




Google Cloud Whitepaper
October 2020

Privileged Access Management in Google Cloud Platform



Google Cloud

Table of contents

Table of contents	1
CIO level summary	2
Disclaimer	2
What is privileged access	3
Privileged access philosophy	3
How Google controls and monitors access to customer data	4
Infrastructure-wide controls	4
Strong authentication for all access	4
Use context instead of trusted networks for all access: BeyondCorp	4
Review and authorize all production software	5
Detect and remediate unauthorized access	6
Detailed logging of all access	6
Incident Detection and Response	6
Signal development	6
Threat Detection	6
Incident Investigation	7
Remediation	7
Google personnel by default can't access customer data	8
Reasons for Google personnel to access customer data	10
Customer initiated support	10
Outage or tool failure	10
Third party legal requests	10
Google initiated review	10
Customer tools to monitor and control Google's access to data	11
Access Transparency	11
Access Approvals	11
Cloud External Key Manager (Cloud EKM)	12
Confidential VMs	12
Key Access Justifications - Preview	12
Assured Workloads for Government - Preview	12

CIO level summary

At Google, protecting the sensitive data that customers and enterprises trust us with is a top priority. Our zero trust-based architecture and least privilege principles includes the industry's strongest authentication protocols, is highly resistant to data exfiltration, and deploys 24/7 advanced monitoring and analytics to restrict the misuse of credentials, detect abnormal employee activity, and automatically respond to new or evolving threats.

With a team of security and privacy experts and inventive software design, we built our least-privilege framework from the ground up, guided by four pillars:

- **Strong authentication and role-based non-unilateral access restrictions:** access to customer data, critical operations on production systems, and modifications of source code are all tightly controlled by manual and automated verification systems and, in some cases, requires the cooperation of multiple employees. Employees can only access the resources necessary to do their jobs and must provide a valid justification to access customer data.
- **End-to-end workload protection:** with Confidential VM, Google Cloud Platform is the first major cloud provider with the ability to provide end-to-end encryption of customer workloads in select services and protect data from unauthorized access at rest, in transit, and in use.
- **Continuous logging and auditing:** access to customer data is logged and intelligent threat detection systems conduct real-time audits, alerting staff when log entries match threat indicators. Internal security teams evaluate alerts and logs to identify and investigate anomalous activities, limiting the scope and impact of any incident. Learn more about incident response [here](#).
- **Transparency and customer control:** Using Customer Managed Encryption Keys (CMEK) customers can manage or supply their own encryption keys, allowing them to revoke access to their data at any time. Access Transparency and Access Approval can be used with certain Google services. Access Transparency provides users with near-real time logs of manual targeted accesses by Google administrators. Access Approval allows customers to approve or deny eligible manual, targeted accesses by Google administrators to certain datasets.

Disclaimer

The content contained herein is correct as of October 2020, and represents the status quo as of the time it was written. Google Cloud's security policies and systems may change going forward, as we continually improve protection for our customers.

What is privileged access

Access to customer data in Google Cloud is often because a customer accesses their own data or the service they are using is accessing data on their behalf. In some cases, Google personnel acting as a privileged administrator access customer data in order to help provide the contracted service (e.g. because the customer needs support or to recover from an outage). The Google administrator's access is referred to as privileged access and is what this paper addresses.

Google only accesses customer data in accordance with the Data Processing Security Terms [here](#). Google does not sell customer data to third parties.

Privileged access philosophy

Our philosophy: One hundred percent perfect security does not exist and highly-privileged employees continue to pose a unique risk (i.e. individuals who are temporarily granted or otherwise acquire elevated permissions that allow them to perform actions without oversight or justification). Our approach to privileged access focuses on reducing the number of possible attack vectors through the use of redundant authentication schemes, limiting data access pathways, logging and alerting actions system-wide, and vigilantly regulating permissions. These approaches help us control and detect internal attacks, limiting the impact of incidents and greatly reducing the risk to your data.

