Google Cloud

Google Security Overview

October 2026

Table of contents

Google's security and privacy-focused culture **5**
Our dedicated security team 5
Collaboration with the security research community 6
Internal security and privacy training 7
Our dedicated privacy team 7
Internal audit and compliance specialists 8
Operational security **8**
Vulnerability management 8
Malware prevention 9
Security monitoring 9
Incident management 10
Technology with security at its core **10**
State-of-the-art data centers 10
Powering our data centers 11
Custom server hardware and software 11
Hardware tracking and disposal 12
Software development practices 12
Key security controls **13**
Encryption 13
Securing data at rest 13
Securing data in transit 13
Software supply chain integrity 14
Securing data in use 14
Security benefits of our global network 14
Low latency and highly available solutions 15
Google Cloud service availability 16
Data access and restrictions **16**
Data usage 16
Administrative access for Google employees 16
Law enforcement data requests 17
Third-party suppliers 17
Support for compliance requirements **18**
Compliance for Google models 18
Risk management and insurance 18
Google Cloud security products and services **19**

Conclusion **20**
What's next **20**

This content was last updated in October 2026 and represents the status quo as of the time it was written. Google's security policies and systems may change going forward, as we continually improve protection for our customers.

Traditionally, businesses have looked to the public cloud to save costs, experiment with new technology, and provide growth capacity. Increasingly, businesses are also looking to the public cloud for their security, realizing that cloud providers can invest more than the businesses can in technology, people, and processes to deliver a more secure infrastructure.

As a cloud innovator, Google understands security in the cloud. Our cloud services are designed to deliver better security than many on-premises approaches. We make security a priority in our operations—operations that serve users across the world.

Security drives our organizational structure, culture, training priorities, and hiring processes. It shapes the design of our data centers and the technology that they house. It's central to our everyday operations and disaster planning, including how we address threats. It's prioritized in the way we handle customer data, our account controls, our compliance audits, and our certifications.

This document describes our approach to security, privacy, and compliance for Google Cloud, which is our suite of public cloud products and services. The document focuses on the physical, administrative, and technical controls that we have deployed to help protect your data.

The following diagram shows how Google integrates culture, infrastructure, operational security, encryption, data governance, and compliance into a unified defense-in-depth architecture:

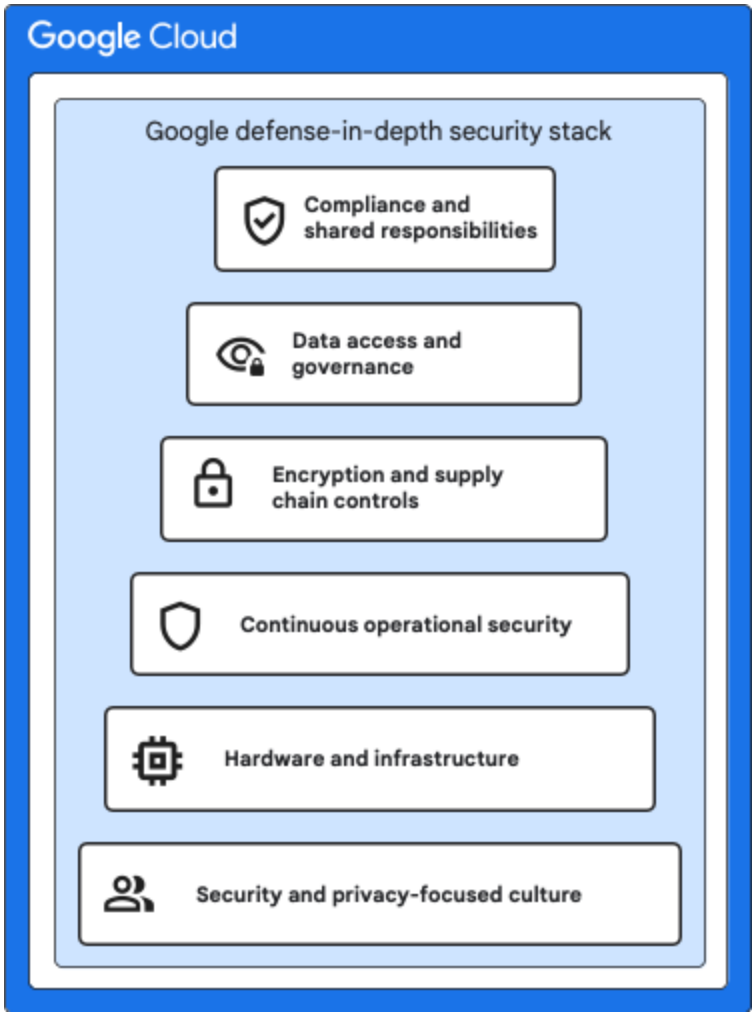

Google’s security and privacy-focused culture

Google's culture stresses the importance of protecting the large volume of information that belongs to our customers. This culture influences our hiring processes and employee onboarding. We advance and reinforce data protection guidelines and technologies through ongoing training and events that are focused on security and privacy.

