ons are supported.

[Vids](#): Create custom, high-quality video clips with only a prompt, without ever leaving Workspace.

[Audio overviews in Docs](#): Consume Docs content in an audio format.

[Help me refine in Docs](#): proactively suggests writing, content, or structural improvements

[Side panel capabilities in Meet](#): In-meeting AI assistant

[AI model in Gmail](#): Holistic AI threat defense model in Gmail

[Gmail DLP & automated classification labels](#): Protect sensitive data from being shared

[Gmail client-side encryption](#): Send private, encrypted emails to any inbox

[Gemini leads the way in compliance](#): A comprehensive set of safety, privacy, and security certifications.

[Data Regions support](#): Sovereign control over your Gemini data.

Infrastructure

Workload-optimized infrastructure,
with a focus on AI and enterprise applications



AI Hypercomputer: Advancements across every layer

GA

A4 and A4X VMs, powered by NVIDIA B200 and GB200 processors

Cluster Director

Pathways on Google Cloud

Preview

Dynamic Workload Scheduler

Hyperdisk Exapools

Pre-announcement

7th Gen TPU Ironwood

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

AI Hypercomputer delivers more intelligence at a consistently lower price:

Advances in workload-optimized hardware:

- **Rapid Storage:** A new Cloud Storage zonal bucket that enables you to colocate your primary storage with your TPUs or GPUs to accelerate AI workloads.
- **7th generation TPU**, Ironwood, coming later this year.
- **Hyperdisk Exapools** allowing you to create and manage up to 1 exabyte of block storage capacity and TiB/s of throughput from a single environment.

Open software enhancements:

- **Pathways for Google Cloud:** Pathways enable a single AI system to generalize across thousands or millions of tasks, to understand different types of data, and to do so with remarkable efficiency.
- **Cluster Director for Slurm:** software designed to make the management of large scale distributed jobs easier
- **vLLM support for TPUs:** enabling PyTorch developers to use TPUs with minimal configuration.
- **Optimized Inference on Google Kubernetes Engine:** consists of two components: GKE Inference Gateway for intelligent routing and GKE Inference Recommendations for intelligent scaling and accelerator selection.

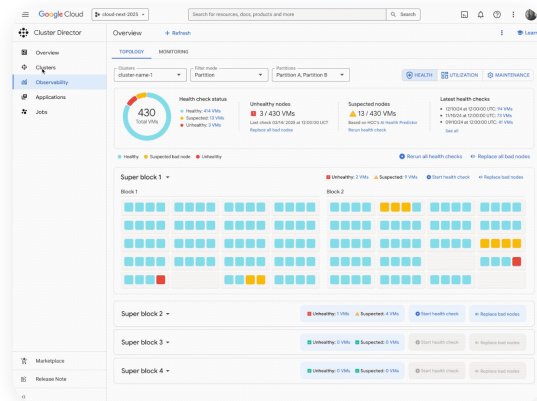
Learn more: [Blog](#), [Website](#)

New flexible consumption options:

- **Dynamic Workload Scheduler**, with TPU and A3 Ultra support

Noteworthy stats

- Announcing new accelerators, including our 7th generation of TPUs. To push the limits of accelerator performance even further, we're announcing Ironwood, with up to **4X more peak compute capacity** and **6X the HBM capacity** compared to Trillium, our 6th generation TPU.
- Rapid Storage provides up to **20x faster random-read data loading** than a Cloud Storage regional bucket.



Google Distributed Cloud (GDC): Gemini, Google Agentspace search

Pre-announcement

Gemini on GDC

Pre-announcement

Google Agentspace search on GDC

Google Distributed Cloud is helping enterprises and governments leverage the best of Google's AI capabilities AI in their own on-premises and edge locations.

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

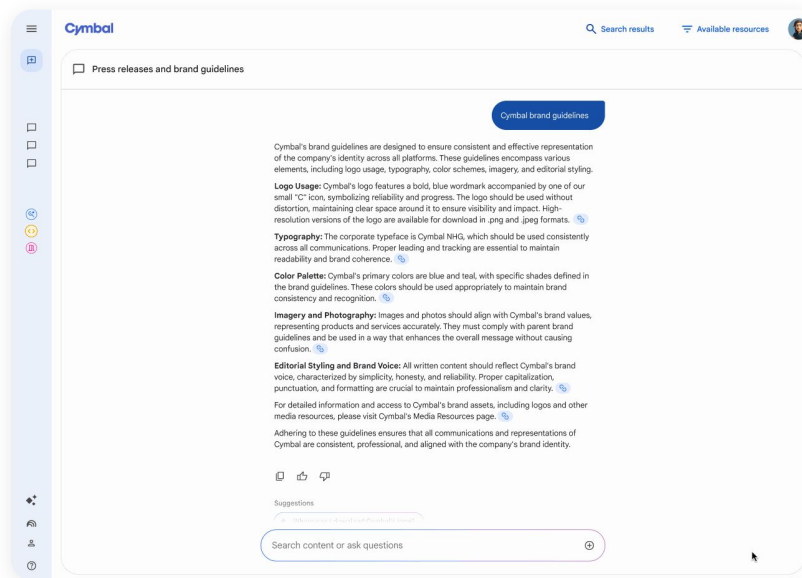
Google Cloud Next

Gemini on-premises

Google Distributed Cloud customers will now be able to run Gemini on their own on-premises infrastructure. Customer can run Gemini models on NVIDIA DGX and HGX platforms that customers can purchase through Google Cloud or directly.

Google Agentspace search

Google Agentspace search brings Google-quality enterprise search on-premises, enabling staff to interact with data and automate business functions. Customers can achieve accurate, and secure answers without compromising on data residency or latency.



Learn more: [website](#), [Next Session 1](#), [Next Session 2](#)

Networking: A network for the AI era

Preview

Cloud WAN

Cloud NGFW (new features)

DNS Armor

At Next '25, we are announcing that Google's global network is now available for large businesses and governments to use with our new Cloud WAN solution.

Cross-Cloud Network updates:

Cloud WAN

Cloud WAN is a fully managed, reliable, and secure enterprise backbone to transform enterprise wide area network (WAN) architectures. Cloud WAN is a Cross-Cloud Network solution.

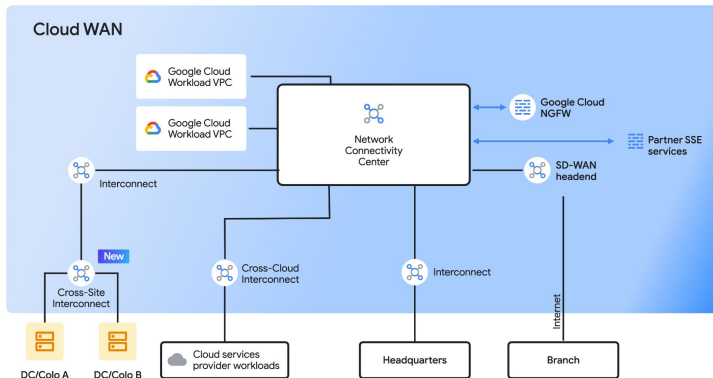
Networking security enhancements:

- Cloud NGFW with Malware Detection:**
 With enhanced features in Cloud NGFW, we are the only vendor to provide protection against 99.5% (PANW WildFire stat) of known and new malware.
- DNS Armor:** detects DNS based data exfiltration and command and controls attacks and it [analyzes 70 Billion events](#) per day and simplifies compliance, protects sensitive data with best-of-breed Data Loss Protection across data-at-rest and data-in-transit

Learn more: [blog](#), [website](#)

Noteworthy customers / metrics

- Metric:** up to 40% faster performance, and up to 40% savings over a customer-managed WAN solution.
- Customers:**



Workload-optimized Storage: For AI and enterprise workloads

GA

Anywhere Cache

Google Cloud Managed Lustre

Storage Intelligence

Preview

Rapid Storage

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

The latest storage advancements:

Cloud Storage Anywhere Cache (GA)

A new strongly consistent cache that works with existing regional buckets to cache data within a selected zone. Anywhere Cache reduces latency up to 70% and 2.5TB/s throughput, accelerating AI workloads and maximizing goodput by keeping data close to GPU/TPUs.

Rapid Storage (Preview)

A new Cloud Storage zonal bucket that provides industry-leading <1ms random read and write latency, 20x faster data access, 6 TB/s of throughput, and 5x lower latency for random reads and writes compared to S3 Express One Zone

Google Cloud Managed Lustre (GA)

A new high-performance, fully managed parallel file system built on DDN EXAScaler, provides petabyte-scale sub-millisecond latency, millions of IOPS, and terabytes-per-second throughput for AI workloads.

Google Cloud Storage Intelligence (GA)

The industry's first offering enabling customers to generate storage insights specific to their environment by querying object metadata at scale, using the power of LLMs.

Learn more: [blog](#), [website](#), [Next Session 1](#), [Next Session 2](#)

Noteworthy customers / metrics

- **Metric:** Anywhere Cache is a strongly consistent object storage zonal cache. It keeps data close to your accelerators, within a zone, in any region, with up to 70% lower read storage latency for training time.
- **Metric:** Rapid Storage delivers 5x lower latency for random reads and writes compared to other hyperscalers.
- **Customer:** Anthropic leverages Anywhere Cache to improve the resilience of their cloud workload by co-locating data with TPUs in a single zone and providing dynamically scalable read throughput up to 2.5TB/s. They also use Storage Intelligence to gain deep insight into their 85+ billion objects allowing them to optimize their storage infrastructure.

Axion processors: New database integrations

Preview

AlloyDB integration

Preview

Cloud SQL integration

Last year we announced the Axion processor, our first custom Arm®-based CPUs.

Since then Axion has solidified its position as the undisputed champion of cost-effective, high-performance cloud computing.

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

Axion proves its superiority in its first year:


Today, we are excited to share with you the latest updates on Axion, bringing further value to our customers and providing more ways to leverage the power of Axion:

The C4A VM family now supports Cloud SQL and AlloyDB for PostgreSQL managed databases, providing significant price-performance advantages for database workloads. To utilize Axion processors today, you can now choose to host your database on a C4A VM from directly within the console.

Noteworthy customers / metrics:

When running on C4A instances, AlloyDB and Cloud SQL, provide nearly 50% better price-performance compared to Compute Engine N series machines for transactional workloads, and up to 2x better transactional throughput performance compared to equivalent Amazon's Graviton 4 offerings.

Learn more: [Blog](#), [website](#)



Compute: New compute innovation for enterprise workloads

Preview

C4D, Z3H, H4D, M4

GA

Hyperdisk Instant Snapshots

At Next '25, we are announcing a number of Compute enhancements that are focused on optimizing for memory, storage, lowering costs, or maximizing performance.

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Compute updates:

C4D VMs - Consistently high performance and controlled maintenance with C4D VMs.