Google Cloud's privileged access management strategy limits the power of a single Googler to view or modify customer data associated with a specific account or impersonate any other account. Google Cloud is built such that limits on privileged access are not an optional feature; but instead an integral part of how products are designed to work.

Google's Privileged Access Philosophy is built around a few key guiding principles:

- **Least-privilege:** All Google employee system access is denied by default. When access is granted, it is temporary and should be no greater than absolutely necessary to perform their role.
- **Limit singular access to data:** It should be extremely difficult for Googlers to singularly access customer data without another individual involved.
- **All access must be justified:** Google personnel by default do not have access to customer data¹; but when Google personnel have privileged access, it is related to one of four valid business justifications (see justifications [here](#)).
- **Monitor and alerting:** Monitoring and response processes exist to identify, triage, and remediation of violations of these principles.

¹ Defined as "data provided by or on behalf of Customer or Customer End Users via the Services under the Account" per Data Processing and Security Terms (Customers) [here](#).

How Google controls and monitors access to customer data

Google built its Google Cloud Platform systems with a focus on protecting customer data using controls and monitoring. Google continues to make this a top priority, with dedicated programs and teams focused on protecting customer data.

Google's internal controls have the following aims:

- Pervasive, **infrastructure-wide control** systems to prevent unauthorized access
- **Detect and remediate unauthorized access** through continuous controls monitoring, violation alerting, and regular audits by an internal audit team and independent third party auditors

To learn more about how Google secures the physical infrastructure and the manual processes for protecting data, see Google Infrastructure Security Design Overview [here](#).

Infrastructure-wide controls

Google has built its infrastructure with security at its core. Since Google's global infrastructure is fairly homogeneous, it allows for Google to leverage automated infrastructure to implement controls and limit privileged access in the following ways.

Strong authentication for all access

Google has strong authentication for all access by both users (like an employee) and roles (like a service). Jobs running in our production environment use these identities to access data stores or remote procedure call (RPC) methods of other services. Multiple jobs may run with the same identity. We restrict the ability to deploy or modify jobs with a particular identity to those responsible for running the service—generally our [Site Reliability Engineers \(SREs\)](#). When a job starts it is provisioned with cryptographic credentials. The job uses these credentials to prove its identity when making requests of other services (using [Application Layer Transport Security](#)). [Learn more](#).

Use context instead of trusted networks for all access: BeyondCorp

BeyondCorp provides user and device based authentication and authorization for Google's core infrastructure. Access decisions are not based solely on static credentials or whether they originate from a corporate intranet. The complete context of a request (user identity, location, device ownership and configuration, and fine-grained access policies) is evaluated to determine its validity and guard against phishing attempts and credential-stealing malware.

- Users: in addition to strong passwords, security tokens and other two-factor authentication protocols are required.

- Devices: machine inventories are securely maintained and the state of each connecting device (OS updates, security patches, device certificates, installed software, virus scans, and encryption status, among other factors) is evaluated for potential security risks.
- Access policies: authenticated users and trusted devices are granted limited, temporary access to necessary resources. All traffic is encrypted.