Our dedicated security team

Our dedicated security team includes some of the world’s foremost experts in information security, application security, cryptography, and network security. This team maintains our

defense systems, develops security review processes, builds security infrastructure, and implements our security policies. The team actively scans for security vulnerabilities using commercial and custom tools. The team also conducts penetration tests and performs quality assurance and security reviews.

Members of the security team review security plans for our networks and services, and they provide project-specific consulting services to our product and engineering teams. For example, our cryptography engineers review product launches that include cryptography implementations.

The security team monitors for suspicious activity on our networks and addresses information security threats as needed. The team also performs routine security evaluations and audits, which can involve engaging outside experts to conduct regular security assessments.

Collaboration with the security research community

We have long enjoyed a close relationship with the security research community, and we greatly value their help with identifying potential vulnerabilities in Google Cloud and other Google products. Our security teams take part in research and outreach activities to benefit the online community. For example, we run Project Zero, which is a team of security researchers who are dedicated to researching zero-day vulnerabilities. Some examples of this research are the discovery of the Spectre exploit, the Meltdown exploit, the POODLE SSL 3.0 exploit, and cipher suite weaknesses. We also pioneered AI-assisted vulnerability discovery with Project Naptime and Big Sleep agent to find zero-day vulnerabilities in critical open-source software.

Google's security engineers and researchers actively participate and publish in the academic security community and the privacy research community. They also organize and participate in open-source projects and academic conferences. Google's security teams have published an in-depth account of our practices and experience in the Building Secure and Reliable Systems book.

Our Vulnerability Reward Program offers rewards in the tens of thousands of dollars for each confirmed vulnerability. The program encourages researchers to report design and implementation issues that might put customer data at risk. In 2025, we awarded researchers over 17 million dollars in prize money. To help improve the security of open-source code, the Vulnerability Reward Program also provides a variety of initiatives to researchers. For more information about this program, including the rewards that we've given, see Bug Hunters Key Stats.

Our world-class cryptographers participate in industry-leading cryptography projects. To protect TLS connections and data against quantum computer attacks, we developed hybrid post-quantum key exchange on Chrome, the Google Front End (GFE), and our Application Layer Transport Security (ALTS) protocol using the ML-KEM (FIPS 203) algorithm. Our

cryptographers developed Tink, which is an open-source library of cryptographic APIs. We also use Tink in our internal products and services.

For more information about how you can report security issues, see How Google handles security vulnerabilities. For more information about our post-quantum cryptography roadmap, see PQC in Plaintext: Google Cloud's post-quantum cryptography roadmap.

Internal security and privacy training

All Google employees undergo security and privacy training as part of the orientation process, and they receive ongoing security and privacy training throughout their Google careers. During orientation, new employees agree to our Code of Conduct, which highlights our commitment to keeping customer data safe and secure.

Depending on their job role, employees might be required to take additional training on specific aspects of security. For example, the information security team instructs new engineers on secure coding practices, product design, and automated vulnerability testing tools. Engineers attend regular security briefings and receive security newsletters that cover new threats, attack patterns, mitigation techniques, and more.

Security and privacy are an ever-changing area, and we recognize that dedicated employee engagement is a key means of raising awareness. We host regular internal conferences that are open to all employees to raise awareness and drive innovation in security and data privacy. We host events across global offices to raise awareness of security and privacy in software development, data handling, and policy enforcement.

Our dedicated privacy team

Our dedicated privacy team supports internal privacy initiatives that help improve critical processes, internal tools, products, and privacy infrastructure. The privacy team operates separately from product development and security organizations. They participate in Google product launches by reviewing design documentation and performing code reviews to help ensure that privacy requirements are followed. The team helps release products that incorporate strong privacy standards around the collection of user data.