Z3 Family upgrades - ideal for storage-intensive workloads, hosting hypervisors, and hybrid cloud platforms.

H4D VMs - Unlocking scalable HPC workloads for faster results.

M4 VMs - double performance for demanding SAP workloads. Backed by Compute Engine's memory-optimized 99.95% single instance SLA, M4 VMs offer up to 65% better price-performance and 2.25x more SAP Application Performance Standard (SAPS) compared to our previous memory optimized M3.

New Hyperdisk functionality - Customers can now store up to 5 PiB of data in a single pool - a 5X increase from before. We're also launching **Hyperdisk Instant Snapshots**, which deliver near-instantaneous, high-frequency, point-in-time checkpoints of disks that can be rapidly restored as needed.

18 additional VMware Engine node shapes for better capacity shaping and TCO optimization, VMware Engine offers 6 times more node shapes than competitors.

Noteworthy metrics:

C4D delivers big performance gains over previous VM families: 30%+ over C3D and 57% over N2D, helping customers optimize their TCO and meet the needs of mission-critical applications.

H4D VMs boost VM performance by 75% on CPU bound HPL benchmarks and 35% on memory bandwidth bound STREAM Triad vs C3D.

Learn more: [Blog](#), [website](#)

What's new with GKE: AI innovation

GA

Cluster Director

Preview

Gemini Cloud Assist Investigations

GKE AI Inference Recommendations

GKE Inference Gateway

Pre-announcement

GKE Autopilot for standard GKE clusters

RayTurbo on GKE

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

GKE improvements to help platform teams succeed with AI:

- Coming in Q3 2025, **GKE Autopilot's dynamic compute platform** will be available for standard GKE clusters, without having to use a specific cluster configuration.
- **Gemini Cloud Assist Investigations**, now in private preview, helps with GKE troubleshooting, decreasing the time it takes to understand, root cause, and resolve issues.
- A new **GKE AI Inference Recommendation experience**, now in public preview, uses Gemini Cloud Assist to configure infrastructure and resources following performance and cost optimization best practices.
- **GKE Inference Gateway**, now in public preview, provides intelligent routing and load balancing for AI inference on GKE.
- Built in partnership with Anyscale, **RayTurbo on GKE** will launch later this year to deliver superior GPU/TPU performance, rapid cluster startup, and robust autoscaling.
- **Cluster Director**, now generally available, lets you deploy and manage large clusters of accelerated VMs with compute, storage, and networking — all operating as a single unit.

Noteworthy customers / metrics:

- **Metric: 30% of all active GKE clusters created in 2024 were created in Autopilot mode**, with customers like Toyota and Contextual AI leveraging GKE Autopilot for their critical workloads.
- **Metric: GKE Inference Gateway**, now in public preview, **reduces serving costs by over 30%, tail latency by 60%, and increases throughput by up to 40%** compared to other managed and open source Kubernetes offerings.
- **Metric: With RayTurbo an optimized version of open-source Ray, achieve 4.5x faster data processing and that requires 50% fewer nodes for serving.**

Learn more: [Blog](#), [website](#)

AI Models, Agents, & Platforms



AI Model Innovation



Gemini: What is new



Gemini 2.5 Pro is state-of-the-art on a wide range of benchmarks and **debuts at #1 on LMArena**.



Since Next 2024, we've seen over **40X growth** in Gemini use on Vertex AI, now with billions of API calls per month.



More than **4 million** developers are using Gemini.

Learn more:

- [Launch blog](#), [Developer blog](#), [Gemini Website](#)
- Try [Gemini 2.5 on Vertex AI](#)



Use Gemini 2.5, our most intelligent AI models yet, to gain a competitive edge

Gemini 2.5 are "thinking models" with advanced reasoning. They show thoughts step-by-step for transparency, tackling complex problems and enabling smarter, context-aware agents with optimized performance.

Benefits of thinking models:



Gemini 2.5 allows organizations to **build agents faster, improve coding, streamline operations**, and more.



Thinking models offer the potential to **solve more complex business challenges, automate sophisticated tasks** requiring reasoning and planning, and **gain deeper, more actionable insights from their data**.

What's new?

Vertex AI is expanding Gemini capabilities:



Gemini 2.5 Pro (Preview): SOTA model for complex reasoning, coding, math now available on Vertex AI.



Gemini 2.5 Flash (Coming Soon): Fast, efficient thinking model, optimized for speed/cost.



Vertex AI Model Optimizer (Experimental): Leverages Google's unique understanding of Gemini to automatically direct queries to the most performant model and tools, based on our customers' quality, speed, and cost preferences.



Live API (Preview): Production-ready API for real-time, interactive audio experiences with enhanced features.

Generative media models

With our [generative media models in Vertex AI](#), enterprises can build a complete, production ready asset starting from a text prompt, to an image to a complete video asset with music and speech.

Did you know?

Vertex AI is the **only platform** with generative media models across all modalities.

Introducing Lyria

Text-to-Music

What's new?

New text-to-music capability allowing customers to create 30 second instrumental sound clips in seconds on Vertex AI.

[Read more](#)

Veo 2

Text-to-video & Image-to-video

What's new?

- **New editing features** for Veo 2, that help you refine and repurpose video content with precision.
- **Camera presets** like pan-left, pan-right, and drone style shots
- **Interpolation**: define first and last frame of video and Veo will fill in
- **Video inpainting and out painting**

[Read more](#)

Chirp 3

Audio generation & understanding model

What's new?

- **HD Voices** for text-to-speech in 30+languages. [Try our latest Chirp 3 HD voices](#)
- **Instant Custom Voice**: a new way to create custom voices with just ten seconds of one input
- **Transcription with Diarization** accurately separates and identifies individual speakers in multi-speaker recordings, significantly improving the clarity and usability of transcriptions for applications like meeting summaries, podcast analysis, and multi-party call recordings.

[Read more](#)

Imagen 3

Text-to-image

What's new?

Over the last year we've made huge improvements to Imagen 3, our highest quality text-to-image model, capable of generating images with even better detail, richer lighting and fewer distracting artifacts than our previous models.

[Read more](#)

Veo 2:

What is new

GA

Veo 2 Text to Video Generation

Veo 2 Image to Video Generation (GA with Allowlist)

Experimental

Advanced editing features including camera presets, interpolation, and video inpainting and out painting.

Ti;dr: Announcing a robust feature set that helps you create videos, edit them, and add visual effects with Veo 2. These features help teams edit and repurpose video content to meet their evolving needs, **transforming Veo on Vertex AI from a generation tool to a comprehensive video creation and editing platform.**

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Veo 2, your comprehensive video creation and editing platform, is expanding with 3 new editing capabilities:



Original video

With harness removed

Video inpainting and out painting

Refine and enhance existing video footage by adding / removing object and expanding the video frame

Camera presets- implement sophisticated cinematic techniques without requiring complex prompting or specialized expertise. For example, you can move the camera in different directions, create a timelapse effect, or generate a drone style shot.



Interpolated video

Interpolation - With interpolation, you can define the beginning and end of a video sequence, allowing Veo to seamlessly generate the connecting frames.

[Learn more on generative media launches](#)

Imagen 3:

What is new

GA

Imagen 3 - text to image

Imagen 3, our highest quality text-to-image model, now has **improved image generation and inpainting capabilities** for reconstructing missing or damaged portions of an image. Our latest update significantly elevates the quality of object removal, delivering a more natural and seamless editing experience.

[Learn more on generative media launches](#)
[Learn more about Imagen 3](#)

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

New improvements to Imagen 3 quality and editing

Images generated with **even better detail, richer lighting and fewer distracting artifacts** than our previous models.



Imagen 3 Editing provides a powerful and user-friendly way to refine and tailor any image.

Imagen 3 inpainting capabilities for reconstructing missing or damaged portions of an image. Our latest update significantly elevates the quality of object removal, delivering a more natural and seamless editing experience. Now you can quickly remove unwanted objects, blemishes, or distractions from your photos.

Creative workflow example



Original image



*Imagen 3 Editing:
background change*



*Imagen 3 Editing:
inpaint (object removal)*



*Imagen 3 Editing:
outpaint*

Chirp 3

GA

Instant Custom Voice (GA with Allowlist)

Preview

Transcription with Diarization

Our groundbreaking audio generation and understanding model, Chirp 3, now includes two new features:

Chirp 3: Instant Custom Voice: generate realistic custom voices from 10 seconds of audio input. Use this to personalize call centers, develop accessible content, and establish unique brand voices. To ensure responsible use, Instant Custom Voice includes built-in safety features and usage permissions.

Chirp 3: Transcription with Diarization: accurately separates and identifies individual speakers in multi-speaker recordings, significantly improving the clarity and usability of transcriptions for applications like meeting summaries, podcast analysis, and multi-party call recordings.

Learn more: [Chirp 3 Documentation](#)

Lyria

Text-to-music model

Private Preview

Introducing Lyria

Lyria, produces high-fidelity audio, meticulously capturing subtle nuances and delivering rich, detailed compositions across a range of musical genres. Lyria on Vertex AI can help enterprises:

Elevate brand experiences: Quickly create soundtracks for marketing campaigns, product launches, or immersive in-store experiences, all tailored to your brand's unique identity.

Streamline content creation: For video production, podcasting, and digital content creation, finding the perfect royalty-free music can be a time-consuming and costly process.



Listen to Lyria in action

Prompt: Craft a high-octane bebop tune. Prioritize dizzying saxophone and trumpet solos, trading complex phrases at lightning speed. The piano should provide percussive, chordal accompaniment, with walking bass and rapid-fire drums driving the frenetic energy. The tone should be exhilarating, and intense. Capture the feeling of a late-night, smoky jazz club, showcasing virtuosity and improvisation. The listener should not be able to sit still.

[Learn more on generative media launches](#)

AI Agents & Platforms



Vertex AI

New ways to build and scale multi-agent systems in Vertex AI Agent Builder



Vertex AI Agent Engine

Deploying AI agents with enterprise confidence

- **Deploy agents built using any agentic framework**, with enterprise-grade controls for seamless and compliant scalability.
- **Enhance agent intelligence**: Equip agents with memory for persistent sessions and leverage comprehensive tools to measure and continuously improve performance.
- **Drive enterprise adoption at scale** with built-in integration to Google Agentspace.



Agent Development Kit (ADK)

Build customizable agents with our open source framework

- Working together we packaged up **agent building best practices** so you can get a sophisticated multi-agent application up and running in **under 100 lines of intuitive code**.
- Leverage familiar **Python** and extensive documentation.
- **Maintain full control and observability** over your agent logic, with a high degree of customizability.



Agent Garden

Discover pre-built agents to jumpstart development

- **Accelerate development** with curated collection of [collection](#) of pre-built agent samples, solutions, tools, and frameworks.
- Access expert guidance and **best practices for building effective AI solutions**.
- **Build agents faster and easier** with proven building blocks.



