

Security Whitepaper: Google Apps Messaging and Collaboration Products

Security Whitepaper

Google Apps Messaging and Collaboration Products

Table of Contents

Introduction.....	2
Overview.....	3
Google Corporate Security Policies.....	3
Organisational Security.....	3
Asset Classification and Control.....	4
Physical and Environmental Security.....	6
Operational Security.....	7
Access Control.....	9
System Development and Maintenance.....	10
Disaster Recovery and Business Continuity.....	12
Regulatory Compliance.....	12
Security Feature Customisations.....	13
Conclusion.....	14

For more information on Google Apps, visit www.google.com/a

Introduction

The security of online services is a topic of increasing interest to enterprises as the number of third party hosted service offerings has expanded in recent years. The emergence of various “cloud computing” concepts and definitions has highlighted not only questions about data ownership and protection, but also how various vendors of cloud computing technologies build and implement their services. Security experts, end-users and enterprises alike are all considering the security implications of the cloud computing model.

Google Apps (comprising Gmail, Google Calendar, Google Docs, and other web applications) provide familiar, easy to use products and services for business settings. These services, characterised by redundant computing environments and dynamic resource allocation, enable customers to access their data virtually anytime and anywhere from Internet-capable devices. This computing environment — often called the “cloud” — allows CPU, memory and storage resources to be shared and utilised by many customers while also offering security benefits.

Google provides cloud services reliably due to its experience with operating its own business, as well as its core services like Google Search, in a similar manner. The security controls that isolate data during processing in the cloud were developed alongside the core technology from the beginning. Security is thus a key component of each of our cloud computing elements, such as compartmentalisation, server assignment, data storage, and processing.




Figure 1: Google's multi-tenant, distributed environment

This paper will explain the ways Google creates a security-based platform for offering its Google Apps products, covering topics like information security, physical security and operational security. This exploration will demonstrate how security is an integral component of Google's cloud computing system, as well as a core element of Google's design and development processes. The policies described in this paper are detailed as of the time of authorship. Some of the specifics may change over time as we regularly innovate with new features and products within Google Apps.

Overview

Google's security vision is formed around a multi-layered security strategy that provides controls at multiple levels of data storage, access, and transfer. The strategy includes the following ten components:

- Google corporate security policies
- Organisational security
- Asset classification and control
- Personnel security
- Physical and environmental security
- Operational security
- Access control
- Systems development and maintenance
- Disaster recovery and business continuity
- Regulatory compliance

Google Corporate Security Policies

Google is committed to the security of all information stored on its computer systems. This commitment is outlined in the Google Code of Conduct, which is posted on Google's website at <http://investor.google.com/corporate/code-of-conduct.html>. Google's Security Philosophy is also outlined at the following page: <http://www.google.com/intl/en/corporate/security.html>.

The foundation of Google's commitment to security is its set of security policies that cover physical, account, data, corporate services, network and computer systems, applications services, systems services, change management, incident response, and data centre security. These policies are reviewed on a regular basis to help ensure their continued effectiveness and accuracy.

In addition to these security policies, with which all persons employed by Google must comply, employees are also given a Security Primer. This document outlines the most important aspects of information security policy, such as the safe use of the Internet, working from remote locations safely, and how to label and handle sensitive data. Additional guidance is routinely given on policy topics of interest, especially in areas of emerging technology, such as the safe use of mobile devices and peer-to-peer software. These supplemental policy documents are written with Google's core ideology of simplicity in mind, knowing that written policies are only effective if their information is consumed.

Organisational Security

Information Security

Google employs a full-time Information Security Team, embedded in the Google Software Engineering and Operations organisation, that is comprised of some of the world's foremost experts in information, application, and network security. This team is responsible for maintaining the company's perimeter defense systems, developing security review processes, and building customised security infrastructure. It also has a key role in the development, documentation, and implementation of Google's security policies and standards.

Specifically, Google's Information Security staff undertakes the following activities:

- Reviews security plans for Google's networks, systems, and services using a rigorous, multi-phase process
- Conducts security design and implementation-level reviews
- Provides ongoing consultation on security risks associated with a given project and possible solutions to security concerns
- Monitors for suspicious activity on Google's networks, and follows formal incident response processes to quickly recognise, analyse, and remediate information security threats

- Drives compliance with established policies through routine security evaluations and internal audits
- Develops and delivers training for employees on complying with Google security policy, especially in the areas of data security and secure programming
- Engages outside security experts to conduct regular security assessments of its infrastructure and applications
- Runs a vulnerability management program to help discover problem areas on the networks, and help ensure known issues that need to be remediated are addressed within expected timeliness