Our products are designed to provide users and administrators with meaningful privacy configuration options. After products are launched, the privacy team oversees ongoing automated processes to verify that data collected by the products is handled appropriately. In addition, the privacy team conducts research on privacy best practices for our emerging technologies. To understand how we stay committed to user data privacy and to compliance with applicable privacy regulations and laws, see our commitment to complying with data protection laws. For more information, see the Privacy Resource Center.

Internal audit and compliance specialists

We have a dedicated internal audit team that reviews our products' compliance with security laws and regulations around the world. As new auditing standards are created and existing standards are updated, the internal audit team determines what controls, processes, and systems are needed to help meet them. This team supports independent audits and assessments by third parties. For more information, see Support for compliance requirements later in this document.

Operational security

Security is an integral part of our cloud operations, not an afterthought. This section describes our vulnerability management programs, malware prevention program, security monitoring, and incident management programs.

Vulnerability management

Our internal vulnerability management process actively scans for security threats across technology stacks. This process uses a combination of commercial, open-source, and purpose-built in-house tools, and includes the following:

- Quality assurance processes
- Software security reviews
- Intensive automated and manual penetration efforts, including extensive Red Team exercises
- External audits

The vulnerability management organization and its partners are responsible for tracking and following up on vulnerabilities. Because security improves only after issues are fully addressed, automation pipelines continually reassess the state of patch deployment to mitigate vulnerabilities and flag incorrect or partial deployment.

To help improve detection capabilities, the vulnerability management organization focuses on high-quality indicators that separate noise from signals that indicate real threats. The organization also fosters interaction with the industry and with the open-source community. For example, they run a Patch Reward Program for the Tsunami network security scanner, which rewards developers who create open-source detectors for vulnerabilities.

We also use AI agents and machine learning models across our vulnerability management lifecycle to scan, verify, and remediate vulnerabilities in our infrastructure. Specialized AI agents such as Big Sleep and CodeMender help our security teams discover zero-day vulnerabilities,

validate exploitability using proof-of-concept exploits in isolated sandboxes, and generate tested patches for human review.

For more information about vulnerabilities that we have mitigated, see Google Cloud security bulletins.

Malware prevention

Google maintains malware protections for our core products (like Gmail, Google Drive, Google Chrome, YouTube, Google Ads, and Google Search) that use a variety of malware detection techniques. To discover malware files proactively, we use web crawling, file detonation, custom static detection, dynamic detection, and machine-learning detection. We also use multiple antivirus engines.

To help protect our employees, we use the built-in advanced security capabilities of Chrome Enterprise Premium and the Enhanced Safe Browsing feature in Google Chrome. These capabilities enable proactive detection of phishing and malware sites as our employees browse the web. We also enable the most rigorous security settings that are available in Google Workspace, such as Gmail Security Sandbox, to proactively scan suspicious attachments. Logs from these capabilities feed into our security monitoring systems, as described in the following section.

Security monitoring

Our security monitoring program is focused on information that's gathered from internal network traffic, from employee actions on systems, and from outside knowledge of vulnerabilities. A core Google principle is to aggregate and store security telemetry data in one location for unified security analysis.

At many points across our global network, internal traffic is inspected for suspicious behavior, such as the presence of traffic that might indicate botnet connections. We use a combination of open-source and commercial tools to capture and parse traffic so that we can perform this analysis. A proprietary correlation system built on top of our technology also supports this analysis. We supplement network analysis by examining system logs to identify unusual behavior, such as attempts to access customer data.

Our security engineers review inbound security reports and monitor public mailing lists, blog posts, and wikis. Automated network analysis and automated analysis of system logs help determine when an unknown threat might exist; if the automated processes detect an issue, they escalate it to our security staff.

Our threat detection systems also use machine learning and AI models to correlate security signals across our infrastructure, detect behavioral anomalies, and analyze audit logs for non-standard access attempts. To help prevent alert fatigue and accelerate response, we use AI models to classify, deduplicate, and rank high-volume alert streams into consolidated events with contextual threat diagnostics. For more information, see Identification.

Incident management

We have a rigorous incident-management process for security events that might affect the confidentiality, integrity, or availability of systems or data. Our security incident-management program aligns with the NIST guidance on handling incidents (NIST SP 800-61 Rev. 3). Key members of our staff are trained in forensics and in handling evidence in preparation for an event, including the use of third-party and proprietary tools.

We test incident response plans for key areas, such as systems that store customer data. These tests consider various scenarios, including insider threats and software vulnerabilities. To help ensure the swift resolution of security incidents, the Google security team is available 24/7 to all employees. During an incident, we also use AI tools to help curate timeline updates from response channels and autopopulate initial retrospective drafts for engineering review. If an incident affects your data, Google or its partners inform you and our team investigates the incident.