Agent2Agent Open Protocol

Connecting agents across your enterprise ecosystem

- **Gives your agents a common, open language to collaborate** – no matter which framework or model they are built on.
- **50+ partners** are committed to working with us on advancing our shared vision of multi-agent systems.
- **Eliminates agent silos** by facilitating efficient communication and data exchange between diverse agents and systems.



Grounding in Google Maps

You can now augment your AI agents by grounding them in relevant location context from Google Maps data (which has +100M updates every day) to create fresh and factual responses.

Vertex AI is the most open and comprehensive AI platform available, with a commitment to openness and interoperability, a robust feature set and strong partner ecosystem.

Vertex AI usage increased

↑ **20x** 

in the last year



Learn more:

[Details on new Vertex AI launches](#)
[Get started in the Vertex AI console](#)
[Explore our documentation](#)

Google Agentspace

Helps organizations put **enterprise search** and **AI Agents in the hands of every employee**; with new agents and capabilities to help enterprises scale AI Agent adoption.



Agents gallery

Lets employees adopt agents at scale

Agents gallery provides a **single view where employees can browse and adopt agents** available in their organization, including Google-built agents like NotebookLM, Idea Generation agent and deep research agent, partner agents via Google Cloud Marketplace, and agents they create.



Support for partner agents

Growing variety of agents to choose from

We are growing the AI Agent Marketplace, where customers can easily browse and purchase AI agents from partners such as Accenture, Deloitte, and more, and make them easily available within Google Agentspace.



Chrome Enterprise integration

Google Agentspace unified search is now available from the Chrome Enterprise search bar, letting employees leverage multimodal agentic intelligence in a space they already work: the Chrome browser.



Agent Designer

Lets employees create custom agents with a no-code builder

With **Agent Designer**, employees can automate business processes and tasks by using a simple chat-based interface and multiple pre-built tools and actions. No coding is required and agents support the open Agent2Agent protocol, ensuring interoperability with not only data sources connected via Google Agentspace, but also agents built on other platforms.



2 New Google-built expert agents

Gives employees new ways to work with data and drive impact


Deep research agent leverages Gemini's multimodal intelligence to help employees turn disjointed data into rich, actionable insights and reports.

Idea Generation agent helps employees to generate and develop impactful ideas. Inspired by the scientific method, it generates ideas, tests them against one another, and ranks top performers.

Learn more: [Launch blog](#), [Website](#)

Ready to build your agentic future?

Google Cloud streamlines the entire agent lifecycle, from build to adopt with enterprise-wide access, helping you get to ROI faster.



Google Agentspace is a platform that empowers all employees to **use** AI agents to access enterprise knowledge and automate work tasks with governance.

- **User:** All employees, regardless of technical level
- **Build:** no-code agents (Agent Designer)
- **Use:** pre-built agents (NotebookLM, Deep Research, Idea Generation), agents in Agents gallery, custom agents that are built in Vertex AI and available for adoption in Google Agentspace

Vertex AI is the platform to **build and deploy** AI agents, with the right models, and least effort of managing the underlying AI infrastructure.

- **User:** Developers, Data Scientists & ML Engineers
- **Build:** custom agents (Agent Builder, Agent Engine Agent Garden), Open Source (ADK and Agent2Agent Protocol)

Customer Engagement Suite: What is new

Launch stage:

GA

Customer Engagement Suite

- Conversational Agents console
- Prebuilt agents / 3P connectors

Preview

Customer Engagement Suite

- AI Coach / AI Trainer

Pre-announcement

Next gen Conversational Agents

TL;DR:
New capabilities that make **building AI agents easier**, customer service representatives **be more productive**, and **streamline** contact center operations.

Customer Engagement Suite is expanding capabilities to yield even **stronger results:**

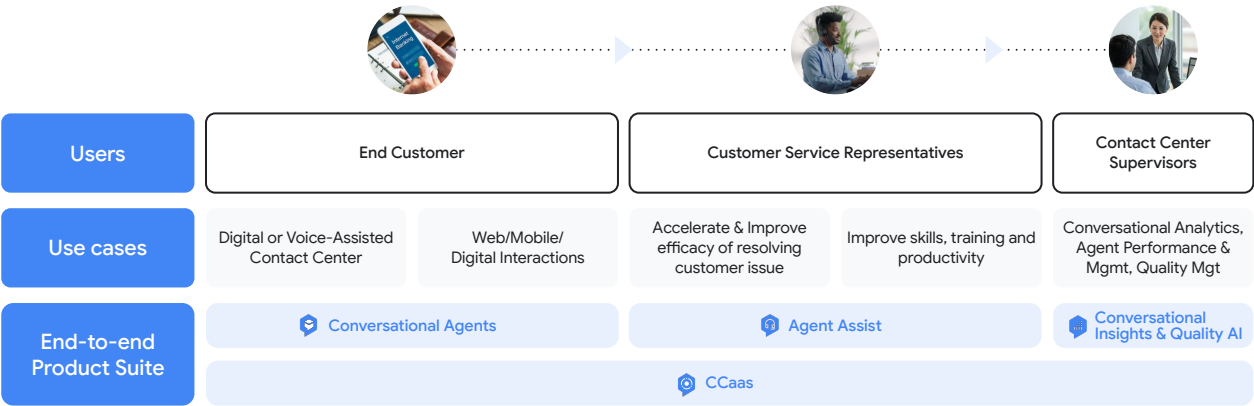
- **Conversational Agents console:** Enables users to build AI agents with both deterministic and generative AI functionality
- **Prebuilt agents:** For common commerce use cases
- **3P connectors:** To most common telephony, data source and CRM systems
- **AI Coach:** Guides customer service reps through complex interactions and upsell activities in realtime
- **AI Trainer:** Provides customer service reps with scenario simulation and personalized training to improve performance.
- **Next gen Conversational Agents:** Human-like voice, streaming video support, and simplified conversational agent building experience

Customer success:

Verizon: 25% faster customer service call resolution

TTEC: 40% reduction in call center escalations

[Read more here](#)



Learn more: [Website](#)

Developer

Firestore: Build full-stack AI apps, all in one place

GA

Firestore App Hosting

Firestore Data Connect

Genkit for Node.js version 1.0

Preview

Firestore Studio & Gemini Code Assist Agents

App Testing Agent

Crashlytics Alerting Revamp

Vertex AI in Firestore: Gemini 2.0 Multimodal Live API support

Genkit support for Python and Go

Introducing transformational capabilities to the Firestore Platform that help professional developers forge new paths for building AI applications across multiple platforms.

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

Driving development innovation:

Firestore Studio [Preview]: Cloud-based, agentic development environment that includes everything full-stack developers need to create production-quality, AI-powered apps quickly and safely, all in one place. Access Gemini Code Assist agents via Firestore Studio [Preview]

Integrating AI capabilities into existing applications:

- **Vertex AI in Firestore:** Lets developers integrate generative AI into their applications by providing a streamlined, serverless, and secure SDK to access models in Vertex AI from Firestore. Announcements include **Gemini 2.0 Multimodal Live API** [Preview] and **Imagen support** [Preview]
- **Genkit [UPDATES]:** Expanded language support for Genkit, which helps reduce the complexity of building, testing, and monitoring your apps' AI, including **Genkit Node.js** [GA], **Genkit Python** [early support] and **Genkit Go** [expanded support]

Optimize development velocity across the entire application lifecycle:

- **Firestore App Hosting** [GA]: App Hosting manages the entire stack for your app– from the build pipeline, to the CDN, to server-side rendering– so that you can get-to-market faster.
- **Firestore Data Connect** [GA]: Developers can now leverage the full power of a managed PostgreSQL database, seamlessly integrated with Firestore, to build sophisticated and scalable applications.
- **App Testing Agent** [Preview]: AI Testing agent in the Firestore Console, which developers can guide to author and run tests using natural language.
- Get started in Firestore faster: **Integrated sample app experiences** and **Gemini-assisted onboarding**.

Gemini Code Assist: What is new

Preview

Gemini Code Assist agents

Gemini Code Assist tools

Coding assistance in Android Studio with Gemini Code Assist Enterprise

Gemini 2.5 is now powering Gemini Code Assist for individuals to use in personal projects

Agents carry out tasks on behalf of developers to help them create quality faster and easier.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Gemini Code Assist is introducing agents and offering more surfaces for coding assistance:

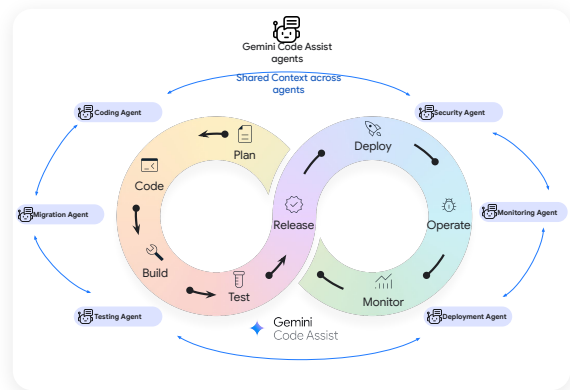
- Code Assist agents to help with the following tasks:
 - Generate new software: Create applications from product specifications.
 - Migrate code: Code transformation and code translation.
 - Implement new features: Read issues assigned to users in GitHub and create code to address the issue.
 - Execute code reviews: Analysis and feedback on quality, best practices, and organizational style.
 - Generate unit tests: Create test cases, execute tests, and report results.
- Coding assistance in Android Studio with enterprise-grade security and privacy available with purchase of Gemini Code Assist Enterprise
- Gemini Code Assist tools: bringing partner technology into the IDE with Snyk, Atlassian, and more

Noteworthy customers / metrics

- “CME Group, which runs the Chicago Mercantile Exchange, states that Code Assist provides most of its developers more than 10.5 hours of productivity gain per month.”
- Customers: Ampere from Renault Group, Broadcom, CME Group, CGI, PayPal, and Wipro

Key Stats

- “Developers are estimated to get common development tasks done over 20% faster with Gemini Code Assist.”
- “Using Gemini Code Assist resulted in a nearly 20 percentage point increase in the task success rate for common development tasks.”
- “Using Gemini Code Assist is estimated to make documenting code over 20% faster.”
- “Using Gemini Code Assist is estimated to help developers create unit tests over 30% faster.”
- “Using Gemini Code Assist resulted in over an estimated 10 percentage point increase in the task success rate for migrating legacy Java code to a modern Java version.”