The Information Security Team also works publicly with the security community outside of Google:

- Publishing new techniques for secure programming to remain current with cutting-edge security trends and issues
- Working with software vendors and maintainers to identify and remediate vulnerabilities in third-party open and closed source software
- Developing worldwide privacy standards
- Providing educational materials for the public on information security issues such as browser security (<http://code.google.com/p/browsersec/wiki/Main>)
- Participating in, and organising, open source projects such as skipfish, a fully automated, active web application security reconnaissance tool (<http://code.google.com/p/skipfish>)
- Building training curricula for top universities
- Running and participating in academic conferences

A list of Security and Privacy related publications by Google employees can be found at <http://research.google.com/pubs/SecurityCryptographyandPrivacy.html>.

Global Internal Audit and Global Compliance

In addition to a full-time information security team, Google also maintains several functions focused on complying with statutory and regulatory compliance worldwide. Google has a Global Compliance function that is responsible for legal and regulatory compliance as well as a Global Internal Audit function responsible for review and auditing adherence to said compliance requirements, such as Sarbanes-Oxley and Payment Card Industry standards (PCI).

Physical Security

Google maintains a global team of staff, headquartered in the United States, dedicated to the physical security of Google's office and data centre facilities. Our security officers are highly qualified and have training in protecting similar high security infrastructure type environments.

Asset Classification and Control

Information Access

Google has extensive controls and practices to protect the security of customer information.

Google applications run in a multi-tenant, distributed environment. Rather than segregating each customer's data onto a single machine or set of machines, Google Apps data from all Google customers (consumers, business, and even Google's own data) is distributed amongst a shared infrastructure composed of Google's many homogeneous machines and located across Google's many data centres.

Google Apps uses a distributed file system designed to store large amounts of data across large numbers of computers. Structured data is then stored in a large distributed database built on top of the file system. Data is chunked and replicated over multiple systems such that no one system is a single point of failure. Data chunks are given random file names and are not stored in clear text so they are not humanly readable. For more information please download the abstract at <http://labs.google.com/papers/gfs.html>

The layers of the Google application and storage stack require that requests coming from other components are authenticated and authorised. Service-to-service authentication is based on a security protocol that relies on a Google system to broker authenticated channels between application services. In turn, trust between instances of this authentication broker is derived from x509 host certificates that are issued to each Google production host by a Google-internal certificate authority.

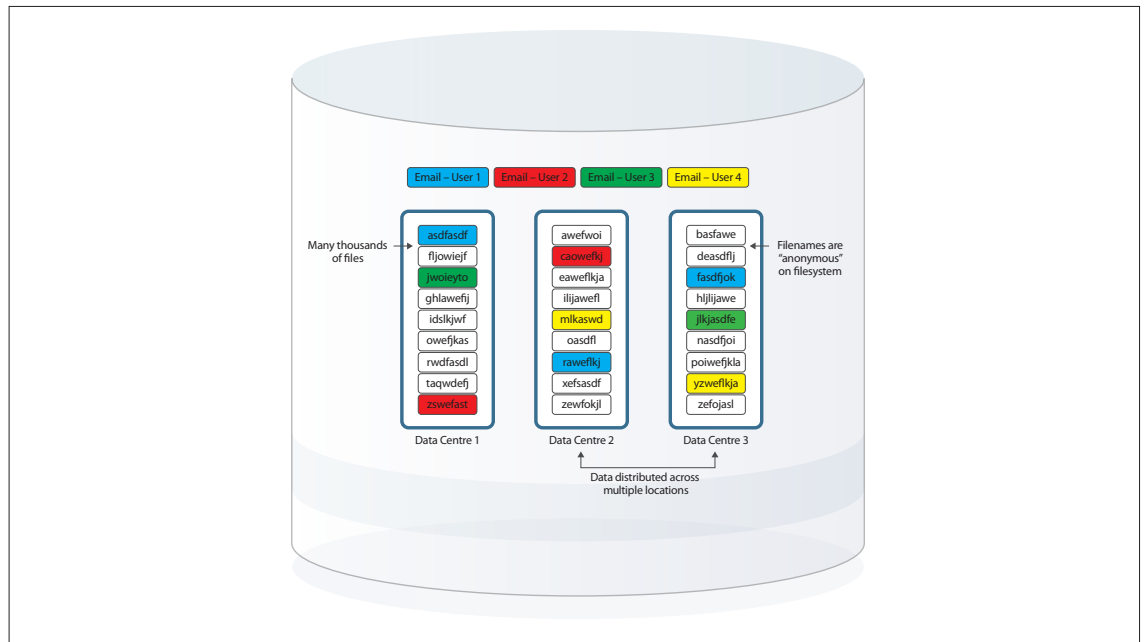


Figure 2: Google File System (GFS) architecture

For example, a Gmail web frontend service would make a remote procedure call to a Gmail backend service to request a message from a particular user's inbox. The Gmail backend would authenticate and process this request only if the requester is indeed a service running under a service identity that is allowed to access Gmail backends. The Gmail backend would in turn authenticate in order to access files in the Google distributed file system, and if successful, would only be granted access in accordance with file access control lists (ACLs).