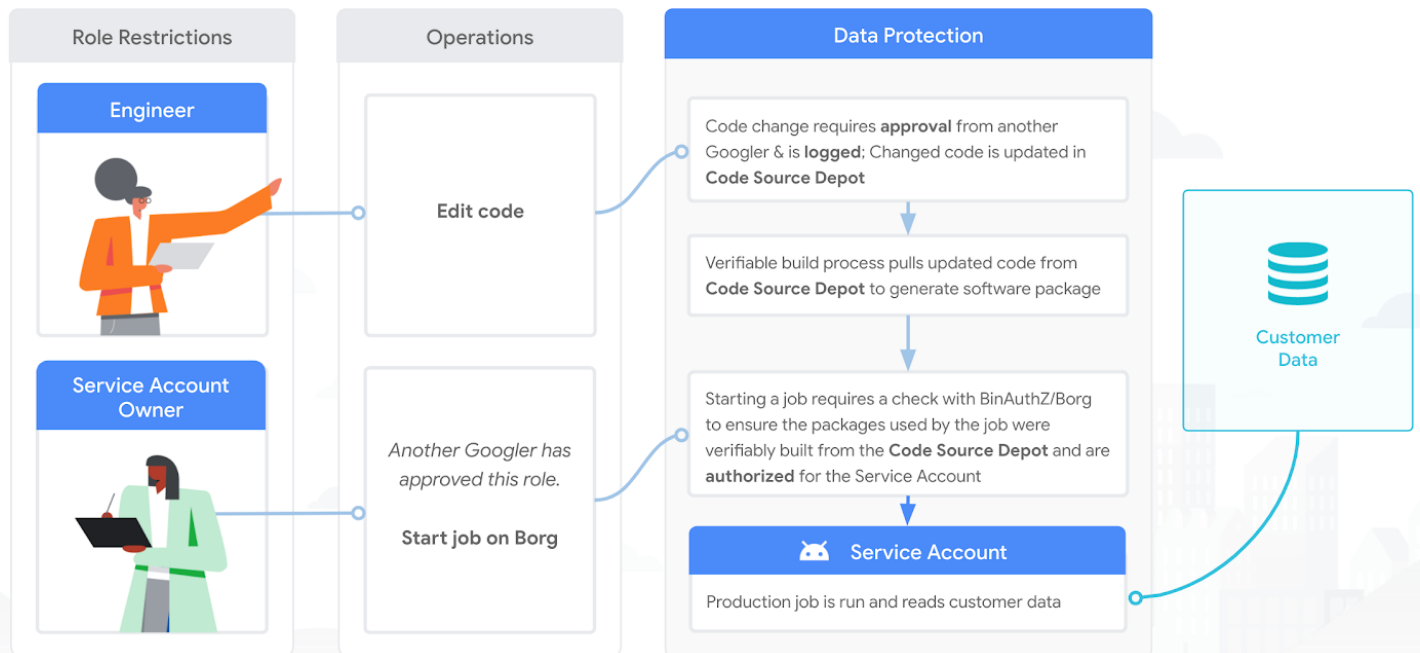
BeyondCorp helps ensure employee credentials are not stolen or misused and that connecting devices are not compromised. By shifting access controls from the network perimeter to the context of individual users and devices, BeyondCorp also allows Googlers to work more securely from virtually any location without the need for a traditional VPN. [Learn more.](#)

Review and authorize all production software

Our infrastructure is containerized, using a cluster management system called [Borg](#). Binary Authorization on Borg (BinAuthz/Borg) is an internal deploy-time enforcement check which ensures that production software and configuration deployed at Google is properly reviewed and authorized. BinAuthz/Borg ensures code and configuration deployments meet certain standards and will alert the service owners when these requirements are not met to trigger remediation. By requiring code to meet certain standards and change management practices before accessing user data, BinAuthz/Borg reduces the potential for a Googler acting alone (or whose account has been compromised) from accessing user data programmatically. [Learn more.](#)

Google Cloud Privileged Access Management

Google Engineer access path has robust controls to protect customer data



Detect and remediate unauthorized access

Detailed logging of all access

Google has detailed logging for any data access and code changes. In general, there are three types of logging:

1. **Customer logs** - Visible to customer in Cloud Audit Logs
2. **Administrative access logs** - Visible to customer using Access Transparency (AXT)
3. **Deployment integrity logs** - Not visible to customers, but exceptions are closely monitored internally by a centralized security team dedicated to auditing access to customer data. Exception monitoring protects sensitive data and boosts production reliability by ensuring that unreviewed or unsubmitted source code does not run in privileged environments, whether accidentally or as a result of a deliberate attack.