For more information about our data incident response process, see Data incident response process.

Technology with security at its core

Google Cloud runs on a technology platform that is designed and built to operate securely. We are an innovator in hardware, software, network, and system management technologies. We design our servers, our proprietary operating system, and our geographically distributed data centers. Using the principles of defense in depth, we've created an IT infrastructure that is more secure and easier to manage than more conventional technologies.

State-of-the-art data centers

Our focus on security and protection of data is among our primary design criteria. The physical security in Google data centers is a layered security model. Physical security includes safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors, and biometrics. In addition, to detect and track intruders, we use security measures such as laser beam intrusion detection and 24/7 monitoring by high-resolution interior and exterior cameras. Access logs, activity records, and camera footage

are available in case an incident occurs. Experienced security guards, who have undergone rigorous background checks and training, routinely patrol our data centers. As you get closer to the data center floor, security measures also increase. Access to the data center floor is only possible through a security corridor that implements multi-factor access control using security badges and biometrics. Only approved employees with specific roles may enter. Very few Google employees ever gain access to one of our data centers.

Inside our data centers, we employ security controls in the *physical-to-logical space*, defined as "arm's length from a machine in a rack to the machine's runtime environment." These controls include hardware hardening, task-based access control, anomalous event detection, and system self-defense. For more information, see How Google protects the physical-to-logical space in a data center.

Powering our data centers

To help keep things running 24/7 and provide uninterrupted services, our data centers have redundant power systems and environmental controls. Every critical component has a primary and alternate power source, each with equal power. Backup generators can provide enough emergency electrical power to run each data center at full capacity. Cooling systems maintain a constant operating temperature for servers and other hardware, which reduces the risk of service outages while minimizing environmental impact. Fire detection and suppression equipment help prevent damage to hardware. Heat detectors, fire detectors, and smoke detectors trigger audible and visible alarms at security operations consoles and at remote monitoring desks.

We are the first major internet services company to get external certification of our high environmental, workplace safety, and energy management standards throughout our data centers. For example, to demonstrate our commitment to energy management practices, we obtained voluntary ISO 50001 certifications for our data centers in Europe. For more information about how we reduce our environmental impact in Google Cloud, see Efficiency.

Custom server hardware and software

Our data centers have purpose-built servers and network equipment, some of which we design ourselves. While our servers are customized to maximize performance, cooling, and power efficiency, they are also designed to help protect against physical intrusion attacks. Unlike most commercially available hardware, our servers don't include unnecessary components such as video cards, chipsets, or peripheral connectors, all of which can introduce vulnerabilities. We vet component vendors and choose components with care, working with vendors to audit and validate the security properties that are provided by the components. We design custom chips, such as Titan, that help us securely identify and authenticate legitimate Google devices at the hardware level, including the code that these devices use to boot up.

In our server fleet, the Titanium hardware security architecture uses purpose-built Titanium offload processors (TOPs) to isolate management, storage, and networking workloads from the host CPU, enforce cryptographic boot verification, and verify firmware security patch levels.

Server resources are dynamically allocated. Dynamic allocation gives us flexibility for growth and lets us adapt quickly and efficiently to customer demand by adding or reallocating resources. This environment is maintained by proprietary software that continually monitors systems for binary-level modifications. Our automated, self-healing mechanisms are designed to let us monitor and remediate destabilizing events, receive notifications about incidents, and slow down potential compromises on the network.

Hardware tracking and disposal

We meticulously track the location and status of equipment within our data centers using barcodes and asset tags. We deploy metal detectors and video surveillance to help make sure that no equipment leaves the data center floor without authorization. If a component fails to pass a performance test at any point during its lifecycle, it's removed from inventory and retired.

Our storage devices, including hard drives, solid-state drives, and non-volatile dual inline memory modules (DIMMs), use technologies like full disk encryption (FDE) and drive locking to protect data at rest. When a storage device is retired, authorized individuals verify that the device is sanitized. They also perform a multiple-step verification process that checks whether data is still on the device. If a device cannot be erased for any reason, it's physically destroyed. Physical destruction is performed using a shredder that breaks the device into small pieces, which are then recycled at a secure facility. Each data center adheres to a strict disposal policy and any variances are immediately addressed. For more information, see Data deletion on Google Cloud.