Learn more: [Blog](#), [website](#), [Session 1](#), [Session 2](#), [Session 3](#), [Session 4](#), [Session 5](#), [Session 6](#)

- Application Design Center Integration (Public Preview)
- Investigations (Private Preview)
- Cost and utilization optimization (Private Preview)
- Enhanced chat experience (Public Preview)
- Deeper integrations with Google Cloud products (Public Preview)

Google Cloud Next

- **What's new 1:** Application infrastructure design and deployment with natural language
- **What's new 2:** Gemini-guided investigations workflow for fast troubleshooting
- **What's new 3:** Cost and utilization optimization with AI assistance
- **What's new 4:** Enhanced Gemini chat with richer context and ability to conduct domain specific tasks
- **What's new 5:** Deeper integrations across many Google cloud products including Storage Insights, Cloud Observability, Firebase, Database Center, Flow Analyzer, FinOps Hub, as well as our security- and compliance-related services.

- Blog, website



Cloud Hub: What is new

Preview

Deployments dashboard (Public Preview)

Health and troubleshooting dashboard (Public Preview)

Optimization dashboard (Private Preview)

Quotas and reservations dashboard (Public Preview)

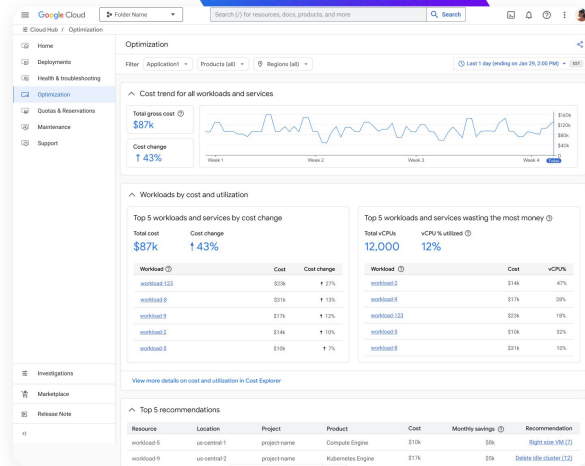
Maintenance dashboard (Public Preview)

Support cases dashboard (Public Preview)

Cloud Hub serves as the central command center for your entire application and project landscape on Google Cloud.

What's new?

- **Deployment dashboard** where you can check your application deployment status and application templates
- **Health and troubleshooting dashboard** which shows the health of the deployments.
- **Optimization dashboard** shows the cost and resource utilization of the applications and its components
- **Quotas and reservations dashboard**
- **Maintenance dashboard**
- **Support cases dashboard** shows your support cases and their status



Learn more:

- [Blog](#), [website](#), [Next Session](#)

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Cloud Run: Run your app on a fully managed platform

GA

Cloud Run GPU

Cloud Run functions

Preview

Vertex AI Studio to Cloud Run

GPU for Cloud Run jobs

High availability for Cloud Run

Cloud Run threat detection

Built-in Identity-Aware Proxy

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Cloud Run is the perfect runtime for AI app development and is capable of running workloads no other serverless platforms can.

AI with Cloud Run: You can run AI frameworks such as Langchain or Firebase Genkit, to orchestrate calls to AI models from Vertex AI, vector databases, and any other APIs. At Next25, we are announcing that you can deploy your gen AI app in a click from Vertex AI Studio to Cloud Run.

Cloud Functions is now Cloud Run functions: Cloud Run functions is now generally available. You can run code, functions or containers with Cloud Run.

Hyper scalable runtime: You can run large variety of workloads including AI inference apps and build highly available architecture leveraging multi-region deployments.

Time-to-First-Token w/ Coldstart

Model	TTFT w/ Coldstart
gemma2:2b	10s
gemma3:4b	19s
deepseek-r1:7b	12s
deepseek-r1:14b	18s
llama3.1:8b	17s

Framework: Ollama
Quantization: 4-bit

Note that these numbers are observed in a controlled lab environment and actual performance numbers may vary depending on a variety of factors.

Coldstart from 0 to an instance with a GPU and drivers installed in...

5s

Learn more: [Next session 1](#), [Next session 2](#)

Google Developer Program: Enhanced benefits and Enterprise tier

GA

Google Developer Program Premium tier with new benefits

Preview

Google Developer Program Enterprise tier

Introducing an upgraded [Google Developer Program](#) that encourages individuals and teams to learn, build and grow with Google

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

AI Enhanced Premium Benefits [GA]

We've integrated Google's latest AI tools directly into the Google Developer Program premium membership, making it your catalyst for growth with advanced tools and dedicated support for every phase of your development journey.

- **Gemini Code Assist Standard:** Access paid Gemini in Firebase features and Gemini Code Assist
- **Expanded Firebase Studio Workspaces:** Manage up to 20 Firebase Studio workspaces
- **\$50 Gen AI Developer Credit:** Exclusive use for Google AI Studio and Google Cloud Vertex AI
- **Google One AI Premium:** 3 month trial to enjoy Gemini Advanced, NotebookLM Plus, and much more

Google Developer Program Enterprise [Preview]

For organizations looking to empower their development teams, the Google Developer Program now has an Enterprise offering, a safe and affordable way to explore Google Cloud and its AI capabilities. This offering is a company managed subscription.

- **Google Cloud Developer Sandboxes:** A risk-managed environment with flexible access to the full range of Google Cloud products, complete with guardrail controls and cost management set by your organization. It includes \$150 USD in Google Cloud credit per developer per month.
- **Gemini Code Assist Enterprise:** A comprehensive AI-powered application development solution, customizable to your private source code repositories, and integrated with many Google Cloud services.
- **Unlimited Google Cloud Skills Boost:** Get unlimited access to over 700 on-demand, hands-on labs and skills badges on Google Skills Boost.

Learn more: [Next session 1](#), [Next session 2](#)

Databases & Data Analytics

The Autonomous Data to AI Platform

Data Analytics



BigQuery platform

GA

BigQuery metastore

Preview

BigQuery multimodal tables

Gemini in BigQuery (new features in preview)

Pre-announcement

BigQuery AI Engine

TL;dr: At Next '25, we launched BigQuery's next evolution: the autonomous data to AI platform to provide customers with a flexible, integrated, and intelligent platform with automation and agentic capabilities for the entire data lifecycle and for all data teams, including data engineers, scientists, and analysts.

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

BigQuery is the autonomous data to AI platform.

What's new?

- **Gemini in BigQuery (new features in Preview)** innovations for data prep, code assist, translation and migration, metadata curation, and data insights
- **BigQuery multimodal tables (Preview)** to unlock new avenues for such feature engineering and model development.
- **AI Query Engine (Pre announce)**, entering preview, gives you the ability to incorporate language models directly into SQL queries to simplify semantic feature extraction, sentiment analysis
- **BigQuery metastore (GA)** simplifies data discovery and querying for external engines, mirroring the open-source experience.
- **BigQuery workload management (GA)** to automatically scales resources, manage workloads, and ensure cost-effectiveness
- **BigQuery Spend Commit (GA)**, providing predictable costs

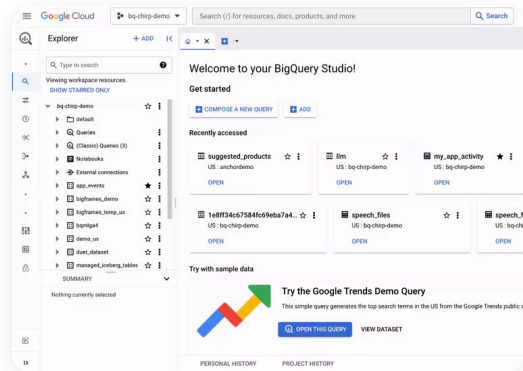


BigQuery has
5x more customers
than any other independent
cloud data warehouse and
data science solution



BigQuery offers up to
2.5x faster performance and
up to
20x more cost efficiency
with BigQuery and Vertex AI.

Learn more: [Blog](#), [Website](#), [Next session 1](#), [Next session 2](#)



We didn't just need a place to store or consume data, we wanted a collaborator that could help us scale the most advanced data management in the industry."

Jaime Montemayor,
Chief Digital & Technology
Officer, **General Mills**

Looker

GA

Studio in Looker (Next)

Preview

Gemini in Looker (Next)

Conversational Analytics API (Next)

At Next '25, we are launching updates to Looker, our AI for BI solution, delivering conversational analytics allowing customers to interact with their data more easily and intuitively, and adding Looker Reports via Studio in Looker.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

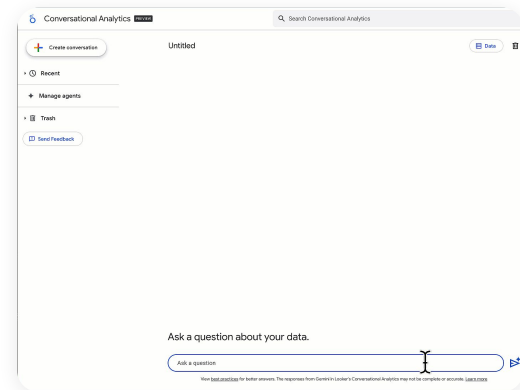
Looker is expanding capabilities for:

- **Conversational Analytics API (Preview)**, enables developers to seamlessly integrate natural language queries into any application or workflow.
- **Gemini in Looker (Preview)** includes visualization, formula, and automated slide generation as well as LookML Code Assistants to simplify the code creation process
- **Looker Reports (GA)**: provides enhanced data storytelling, streamlined exploration, and broader data connectivity to Looker customers via Studio in Looker unification. Includes enterprise reportings as 1st party content, direct connection to Microsoft Excel, and ad-hoc data source access.

Noteworthy customers / metrics

- Looker's semantic layer reduces data errors in gen AI natural language queries by as much as two thirds.
- Unilever processes 75,000 orders daily (€2.5 billion annually) with their BigQuery and Looker powered agent solution for distributed trade transformation.

Learn more: [Blog](#), [Website](#), [Next session 1](#), [Next session 2](#)



Gemini in Looker conversational analytics

Databases



Databases: Transform your business with AI

Preview

New gen AI capabilities in AlloyDB

Firestore with MongoDB compatibility

Arm support for Cloud SQL and AlloyDB

Expanded Oracle & Google Cloud services

SQL Server modernization

At Next '25, we announced Google Cloud database capabilities that help you to build gen AI applications and agents, enjoy freedom of choice as a developer, and migrate and modernize your database estate.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

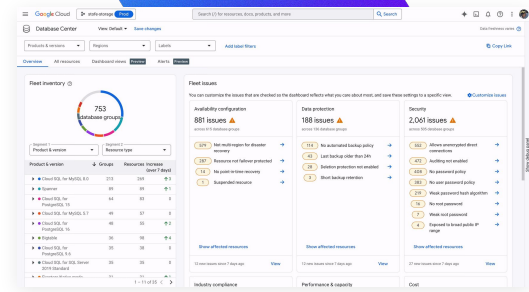
Key database product announcements:

- **New gen AI capabilities in AlloyDB:** Google Agentspace integration, natural language support, optimized SQL and vector search, 3 new gen AI models on Vertex AI, and an AI query engine.
- **Firestore with MongoDB compatibility:** Firestore developers can now take advantage of MongoDB's API portability along with Firestore's industry-leading capabilities.
- **Arm support for Cloud SQL and AlloyDB:** C4A instances are optimized for data-intensive applications and are built on Google Axion processors, delivering superior price-performance and transactional throughput performance.
- **Expanded Oracle and Google Cloud services:** Oracle Base Database Service, a flexible, controllable way to run Oracle databases in the cloud; global Exadata X11M availability; and 11+ new regions planned in 2025, for a total of 15+ regions.
- **SQL Server modernization:** Database Migration Service supports SQL Server to PostgreSQL migrations for Cloud SQL and AlloyDB, featuring low-downtime data migration and a rules-based conversion engine augmented by a specially-trained Gemini model.