Incident Detection and Response

As detailed above, Google deploys robust internal controls to prevent the exfiltration of user data and abuse of privileged access to storage and production services. However, when violations occur, Google uses expert internal investigation teams and a suite of manual and automated controls that combine machine learning, advanced data processing pipelines, and threat intelligence to quickly detect and respond to incidents.

Signal development

The core of Google's detection and response capabilities is threat intelligence, fortified by continuous analysis of past incidents, network traffic, internal data, system access logs, anomalous behavior patterns, the results of offensive security exercises, and many more proprietary alerts. This data is analyzed by dedicated teams that produce a dynamic database of signals, or threat indicators, that touch on all parts of Google. Engineering teams use threat indicators to develop specialized detection systems to monitor internal systems for malicious activity, alert appropriate staff, and implement automated responses (e.g. revoking access to a resource).

Threat Detection

Threats are primarily detected by scanning logs and matching log entries to threat indicators. As a result of strong authentication, Google can delineate between human, service, and humans impersonating service events in the logs to prioritize investigations into actual human access. Activities that involve the access of user data, source code, and sensitive information are logged and a business justification or exception is required. Threats can include an individual attempting to take unilateral action on sensitive systems or accessing user data without a valid business reason. Violations of these and other policies have defined alerting procedures.

Incident Investigation

When policy violations are detected, security teams, separate from the core engineering and operations teams, provide independent oversight and will conduct an initial investigation. The teams will review details of the incident and determine if the access was intentional, unintentional, accidental, caused by a bug or misconfiguration, or the result of inadequate controls (e.g. an external attacker successfully stealing and using the credentials of a compromised employee). If the access is unintentional or accidental (e.g. a Googler was unaware of, or mistakenly violated, access protocols) the teams can take immediate steps to remediate the issue (e.g. recovering intellectual property). However, if malicious behavior is suspected, the incident is escalated and additional information may be collected, including data and system access logs, to determine the scope and impact of the incident. Depending on the results of that inquiry, incidents are submitted for additional investigation, documented and resolved, or, in extreme cases, referred to outside authorities or law enforcement.

Remediation

Past incidents are used to identify and close vulnerabilities and improve detection capabilities. All incidents are documented and metadata is extracted to identify specific tactics, techniques, and procedures for each exploit. That data is used to develop new threat indicators, reinforce existing protections, or make feature requests for security improvements.



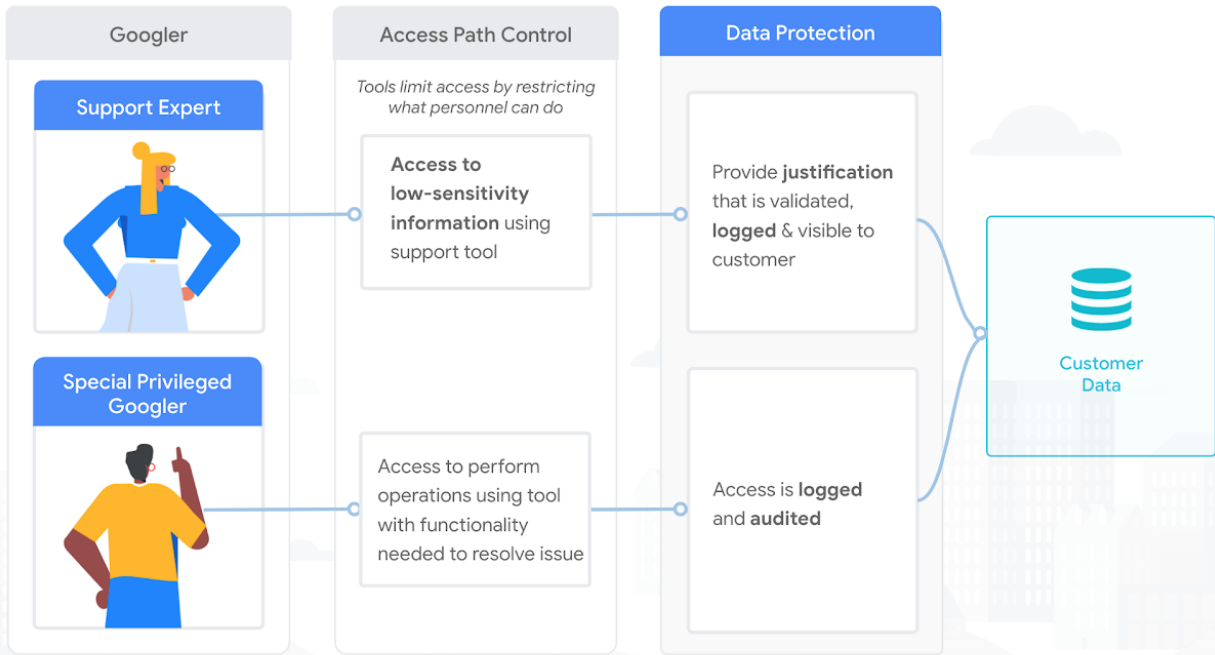
Google personnel by default can't access customer data