Access by production application administrative engineers to production environments is similarly controlled. A centralised group and role management system is used to define and control engineers' access to production services, using an extension of the above-mentioned security protocol that authenticates engineers through the use of a personal x509 certificate that is issued to them.

Policy requires that administrative access to the production environment for debugging and maintenance purposes be based on secure shell (ssh) public key authenticated connections. For both scenarios, group memberships that grant access to production services or accounts are established on an as-needed basis.

The security controls described above rest on the foundation of the integrity of the Google production platform. This platform in turn is founded on:

- Physical security protections of Google's data centre environment
- Integrity of the Google production operating system environment
- Limited, as-needed system administrator (root) level access to production hosts granted to a specialised group of employees whose access is monitored

These aspects of Google's security practices are covered in more detail in subsequent sections of this document.

Deleted Data

After a Google Apps user or Google Apps administrator deletes a message, account, user, or domain, and confirms deletion of that item (e.g., empties the Trash), the data in question is removed and no longer accessible from that user's Google Apps interface.

The data is then deleted from Google's active servers and replication servers. Pointers to the data on Google's active and replication servers are removed. Dereferenced data will be overwritten with other customer data over time.

Media Disposal

When retired from Google's systems, disks containing customer information are subjected to a data destruction process before leaving Google's premises. First, policy requires the disk to be logically wiped by authorised individuals. The erasure consists of a full write of the drive with all zeroes (0x00) followed by a full read of the drive to ensure that the drive is blank.

Then, another authorised individual is required to perform a second inspection to confirm that the disk has been successfully wiped. These erase results are logged by the drive's serial number for tracking.

Finally, the erased drive is released to inventory for reuse and redeployment. If the drive cannot be erased due to hardware failure, it must be securely stored until it can be destroyed. Each facility is audited on a weekly basis to monitor compliance with the disk erase policy.

Personnel Security

Google employees are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards.

Upon hire, Google will verify an individual's education and previous employment, and perform internal and external reference checks. Where local labor law or statutory regulations permit, Google may also conduct criminal, credit, immigration, and security checks. The extent of background checks is dependent on the desired position.

Upon acceptance of employment at Google, all employees are required to execute a confidentiality agreement and must acknowledge receipt of and compliance with policies in Google's Employee Handbook. The confidentiality and privacy of customer information and data is emphasised in the handbook and during new employee orientation.

Employees are provided with security training as part of new hire orientation. In addition, each Google employee is required to read, understand, and take a training course on the company's Code of Conduct. The code outlines Google's expectation that every employee will conduct business lawfully, ethically, with integrity, and with respect for each other and the company's users, partners, and even competitors. The Google Code of Conduct is available to the public at <http://investor.google.com/corporate/code-of-conduct.html>.

Depending on an employee's job role, additional security training and policies may apply. Google employees handling customer data are required to complete necessary requirements in accordance with these policies. Training concerning customer data outlines the appropriate use of data in conjunction with business processes as well as the consequences of violations.

Every Google employee is responsible for communicating security and privacy issues to designated Google Security staff. The company provides confidential reporting mechanisms to ensure that employees can anonymously report any ethics violation they may witness.

Physical and Environmental Security

Security Controls

Google's data centres are geographically distributed and employ a variety of physical security measures. The technology and security mechanisms used in these facilities may vary depending on local conditions such as building location and regional risks. The standard physical security controls implemented at each Google data centre are composed of well-known technologies and follow generally accepted industry best practices: custom designed electronic card access control systems, alarm systems, interior and exterior cameras, and security guards. Access to areas where systems, or system components, are installed or stored are segregated from general office and public areas such as lobbies. The cameras and alarms for each of these areas are centrally monitored for suspicious activity, and the facilities are routinely patrolled by security guards who may use bicycles, Segways and T3 motion scooters.

Google's facilities use high resolution cameras with video analytics and other systems to detect and track intruders. Activity records and camera footage are kept for later review, should it become necessary. Additional security controls such as thermal imaging cameras, perimeter fences and biometrics may be used when necessary.