Noteworthy customers / metrics

- Customers highlighted at Next include Nuro, Bayer, Causal, Cynergy Bank, Character.AI, Target, Walmart, Mercari, Macy's
- Vector search adoption in AlloyDB increased nearly 7x in 2024.
- Compared to the HNSW index in standard PostgreSQL, AlloyDB AI's ScaNN index delivers up to 4x faster vector search queries and up to 10x faster filtered vector search queries.

Learn more: [Blog](#), [website](#), [Next session 1](#), [Next session 2](#)
Free trials: Try [Cloud SQL](#), [AlloyDB](#), or [Spanner](#)



Databases: AlloyDB AI innovations

Preview

Integration with Google Agentspace

AlloyDB AI natural language

Optimized SQL and vector search

3 new models in AlloyDB

AlloyDB AI query engine

At Next '25, we announced several new capabilities to make it easier to build intelligent applications and agents.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Key AlloyDB AI announcements:

- **Enable Google Agentspace to search structured data:** Query and search data in AlloyDB with Gemini's advanced reasoning & Google-quality search
- **AlloyDB natural language:** Query data securely & accurately with supplied context & interactive intent clarification
- **Optimized SQL and vector search:** Span vector search & structured filters and joins, powered by AlloyDB analytics accelerator
- **3 new models in AlloyDB:** Cross-attention based reranker model, multimodal model to support text, images, and videos, and new Gemini embedding model
- **AI Query Engine:** enables developers to freely use natural language expressions and constructs right within their SQL queries.

Noteworthy metrics

- Customers highlighted at Next include Target, Linear, Intelligentsia AI, Fynd, Regnology, Netrise, Neuropace, Opentext, Lightricks, Nuro
- Compared to the HNSW index in standard PostgreSQL, AlloyDB AI's ScaNN index delivers up to 10x faster index creation, up to 4x faster vector search queries, up to 10x faster filtered vector search queries, and typically uses 3x less memory.

Learn more: [Blog](#), [website](#)



AlloyDB AI

Vector processing

Model management

Natural language support

Security

Make Google part of your security team

Google Unified Security: Superior security outcomes

GA

Google Unified Security

Google brings together best-in-class security products for **threat intelligence**, **security operations**, **cloud security**, and **secure enterprise browsing**, along with **Mandiant Expertise solution** — into a converged security solution, powered by AI.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Google Unified Security is a converged security solution which brings integrated product experiences, exemplified by:

- Browser telemetry and asset context from Chrome Enterprise integrated into Google Security Operations to power threat detections and remediation actions
- Google Threat Intelligence integrated with security validation to proactively understand exposures and test security controls against the latest observed threat actor activity
- Cloud risks and exposures from Security Command Center, including those impacting AI workloads, enriched with integrated Google Threat Intelligence

Noteworthy customers / metrics

- “I see Google and its security suite as one of the top partnerships that I have within my organization.” - Craig McEwen, Deputy CISO, Unilever

Learn more : [Blog](#), [website page](#)



Security Agents & Gemini: Agentic AI powering security

Preview

Alert triage agent [2H]

Malware analysis agent [2H]

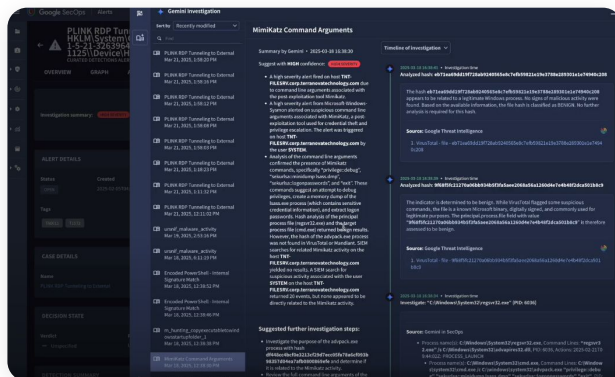
Agentic AI is powering a fundamental shift in how security operations are conducted. Our vision is a future where intelligent agents work alongside human analysts, offloading routine tasks, augmenting their decision-making, and freeing them to focus on complex issues.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

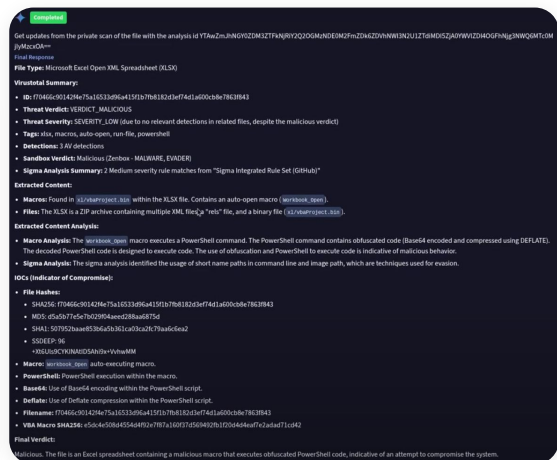
Google Cloud Next

Security Agents & Gemini is transforming security operations:

- **Alert triage agent:** Performs dynamic investigations of alerts and renders a verdict.
- **Malware analysis agent:** Autonomously identify security issues and recommend whether code is malicious or not.



Alert triage agent



Malware analysis agent

Google Security Operations: Advanced capabilities

GA

Data pipeline management

Mandiant Threat Defense

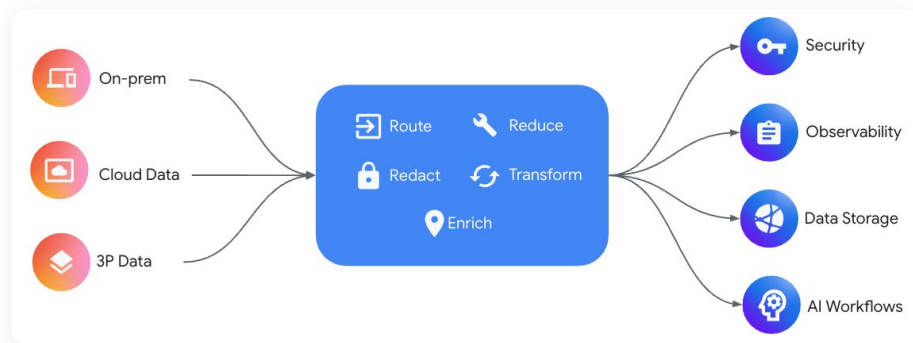
Customers can now leverage data pipeline management tools for improved scale, cost reduction, and compliance, along with the new Mandiant Threat Defense service to get expert-driven, AI-assisted threat detection and response.

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Google security operations data pipeline management and threat defense:

- **Data pipeline management:** Better manage scale, reduce costs, and satisfy compliance mandates.
- **Mandiant threat defense:** AI assisted threat detection & response with incident prioritization.



Noteworthy customers / metrics

- “The benefits of a purpose-built data pipeline tool are to reduce the data preparation necessary to interpret the streams of data and events specific to security insights.”
Allie Mellen, Forrester, [If You're Not Using Data Pipeline Management for Security and IT, You Need To](#)

Learn more: [Blog](#), [Next session](#)

Security Command Center:

Comprehensive security for Google Cloud

GA

AI protection

Preview

Compliance Experience [June 2025]

Pre-announcement

Data Security Posture Management [June 2025]

Security Command Center now includes AI Protection to secure AI systems and sensitive AI training data, along with a Compliance Hub to simplify regulatory management. These updates aim to provide comprehensive security and compliance visibility across Google Cloud

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

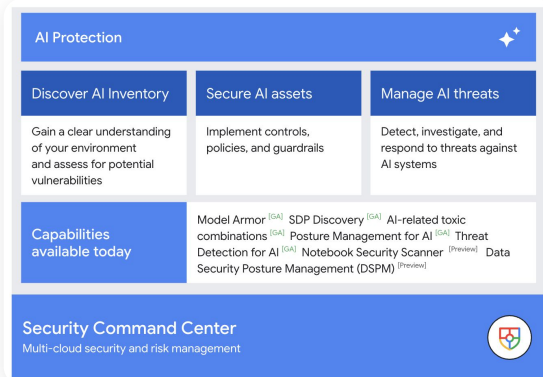
Enhancements integrate:

- **AI protection:** Manage risk across the AI lifecycle including Model Armor which applies content safety and security controls to prompts and responses
- **Data Security Posture Management (DSPM):** Enable discovery, security, governance, and monitoring of sensitive data including AI training data.
- **Compliance experience:** Combine policy definition, control configuration, enforcement, monitoring, and audit into a unified workflow.
- **Other enhancements:** New security risk dashboards and integration with Snyk's developer security platform are additional enhancements.
- **Risk Protection Program:** Expanding to offer discounted cyber-insurance, including AI and quantum-related coverage, through new partnerships with Beazley and Chubb

Noteworthy customers / metrics

- **Customers using Security Command Center:** Yahoo!, Deutsche Brose, Cardinal Health, Globo, and Ericsson

[Learn more: AI protection blog](#), [Secure your AI website page](#), [Unified Security blog](#), [Security Command Center website page](#)



Chrome Enterprise Premium: Enhanced security features

GA

Employee phishing protections

Data masking

Android Enhancements

Chrome Enterprise Premium enhances organizational security by providing advanced phishing protection, data masking for testers, expanded Android security features, and integration with Google's unified security platform

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

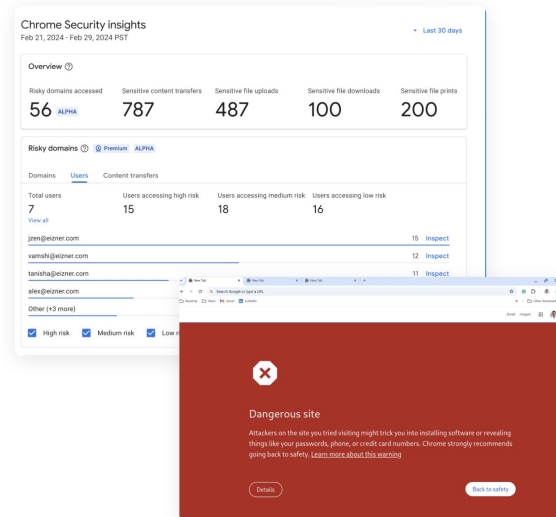
Chrome Enterprise Premium give organization more options to protect employees and data:

- **Employee phishing protections:** Protect employees against lookalike sites and portals.
- **Data masking:** Enhanced data masking capabilities
- **Expanded Mobile Security:** Key protections now on Android including copy/ paste controls and URL filtering
- **Google Unified Security:** Integration with Google Security's industry leading security solution to drive telemetry, safer browsing, and data analysis

Noteworthy customers / metrics

- **Metric**
- **Customers:** Telus, Snap

Learn more: [Protect your Business on the Web with Chrome Enterprise Premium](#)



Employee phishing protection

Mandiant Consulting: Added services

GA

Mandiant retainer

New partnerships with Rubrik and Cohesity

Mandiant now offers flexible incident response retainers and partners with Rubrik and Cohesity for faster, cheaper cloud recovery after cyberattacks.