Software development practices

We proactively seek to limit the opportunities for vulnerabilities to be introduced by using source control protections and two-party reviews. We also provide libraries that prevent developers from introducing certain classes of security bugs. For example, we have libraries and frameworks that are designed to eliminate XSS vulnerabilities in web apps. We also have automated tools that detect security bugs; these tools include fuzzers, static analysis tools, and web security scanners.

Through our Safe Coding strategy, we prioritize memory-safe languages such as Rust and deploy spatial memory safety checking in C++. Developers also use AI-powered coding agents and IDE integrations, such as CodeMender, Gemini Code Assist, and Google Antigravity, to detect, test, and repair security issues before code is submitted.

For more information, see Safe software development.

Key security controls

Google Cloud services are designed to deliver better security than many on-premises solutions. This section describes the main security controls that we use to help protect your data.

Encryption

Encryption adds a layer of defense for protecting data. Encryption helps ensure that if an attacker gets access to your data, the attacker cannot read the data without also having access to the encryption keys. Even if an attacker gets access to your data (for example, by accessing the wire connection between data centers or by stealing a storage device), they won't be able to understand or decrypt it.

Encryption provides an important mechanism in how we help protect the privacy of your data. It lets systems manipulate data—for example, for backup—and engineers to support our infrastructure, without providing access to content for those systems or employees.

Securing data at rest

By default, Google Cloud uses several layers of encryption to protect user data that's stored in Google production data centers. Encryption is applied at the application layer, the storage device layer, or both layers.

For more information about encryption at rest, including encryption key management and Keystore, see Encryption at rest in Google Cloud.

Securing data in transit

Data can be vulnerable to unauthorized access as it travels across the internet or within networks. Traffic between your devices and the Google Front End (GFE) is encrypted using strong encryption protocols such as TLS.

Our core cryptographic module for TLS, BoringCrypto (part of BoringSSL), is validated to FIPS 140-3 Level 1. Cloud service providers such as Google are required to implement FIPS security controls to satisfy the requirements of cloud computing for the US government, the Canadian government, and their contractors and vendors. Google Cloud uses FIPS 140-validated modules to process data in transit, including SSH connections, data center traffic, service-to-service connections, and external interfaces that use TLS 1.2 or later. For more information about our compliance with FIPS 140-3, see FIPS 140 Requirements.

In our data centers, we help protect network traffic using ALTS, which supports hybrid post-quantum key exchange using ML-KEM (FIPS 203) and hardware-offloaded encryption protocols such as PSP on Titanium network adapters.

For more information, see Encryption in transit in Google Cloud.

Software supply chain integrity

Software supply chain integrity verifies the underlying code and binaries for the services that process your data by performing attestation tests. In Google Cloud, we developed Binary Authorization for Borg (BAB) to review and authorize production software that we deploy. BAB helps ensure that only authorized code can process your data. In addition to BAB, we use hardware security chips (called Titan) that we deploy on servers, devices, and peripherals. These chips offer core security features such as secure key storage, root of trust, and signing authority.

To help secure your software supply chain, you can implement Binary Authorization to enforce your policies before deploying your code. For information about securing your supply chain, see SLSA.

Securing data in use

Google Cloud supports data encryption for data in use with Confidential Computing. Confidential Computing provides hardware-enforced memory encryption, isolation, and attestation using a Trusted Execution Environment (TEE). Confidential Computing protects workloads by performing computation in cryptographic isolation, which helps to ensure confidentiality in a multi-tenant cloud environment. This type of cryptographically isolated environment helps prevent unauthorized access or modifications to applications and data while the applications and data are in use. TEE provides independently verifiable attestations that attest to the system state and code run. Confidential Computing might be a good option for organizations that manage sensitive and regulated data and that need verifiable security and privacy assurances.

For multi-party data collaboration and privacy-preserving AI workloads, Confidential Space helps ensure that confidential data is visible only to authorized, attested workloads.

Security benefits of our global network

In other cloud services and on-premises solutions, customer data travels between devices across the internet in paths known as *hops*. The number of hops depends on the optimal route between the customer's ISP and the data center. Each additional hop introduces a new opportunity for data to be attacked or intercepted. Because our global network is linked to most

ISPs in the world, our network limits hops across the internet, and therefore helps limit access to that data by bad actors.