By default Google personnel do not have access to Google Cloud customer data. Google personnel must be a member of relevant access control lists to gain access to sensitive data. In addition, personnel must read and acknowledge Google's data access policies regularly.

Gaining access to control lists is only a precondition of actually accessing data. To access the data, Google personnel must use a trusted device and multi-factor authentication via [Security Keys](#), which also minimizes the risk of credentials being phished. From this trusted device, they are able to access tooling that evaluates the justification provided (ie. support ticket, issue ID, etc), the user's role and context to determine if they can obtain access to the customer data. Some tools will require Google personnel to obtain authorization from another qualified Googler to access the data. For customers who have enrolled in Access Approvals, access to their data will require customer approval prior to CSR and TSE access.

Google Cloud Privileged Access Management

Access happens through secure tools with checks and balances for protecting customer data



Google has multiple layers of support and engineering, each with different access privileges granted to them:

- **Customer Service Representatives (CSRs)** provide frontline support via phone and email
- **Technical Support Engineers (TSEs)** assist in dealing with more complex issues that our CSRs are unable to solve
- **Software Engineers (SWEs)** work on building the actual service (e.g. Google Cloud Storage). Together with the SREs they deal with the most complex issues that originate from genuine engineering faults
- **Site Reliability Engineers (SREs)** work with the Software Engineering teams and are focused on improving service reliability

When Google works with third parties, like customer support vendors, to provide Google services, we conduct an assessment to ensure they provide the appropriate level of security and privacy needed to receive access to your business' data. Google Cloud publishes a [list of all subprocessors](#) used to assist in providing the service.



Reasons for Google personnel to access customer data

Although Google Cloud has been engineered to automate, minimize, or eliminate the need for Google personnel to access customer data, there are still some cases where Google personnel may access customer data:

Customer initiated support

Google personnel access to customer data in Access Transparency enabled services are overwhelmingly the result of customer initiated events such as contacting customer support. The driver of this access is the customer contacting a Technical Support Experts to resolve an issue, such as losing access to a bucket. Support experts only have access to low-sensitively data to resolve customer issues when providing support (e.g. bucket names).

Outage or tool failure

During outages or tool failures, Google personnel can access customer data to perform a backup or recovery as needed. In these situations, Google personnel use tooling that can directly access customer data to maximize efficiency and resolve the issue in a timely manner. This tooling logs these accesses and the justifications provided by the engineers. These accesses are also audited by the Google security response team and logged. Google will still generate Access Transparency logs that are visible to the customer during an outage if access to specific customers' data occurs. Engineers are not able to bypass the ACL for the resource in question under any circumstance, however they will be able to access the data without customer approval.