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Mandiant CyberSecurity Consulting is improving services to deliver more flexibility and expertise:

- **Mandiant retainer:** On-demand access to Mandiant experts with pre-negotiated terms and two-hour incident response times.
- **New partnerships:** Mandiant Consulting is partnering with Rubrik and Cohesity on data backup and recovery to minimize downtime and recovery costs after a cyberattack



Trusted Cloud: New Security Controls & Capabilities

GA

Sensitive Data Protection

Network Security Integration

Preview

S3NS [Available Now]

Managed Workload Identities [Available Now]

HSM [Available Now]

SWP [Available Now]

Pre-announcement

Unified Access Policies

Confidential Computing

New security controls and capabilities on our cloud platform to help organizations meet evolving policy, compliance, and business objectives.


Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next


Trusted cloud helps organizations meet evolving policy, compliance, and business objectives:

- **Sovereign Cloud:** S3NS Trusted Cloud is the first sovereign cloud offering built on technology from a leading hyperscaler that is operated, majority-owned and fully controlled by a European organization.
- **Identity and Access Management:** Unified Access Policies create a single definition for IAM, while managed workload identities can provision a Secure Production Identity Framework For Everyone (SPIFFE).
- **Data security:** Expanding Confidential Computing, adding Sensitive Data Protection discovery services, and implementing a single tenant Cloud Hardware Security Module (HSM) for enhanced data security.
- **Network Security:** Additional enhancements available for Network Security Integration, DNS Armor, Cloud Armor Enterprise, Cloud NGFW Enterprise and Secure Web Proxy (SWP).

Learn more: [Next session 1](#), [Next session 2](#)



Workspace



Google Workspace Flows: **What is new**

Pre-announcement

Coming to preview in Q2

Build automations using 50+ available triggers and integrate with Gems and 3P connectors to automate tasks like approvals, email organization, and data collection.

Industries / use cases:

- All industries
- Both office workers and frontline workers

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

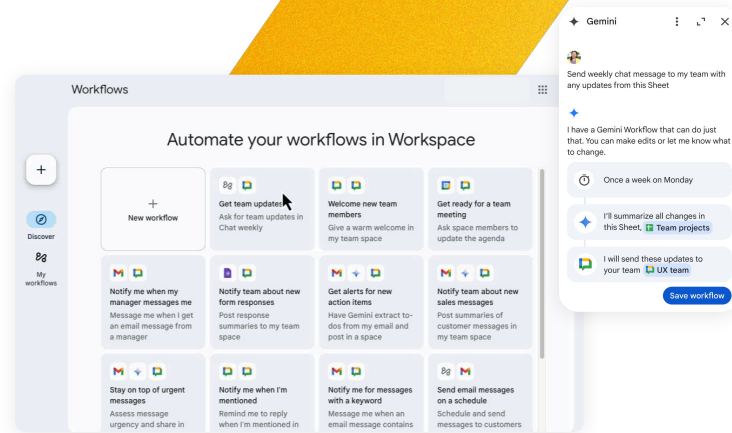
Google Cloud Next

Google Workspace Flows, accelerate your work with AI agent-powered workflows

- Create and run workflows that automate work across the Google Workspace suite
- Create agentic workflows using Gemini and Gems
- Connect to third party data and apps

Noteworthy customers / metrics

- Multiple studies have shown “workflows” as a key value driver for existing / potential customers



3P Data Connector framework

Pre-announcement

Coming to preview in Q2

New ways to integrate and connect other SaaS workloads with Google Workspace and its AI and workflow capabilities

Industries / use cases:

- All industries
- Both office workers and frontline workers

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

The new 3P Data Connector framework is coming to Workspace:

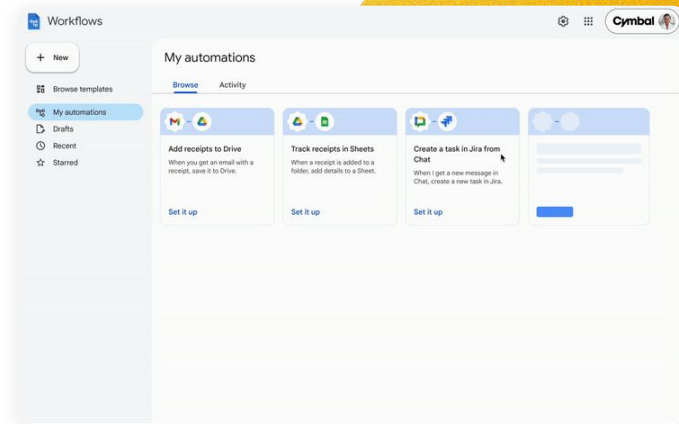
Easily build integrations and access to 3rd party data:

- Use in custom Gemini extensions
- Import data from third parties into Sheets as tables
- Incorporated into Google Workspace Flows
- Take third-party actions in Chat

Noteworthy customers / metrics

Multiple studies have shown “workflows” as a key value driver for existing / potential customers

Learn more: [Next Session](#), [Blog](#)



Sheets: AI function in Sheets

GA

Available in H2 2025

Preview

Currently available to Workspace Labs users

The AI function brings the power of Gemini directly into any and every cell in your spreadsheet. Input your specific prompt describing what you would like to do using simple language, and Gemini in Sheets will carry out your instruction.

Target industries / use cases:

- Ad copy and subject line development
- Managing product reviews and incorporating feedback
- Understanding user sentiment

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

Gemini in Sheets is expanding

Use the AI function to:

- Generate text that is tailored to your specific data
- Summarize information including important highlights
- Categorize data into meaningful groups so that you can detect patterns, including sentiment analysis



Learn more: [Website](#), [Blog](#)

Help me analyze: What is new

Pre-announcement

Coming in 2025

Help me analyze allows you to simply click to analyze data and Gemini will explore the data in Sheets and proactively surface key insights in an interactive experience - all without writing a single formula or building a chart or pivot table yourself.

Example enterprise use cases:

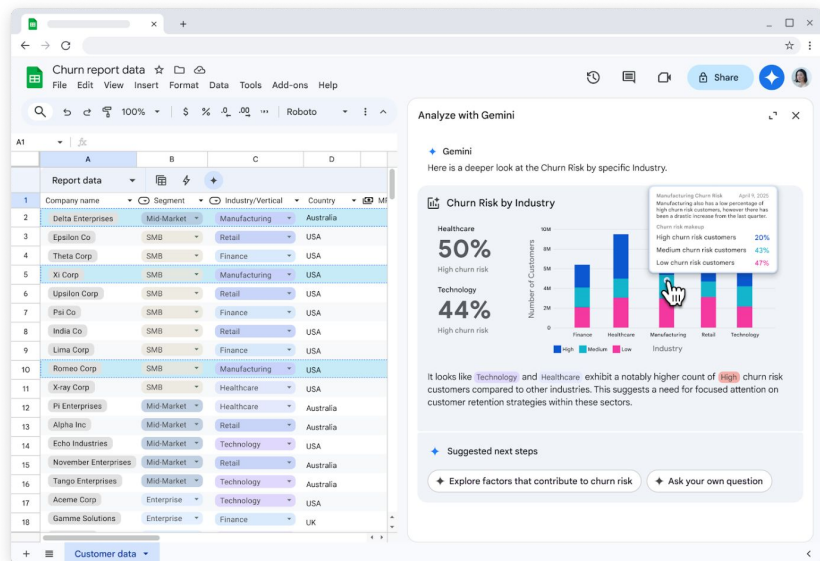
- Understand seasonal trends in consumer purchase behavior
- Identify areas to reduce production costs
- Identify reasons for slow in revenue growth

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Gemini in Sheets is expanding

- Gemini proactively surfaces insights
- Gemini proactively generates data visualizations including highlights for key data points
- Gemini will generate interactive data visuals



Learn more: [Blog](#)

Start a Chat from Gmail (aka Airmail):

What is new

Alpha

Public alpha end of Q2

Makes it easy to move real-time conversations out of your inbox (where they can be chaotic and hard to follow) and into Google Chat.

Industries / use cases:

- Knowledge workers
- Customer service & some frontline (e.g. case management)

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

Gmail and Chat are streamlining collaboration:

- **Start a chat from Gmail:** Move conversations from Gmail to Chat by clicking a button. The attendees and message content come with you.
- **Use the full power of Chat:** In Chat, teams have access to a rich set of collaboration tools and can add Gemini as a conversation partner
- **Threads are cross-connected:** Two-way context is preserved, with persistent cross-links from Gmail to Chat and vice versa.



Learn more: [Blog](#), [Next Session](#)

Gemini in conversations:

What is new

Preview

Labs by end of Q2

Provides tools to improve teamwork and drive alignment. Members of the conversation can create and share discussion summaries, lists of action items, and more in real-time with the entire team.