Our network uses multiple layers of defense—defense in depth—to help protect the network against external attacks. Only authorized services and protocols that meet our security requirements are allowed to traverse it; anything else is automatically dropped. To enforce network segregation, we use firewalls and access control lists. Traffic is routed through GFE servers to help detect and stop malicious requests and distributed denial-of-service (DDoS) attacks. Logs are routinely examined to reveal any exploitation of programming errors. Access to networked devices is restricted to only authorized employees.

Our global infrastructure lets us run Project Shield. Project Shield provides free, unlimited protection to websites that are vulnerable to DDoS attacks that are used to censor information. Project Shield is available for news websites, human rights websites, and election-monitoring websites.

Low latency and highly available solutions

Our IP data network consists of our own fiber, of publicly available fiber, and of undersea cables. This network lets us deliver highly available and low-latency services across the globe.

We design the components of our platform to be highly redundant. This redundancy applies to our server design, to how we store data, to network and internet connectivity, and to the software services themselves. This “redundancy of everything” includes exception handling and creates a solution that isn’t dependent on a single server, data center, or network connection.

Our data centers are geographically distributed to minimize the effects of regional disruptions on global products, such as when natural disasters or local outages occur. If hardware, software, or a network fails, platform services and control planes are designed to shift automatically and swiftly from one facility to another so that platform services can continue without interruption.

Our highly redundant infrastructure also helps you protect your business from data loss. You can create and deploy Google Cloud resources across multiple regions and zones to build resilient and highly available systems. Our systems are designed to minimize downtime or maintenance windows for when we need to service or upgrade our platform. For more information about how Google Cloud builds resilience and availability into its core infrastructure and services, from design through operations, see the Google Cloud infrastructure reliability guide.

Google Cloud service availability

Some Google Cloud services aren't available in all geographies. Some service disruptions are temporary (due to an unanticipated event, such as a network outage), but other service limitations are permanent due to government-imposed restrictions. Our comprehensive Transparency Report and status dashboard show recent and ongoing disruptions of traffic and availability of Google Cloud services. We provide this data to help you analyze and understand the availability of services.

Data access and restrictions

This section describes how we restrict access to data and how we respond to data requests from law enforcement agencies.

Data usage

Data that you store on our systems is yours. We don't scan your data for advertising purposes, we don't sell it to third parties, and we don't use it to train our AI models without your permission. The Data Processing Addendum for Google Cloud describes our commitment to protecting your data. That document states that we won't process data for any purpose other than to meet our contractual obligations. If you choose to stop using our services, we provide tools that let you take your data with you, without penalty or additional cost. For more information about our commitments for Google Cloud, see our trust principles.

Administrative access for Google employees

Our infrastructure is designed to logically isolate each customer's data from the data of other customers and users, even when it's stored on the same physical server. Only a small group of employees have access to customer data. Access rights and levels are based on an employee's job function and role, using the principles of least privilege and need-to-know that match access privileges to defined responsibilities. Our employees are granted only a limited set of default permissions to access company resources, such as employee email and Google's internal employee portal. Requests for additional access must follow a formal process that involves a request and an approval from the data or system owner, manager, or other executives, as dictated by our security policies.

Approvals are managed by workflow tools that maintain audit records of all changes. These tools control both the modification of authorization settings and the approval process to help ensure that approval policies are consistently applied. An employee's authorization settings are used to control access to resources, including data and systems for Google Cloud products.

Support services are provided only to authorized customer administrators. Our dedicated security teams, privacy teams, and internal audit teams monitor and audit employee access, and we provide audit logs to you through Access Transparency for Google Cloud. Also, when you enable Access Approval, our support personnel and our engineers require your explicit approval to access your data.

Under our Zero Touch Prod and No Persons (NoPe) security principles, production changes are executed through automated tooling and auditable safe proxies. Engineers don't have permanent administrative privileges by default. Elevated access requires temporary, business-justified Access on Demand (AoD) grants and multi-party authorization (MPA) approval from another qualified individual. For more information, see Privileged access in Google Cloud.

Law enforcement data requests

As the data owner, you are primarily responsible for responding to law enforcement data requests. However, like many technology companies, we receive direct requests from governments and courts to disclose customer information. Google has operational policies and procedures and other organizational measures in place to help protect against unlawful or excessive requests for user data by public authorities. When we receive such a request, our team reviews the request to make sure that it satisfies legal requirements and Google's policies. Generally speaking, for us to comply, the request must be made in writing, issued under an appropriate law, and signed by an authorized official of the requesting agency.