Third party legal requests

Third party legal requests are rare ([learn more](#)), and only the legal team can generate a valid legal access justification. The legal team will review the request to ensure it meets legal requirements and Google policies, provide notification to the customer when legally permitted, and consider objections to disclose the data to the extent that the law allows. See [Government requests for cloud customer data](#) whitepaper.

Google initiated review

Google initiated reviews are also very rare, but when they do occur, they are to ensure that customer data is safe and secure, and hasn't been compromised. The drivers of these reviews are for security, fraud, abuse, or compliance purposes. For example, if automated bitcoin mining detectors detect that a VM is being used for bitcoin mining. Google looks into it and confirms that malware on a VM device is exhausting the VM's capacity. Google removes the malware so that VM usage returns to normal.

Customer tools to monitor and control Google's access to data

Google's committed to being fully transparent about Google personnel accessing customer data and to building tools to support this commitment.

On Google Cloud Platform, customers can control access to their data by Google personnel. Access Approval allows customers to require explicit approval before Google personnel access data or configurations on Google Cloud Platform (unless those accesses are necessary to resolve a current service disruption or security incident or required by law). This complements the visibility provided by Access Transparency, which generates near real-time logs when Google administrators interact with customer data, including the office location of the administrator and the reason for the access.

Through Assured Workloads for Government (Preview), customers are able to enforce specific attributes of Google administrators accessing their data—including the geographic region from which they are operating and other compliance-relevant attributes.

Key Access Justifications works alongside External Key Manager to provide a detailed justification each time a customer key is requested to decrypt data, along with a mechanism for a customer to approve or deny key access, using an automated policy that they set.

These products and features provide the customer with a large range of options to have visibility and control over Google's access to their data. As a result, customers are the ultimate arbiter of access to their data—a level of control not available from any other cloud provider.

Google continues to advance its data access control suite, so some of the products below are mature while others are new and currently in preview.

Access Transparency

Access Transparency (AXT) represents Google's long-term commitment to security and transparency by providing customers with logs that capture the actions Google personnel take when accessing customer content. Cloud Audit Logs help answer "who did what, where, and when?" in Google Cloud projects. While Cloud Audit Logs provides these logs about the actions taken by members within your own organization, Access Transparency provides logs of the actions taken by Google personnel. [Learn more.](#)

Access Approvals

Access Approvals gives customers the ability to pre-approve accesses to their data by Google personnel before those accesses occur, unless those accesses occur as part of a legal access, outage or security

incident. Access Approvals is an additional control built on top of our existing controls which apply to all internal users. [Learn more.](#)

Cloud External Key Manager (Cloud EKM)

Cloud External Key Manager (Cloud EKM) lets customers encrypt data with encryption keys that are stored and managed in a third-party key management system that's deployed outside Google's infrastructure. External Key Manager allows customers to maintain separation between data at rest and encryption keys while still leveraging the power of cloud for compute and analytics. [Learn more.](#)

Confidential VMs

Confidential VMs leverage the Secure Encrypted Virtualization (SEV) feature of 2nd Gen [AMD](#) EPYC™ CPUs, radically increasing isolation from Google insiders, and potential infrastructure compromises. Data is kept encrypted in-use when processed on GCP (in addition to the existing capabilities to keep data encrypted at-rest and in-transit). [Learn more.](#)

Key Access Justifications - Preview

Key Access Justifications gives customers a justification every time their externally hosted keys are used to decrypt data. It works in conjunction with Cloud External Key Manager to greatly advance the control that GCP customers have over their data. What this means is that there is no way for Google to decrypt customer data-at-rest without customer approval, which our customers can withhold for any reason. [Learn more.](#)

Assured Workloads for Government - Preview

Assured Workloads for Government is a set of data location and administrative access controls automatically enforced on a customer's Google Cloud environment. Government customers can use Assured Workloads to control where content is stored and processed, and ensure compliance with regulatory requirements. [Learn more.](#)