Industries / use cases:

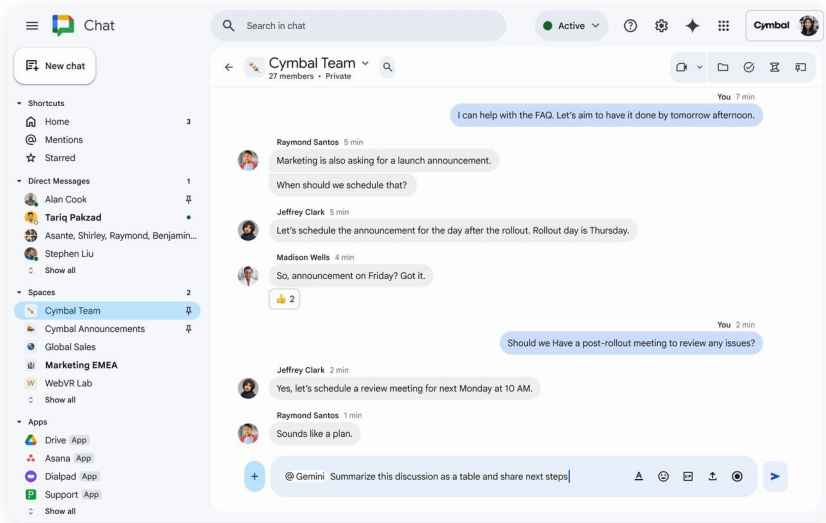
- Knowledge workers

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Gemini in Chat is becoming your team's collaboration partner:

- Action items:** Generate a detailed list of open actions with names & dates
- Refocus the conversation:** Provide a summary of the active conversation visible to the whole team
- Find key details:** Gemini can track down essential data from deep in the conversation and share with the group



Learn more: [Blog](#), [Next Session](#)

Unified Workspace apps: **What is new**

GA

End of Q2

Enables developers to deliver a single app that will work across Workspace, wherever add-ons are supported.

Industries / use cases:

- Knowledge workers
- Developers

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

With new Workspace apps:

- **Developers:** Can deliver a single app that will run anywhere Workspace add-ons are supported
- **Admins:** Can deploy a single app using centralized installation from the Workspace Admin console



NoteBot for Workspace

NoteBot Inc.

Easily and quickly take notes on all your tasks and lists and share them across Google Wo...

★ 4.4 ± 59M+



NoteBot for Google Chat

NoteBot Inc.

Easily and quickly take notes on all your tasks and lists and share them in Google Chat.

★ 4.6 ± 888K



Learn more: [Blog](#), [Next Session](#)

Vids: Generate a video

GA

Preview

Q2 '25

Q2 '25 [Alpha]

Pre-announcement

April 9, 2025

'Generate a video' helps you create custom, high-quality video clips with only a prompt, without ever leaving Workspace.

Enterprise use cases:

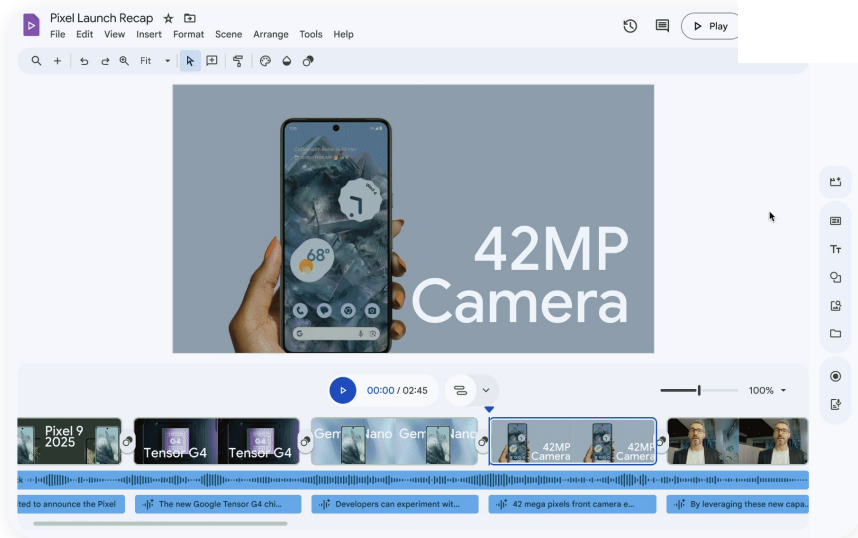
- **Industry Specific Training:** "A safety worker wearing a yellow hard hat on an oil rig while reading OSHA instructions for the job site."
- **Corporate Event Video Invite:** "Executive presenting about AI Innovation to a room full of clients in a modern setting. Single-shot, realistic 4K camera."
- **Internal recaps:** "Realistic shots of an airplane flying to Vegas"
- **Sales pitches:** e.g. "A photo realistic person in first person pushing a sled with 'Nike' written on the shorts"

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Vids expands video capabilities:

Generate custom video clips with Veo 2, Google's latest state of the art video generation model directly in Google Vids



Learn more: [Blog](#), [Google Vids product page](#), [Next Session 1](#), [Next Session 2](#)

Audio overviews in Docs: **What is new**

GA

H2 2025

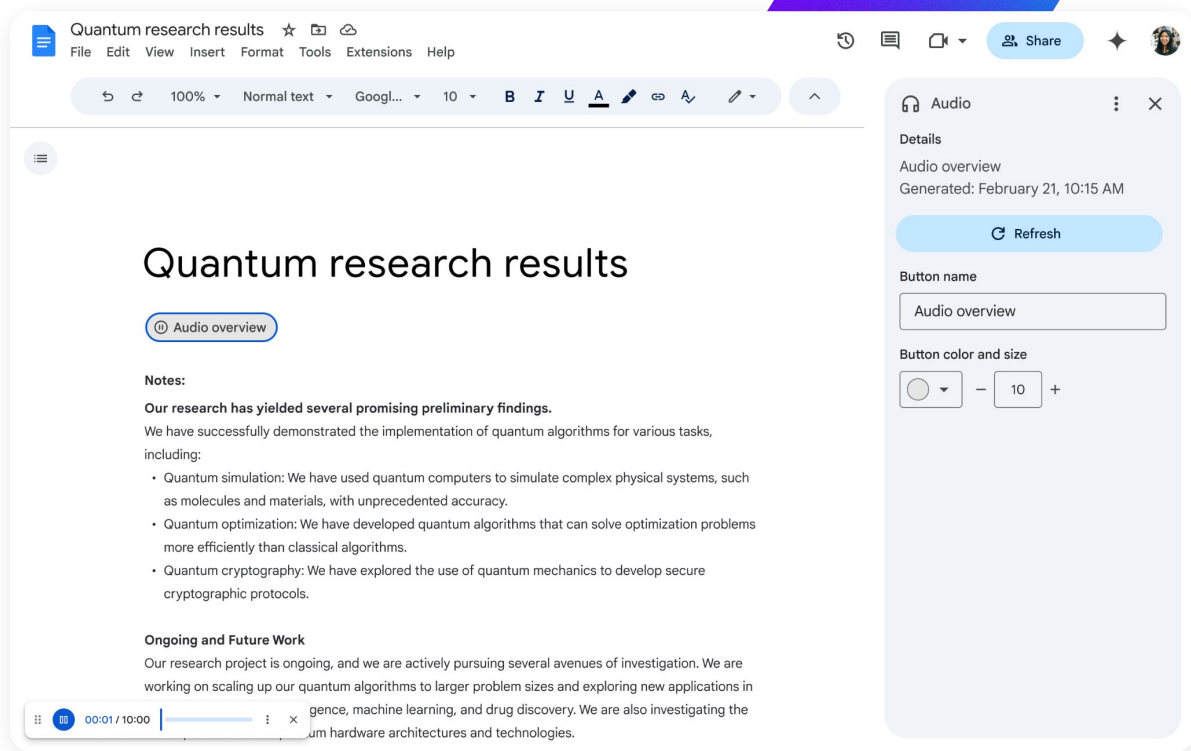
Audio overviews give users flexibility and choice in how they want to consume content. Whether you're multitasking or prefer to process information audibly, audio overviews enables you to consume content in your preferred format.

Example enterprise use cases:

- Get a summary of a Doc to hear the latest updates
- Listen to a long white paper I need to incorporate into my research; start and stop whenever I need to while staying engaged with natural voices

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next



Learn more: [Blog](#)

Help me refine in Docs: **What is new**

GA

June 2025

Help me refine goes way beyond spelling and grammar, delivering suggestions intelligently tailored to your document's goals. Help me refine proactively suggests meaningful edits to help strengthen arguments, to make your documents beautiful, and ultimately improve your writing. Users stay in control by accepting the suggestions that suit their needs.

Example enterprise use cases:

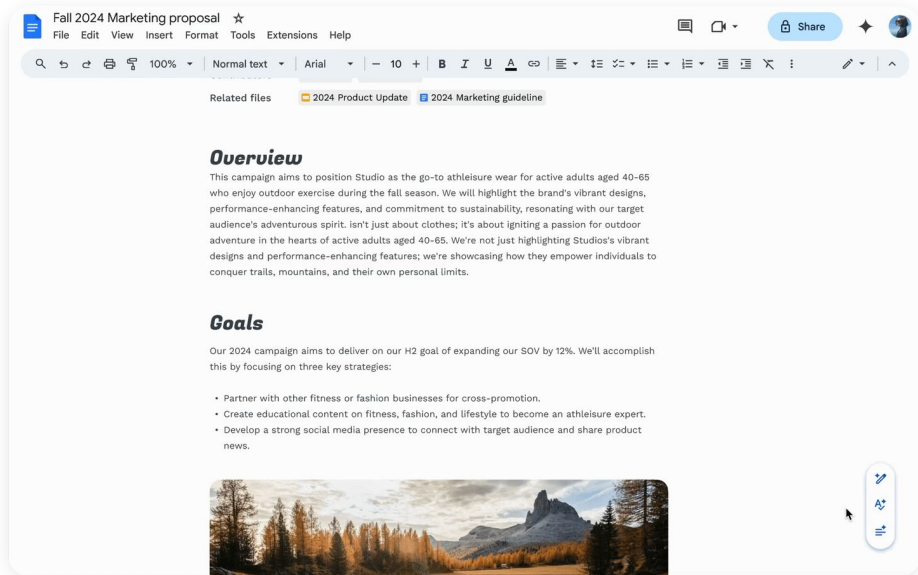
- **Marketing:** refining blog posts, press releases, ad, and web copy
- **Sales:** refining training materials and custom pitches
- **HR:** refining job descriptions and employee handbooks

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Google Cloud Next

Gemini in Docs proactively suggests:

- Improvements for writing (e.g. conciseness, word choice)
- Improvements for content (e.g. ways to strengthen an argument)
- Improvements for document structure (e.g. adding imagery, cleaning up fonts and formatting)



Learn more: [Blog](#)

Side panel capabilities in Meet: What is new

GA

June 2025

Preview

April 2025

Side panel in Google Meet is an in-meeting AI assistant that can help people perform at their best in meetings and drive successful outcomes.