We believe that the public deserves to know the full extent to which governments request information from us. We became the first company to start regularly publishing reports about government data requests. Detailed information about data requests and our response to them is available in our Transparency Report. It's our policy to notify you about requests for your data unless we are specifically prohibited by law or court order from doing so. For more information, see Government Requests for Cloud Customer Data.

Third-party suppliers

For most data-processing activities, we provide our services in our own infrastructure. However, we might engage some third-party suppliers to provide services related to Google Cloud, including customer support and technical support. Before onboarding a supplier, we assess their security and privacy practices. This assessment checks whether the supplier provides a level of security and privacy that is appropriate for their access to data and for the scope of the services that they are engaged to provide. After we have assessed the risks that are presented by the third-party supplier, the supplier is required to enter into appropriate security, confidentiality, and privacy contract terms.

For more information, see the Supplier Code of Conduct.

Support for compliance requirements

Google Cloud regularly undergoes independent verification of its security, privacy, and compliance controls, and receives certifications, attestations, and audit reports to demonstrate compliance. Our information security includes specific customer data privacy-related controls that help keep customer data secure.

Some key international standards that we are audited against are the following:

- ISO/IEC 27001 (Information Security Management)
- ISO/IEC 27017 (Cloud Security)
- ISO/IEC 27018 (Cloud Privacy)
- ISO/IEC 27701 (Privacy)

In addition, our SOC 2 and SOC 3 reports are available to our customers.

We also participate in sector and country-specific frameworks, such as FedRAMP (US government), BSI C5 (Germany), and MTCS (Singapore). We provide resource documents and mappings for certain frameworks where formal certifications or attestations might not be required or applied.

If you operate in regulated industries, such as finance, government, healthcare, or education, Google Cloud provides products and services that help you be meet numerous industry-specific requirements. See Assured Workloads overview for information about how you can implement regulatory requirements in Google Cloud.

For a complete listing of our compliance offerings, see the Compliance resource center.

Compliance for Google models

Google Cloud services can include Google models. The certifications that a Google Cloud service obtained also apply to the Google models that the service includes, unless explicitly documented otherwise.

Risk management and insurance

We maintain a robust insurance program for many risk types, including cyber and privacy liability insurance coverage. These policies include coverage for Google Cloud in events such as unauthorized use or access of our network; regulatory action where insurable; failure to

adequately protect confidential information; notification costs; and crisis management costs, including forensic investigation.

Google Cloud security products and services

Security is a shared responsibility. Generally, you are responsible for securing what you bring to the cloud, whereas we are responsible for protecting the cloud itself. Therefore, while you're always responsible for securing your data, we are responsible for securing the underlying infrastructure. The following image visualizes this relationship as the shared responsibility model, which describes the responsibilities that we and you have in Google Cloud.