Industries / use cases:

- This feature can be useful to knowledge workers in all industries

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Side panel in Google Meet offers the following new capabilities:

- Help meeting attendees summarize and understand ongoing discussions
- Help attendees quickly catch up on what they missed if they were multi-tasking, or joined late (only if Take notes with Gemini was on)
- Help attendees identify key takeaways, decisions, and action items to inform next steps after the meeting

Noteworthy customers / metrics

This is a brand new feature so we don't yet have metrics

Learn more: [Blog](#), [Website page](#), [Transform team collaboration with Gemini in Google Chat and Meet](#), [Reduce costs and improve collaboration by moving to Google Chat and Meet](#)

AI model: Holistic AI threat defense model in Gmail

GA

April 2025

Stronger protections to instantly identify risky content and block new threats as quickly as they emerge

Industries / use cases:

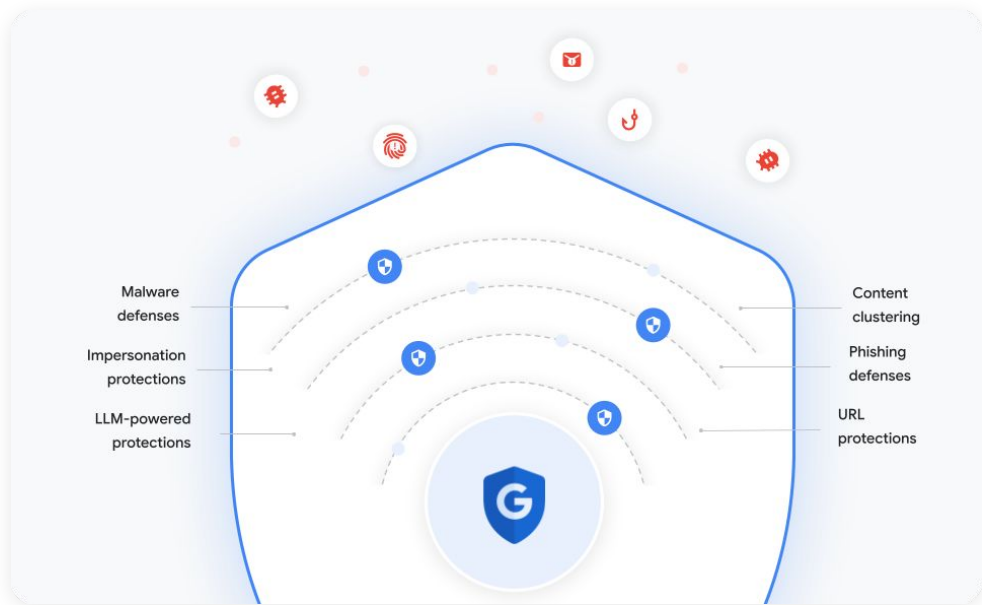
- Beneficial for all by reducing spam, phishing, and malware from reaching users inboxes

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

More AI-powered threat defenses in Gmail

What's new: Holistic model acts as an “AI shield” to bolster existing defenses & protections by evaluating thousands of actor, behavior, and content signals



Learn more: [Blog](#), [Google Workspace threat prevention](#), [Next Session](#)

Gmail DLP & automated classification labels: What is new

GA

February 2025 (DLP)

April 2025 (Labels)

Automated data loss prevention (DLP) with classification labels in Gmail

Industries / use cases:

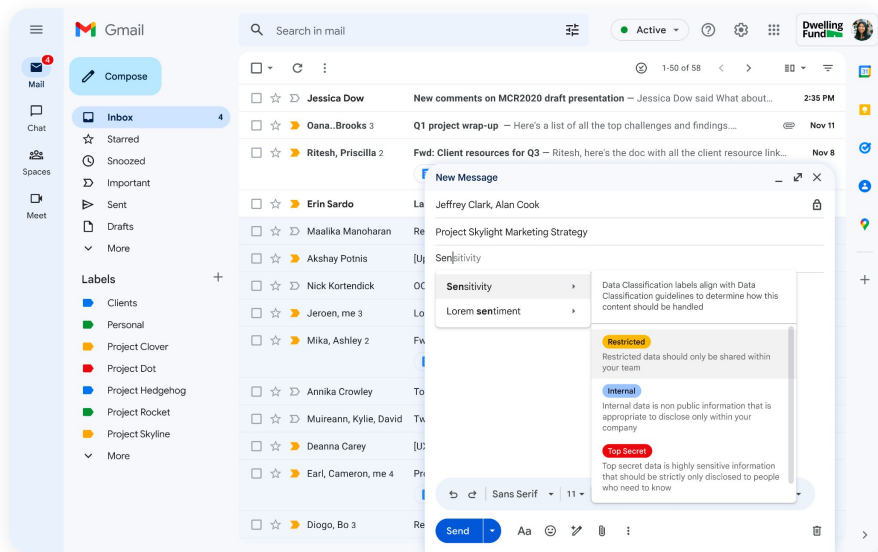
- Beneficial for all by reducing data exfiltration and protecting sensitive business data

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Data protection in Gmail made easier

- What's new:** Protect sensitive data from being shared in email using data protection controls consistent with Drive & Chat policies
- What's new:** Make your IT team more efficient and educate users by simplifying and automating data protection with classification labels



Learn more: [Blog](#), [Google Workspace zero trust](#), [Gmail DLP GA](#), [Gmail labels beta](#), [Next Session](#)

Gmail: Send private, encrypted emails to any inbox

Preview

April 2025

Send end-to-end encrypted emails to anyone, with encryption keys you control in a location of your choice to keep your data inaccessible to any third party including Google or foreign governments.

Industries / use cases:

- Governments and defense
- Financial Services
- Heavily regulated commercial industries (telecommunications, aerospace)

Note: [GA \(Generally Available\)](#), [Preview](#), and [Pre-announcement](#) refer to the launch stage of the product.

Google Cloud Next

Client-side encryption in Gmail is expanding capabilities to enable users to send private, end-to-end encrypted emails to any inbox in the world.

- Businesses can already leverage Client-side encryption in Gmail, enabling their users to send encrypted emails with S/MIME.
- Starting today, we are gradually expanding this capability to enable users to send encrypted emails to any inbox in the world, irrespective of whether their and the recipient's organizations have enabled S/MIME.
- Business users will have the ability to revoke access or set expiration dates to emails.

Learn more: [Blog](#), [Workspace Sovereign AI](#), [Workspace: Safe and compliant collaboration](#)

Gemini leads the way in compliance

Gemini's compliance evaluated and validated by trusted, independent auditors:

Target industries / use cases:

- Governments and defense
- Financial Services
- Heavily regulated commercial industries (telecommunications, aerospace)

Note: **GA** (Generally Available), **Preview**, and **Pre-announcement** refer to the launch stage of the product.

Gemini in Google Workspace and the Gemini app are the first enterprise-ready, certified AI assistants for productivity and collaboration.

- **What's new 1:** Gemini is the first productivity and collaboration AI assistant to have received a comprehensive set of safety, privacy, and security certifications including ISO 42001 - the world's first international standard for AI Management Systems.
- **What's new 2:** Gemini can already help meet HIPAA compliance.
- **What's new 3:** Gemini is also the first generative AI assistant to achieve the German BSI C5 certification.
- **What's new 4:** Gemini is also the first generative AI assistant to achieve the US FedRAMP High attestation.

Gemini is the first enterprise-ready, certified AI assistant

	Gemini in Workspace	Gemini app gemini.google.com	Microsoft 365 Copilot*	ChatGPT Enterprise*
SOC 1	✓	✓		
SOC 2/3	✓	✓		✓
ISO 27001	✓	✓		
ISO 27017	✓	✓		
ISO 27018	✓	✓	✓	
ISO 27701	✓	✓		
ISO 9001	✓	✓		
ISO 42001	✓	✓		
CSA STAR	✓	✓		✓
GDPR	✓	✓	✓	✓
CCPA	✓	✓	✓	✓
HIPAA	✓	✓	✓	
New FedRAMP High	✓	✓		
New BSI C5	✓	✓		

* Based on analysis of publicly available Microsoft and OpenAI documentation as of 03/14/2025

Learn more: [Blog](#), [FedRAMP High Announcement Blog](#), [Workspace Sovereign AI](#), [Workspace: Safe and compliant collaboration](#)

Gemini in Workspace: Data Regions support

GA

February 2025

Sovereign control over your Gemini data.

Industries / use cases:

- EU companies adhering to the Privacy Shield
- US Public Sector organizations

Note: GA (Generally Available), Preview, and Pre-announcement refer to the launch stage of the product.

Google Cloud Next

Gemini now adheres to data regionalization requirements.

- **What's new 1:** We offer EU and US customers advanced data residency controls if they want to stay ahead of evolving regulations and keep their Gemini data at rest and in inference in region
- **What's new 2:** Administrators have flexibility to set up sovereign controls for EU processing, US processing or a mix of both, and powerful controls with granular settings down to organizational units.
- **What's new 3:** This can help meet data residency regulations such as Google Developer Program R and ITAR.

Noteworthy customers

- Schwarz Group
- Air Liquide
- Randstad
- Iron Mountain

Learn more: [Blog](#), [Workspace Data Regions blog](#), [Workspace Sovereign AI](#), [Workspace: Safe and compliant collaboration](#)

04. Offers

Select Next '25 Offers that are available via external links that can be found on each slide.

Source: Individual links to externally available claim information can be found on each slide.

Get early access to Google Workspace Flows

Discover how to automate your workflows using AI

Do you ever find yourself in an inefficient workflow, wishing it could be more automated? Easily create agentic workflows using Google Workspace Flows and Gemini in Google Workspace — and connect to third-party data and apps. Automate inefficient tasks like gaining approvals, organizing email, and collecting data across several different applications.

Don't have Google Workspace?

No problem! Our team can help you get set up with a free trial.

**Restrictions apply. Once you fill out the form, a member of your account team will be in touch with next steps.*

More details: [Link](#)

Start implementing what you learned at Next 25 – with help from Google Cloud

Get expert guidance for your next move

Put those insights from Next '25 to work, right away. Let our team create a personalized roadmap for your projects, including tailored Google Cloud offers and comprehensive support.

**Restrictions apply. Once you fill out the form, a member of your account team will be in touch with next steps.*

More details: [Link](#)

Start the journey to transform how you work

Unleash the potential of every employee with AI

With the introduction of new AI technologies, how you work matters more than ever — and Google is here to support you with this transition.

Find out if your company qualifies for special offers, migration services, and how to get started with a trial of Google Workspace with Gemini.

**Restrictions apply. Once you fill out the form, a member of your account team will be in touch with next steps.*

More details: [Link](#)

Make the most of Google Agentspace

Build AI agents on Google Agentspace

Your Next '25 offer includes:

- **Help identifying impactful use cases:** Collaborate with our experts to pinpoint the most valuable applications for your business.
- **Priority access to Google Agentspace Accelerator program:** Fast-track your team's decision-making in a few short weeks with quick, secure data delivery.
- **Early access to new features:** Be among the first to access the [latest Google Agentspace features](#).
- **Invitations to exclusive events:** Get invited to upcoming Google Agentspace events, interactive labs, and more.

**Restrictions apply. Once you fill out the form, a member of your account team will be in touch with next steps.*

More details: [Link](#)



Thank you