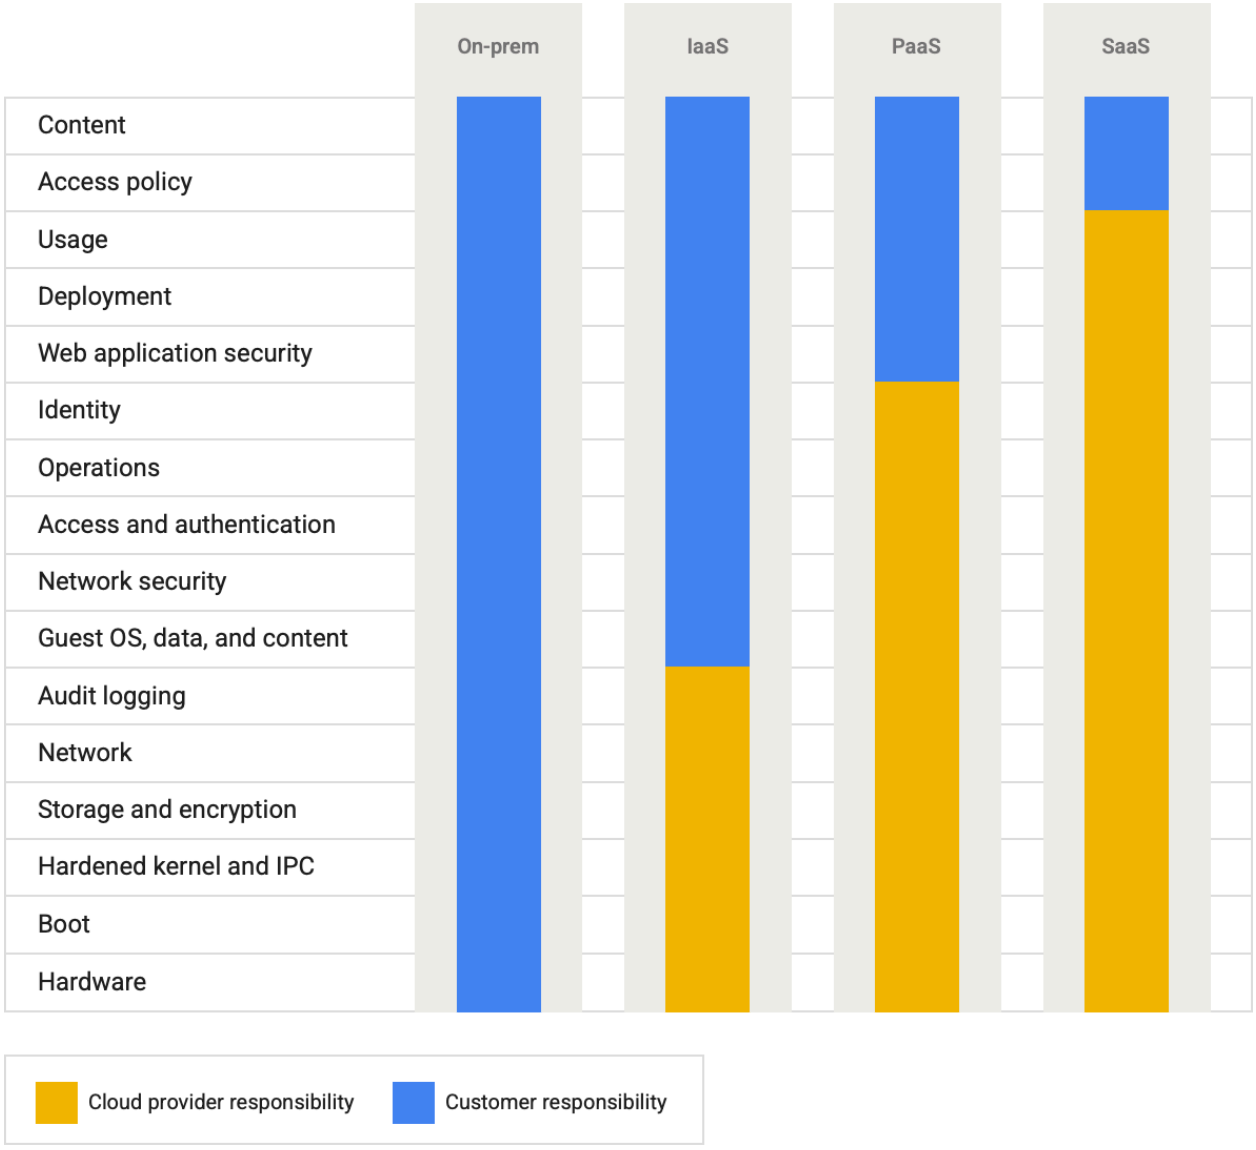

In the infrastructure as a service (IaaS) model, only the hardware, storage, and network are our responsibility. In the software as a service (SaaS) model, the security of everything except the data and its access and usage are our responsibility.

Google Cloud offers a range of security services that you can take advantage of to secure your cloud environment at scale. For more information, see Security and identity products in Google Cloud. You can also find more information in our security best practices center.

Conclusion

The protection of your data is a primary design consideration for our infrastructure, products, and operations. Our scale of operations and our collaboration with the security research community help us address vulnerabilities quickly, and often prevent them. We run our own services, such as Search, YouTube, and Gmail, on the same infrastructure that we make available to our customers, who benefit directly from our security controls and practices.

We offer a level of protection that few public cloud providers or private enterprise IT teams can match. Protecting data is core to our business, so we make extensive investments in security, resources, and expertise at a scale that others cannot. Our investment frees you to focus on your business and innovation. Our strong contractual commitments help you maintain control over your data and how it's processed. We don't use your data for advertising or any purpose other than to deliver Google Cloud services.

Many innovative organizations trust us with their most valuable asset: their data. We continually invest in the security of Google Cloud services to let you benefit from our services in a secure and transparent manner.

What's next

- To learn more about our security culture and philosophy, read Building Secure and Reliable Systems (O'Reilly book).
- For information about our novel approach to cloud security, read BeyondProd, which describes how to protect code change and access to user data in microservices.
- To adopt similar security principles for your own workloads, deploy the enterprise foundations blueprint.
- To learn more about Google Workspace security, see Google Workspace security.