Access to all data centre facilities is restricted to authorised Google employees, approved visitors, and approved third parties whose job it is to operate the data centre. Google maintains a visitor access

policy and set of procedures stating that data centre managers must approve any visitors in advance for the specific internal areas they wish to visit. The visitor policy also applies to Google employees who do not normally have access to data centre facilities. Google audits who has access to its data centres on a quarterly basis to help ensure that only appropriate personnel have access to each floor.

Google restricts access to its data centres based on role, not position. As a result, even most senior executives at Google do not have access to Google data centres.

Environmental Controls

Google's computing clusters are architected with resiliency and redundancy in mind, helping minimise single points of failure and the impact of common equipment failures and environmental risks. Dual circuits, switches, networks, and other necessary devices are utilised to provide redundancy. Facilities infrastructure at the data centres has been designed to be robust, fault tolerant, and concurrently maintainable.

Power To support Google's continuous, 24x7 operations, Google data centre electrical power systems include redundant systems. A primary and alternate power source, each with equal capacity, is provided for every critical component in the data centre. Upon initial failure of the primary electrical power source — due to causes such as a utility brownout, blackout, over-voltage, under-voltage, or out-of-tolerance frequency condition — an uninterruptible power supply (UPS) is intended to provide power until the backup generators can take over. The diesel engine backup generators are capable of providing enough emergency electrical power to run the data centre at full capacity for a period of time.

Climate and temperature Air cooling is required to maintain a constant operating temperature for servers and other computing hardware. Cooling prevents overheating and reduces the possibility of service outage. Computer room air conditioning units are powered by both normal and emergency electrical systems.

Fire detection and suppression Automatic fire detection and suppression equipment helps prevent damage to computing hardware. The fire detection systems utilise heat, smoke, and water sensors located in the data centre ceilings and underneath the raised floor. In the event of fire or smoke, the detection system triggers audible and visible alarms in the affected zone, at the security operations console, and at the remote monitoring desk. Manually operated fire extinguishers are also located throughout the data centres. Data centre technicians receive training on fire prevention and incipient fire extinguishment, including the use of fire extinguishers.

More Information

More information and a video tour about Google's data centres can be found at <http://www.google.com/corporate/green/datacentres/summit.html>.

Operational Security

Malware Prevention

Malware poses a significant risk to today's IT environments. An effective malware attack can lead to account compromise, data theft, and possibly additional access to a network. Google takes these threats to its networks and its customers very seriously and uses a variety of methods to prevent, detect, and eradicate malware.

This strategy begins with infection prevention by using manual and automated scanners to scour Google's search index for websites that may be vehicles for malware or phishing. More information about this process is available at <http://googlewebmastercentral.blogspot.com/2008/10/malware-we-dont-need-no-stinking.html>. The blacklists produced by these scanning procedures have been incorporated into various web browsers and Google Toolbar to help protect Internet users from suspicious websites and sites that may have become compromised. These tools, available to the public, provide protection to Google employees as well.

Secondly, Google makes use of multiple anti-virus engines in Gmail, on servers, and on workstations to help catch malware that may be missed by anti-virus signatures. Support staff are trained to identify and eradicate malware that might infect the Google network, and they will escalate unusual cases through the incident response team.

Monitoring

Google's security monitoring program is focused on information gathered from internal network traffic, employee actions on systems, and outside knowledge of vulnerabilities.

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. This analysis is performed using a combination of open source and commercial tools for traffic capture and parsing. A proprietary correlation system built on top of Google technology also supports this analysis. Network analysis is supplemented by examining system logs to identify unusual behavior, such as unexpected activity in former employees' accounts or attempted access of customer data.

Google Security engineers place standing search alerts on public data repositories to look for security incidents that might affect the company's infrastructure. They actively review inbound security reports and monitor public mailing lists, blog posts, and web bulletin board systems. Automated network analysis helps determine when an unknown threat may exist and escalates to Google Security staff, and network analysis is supplemented by automated analysis of system logs.

Vulnerability Management

Google employs a full-time team that is dedicated to helping ensure that vulnerabilities are managed in a timely manner. The Google Security Team actively scans for security threats using commercial tools, intensive automated and manual penetration efforts, quality assurance (QA) processes, software security reviews, and external audits. The vulnerability management team is responsible for tracking and following up on vulnerabilities.

Once a legitimate vulnerability requiring remediation has been identified by the Security Team, it is logged, prioritised according to severity, and assigned an owner. The vulnerability management team tracks such issues and follows up frequently until they can verify that they have been remediated.

Google also maintains relationships and interfaces with members of the security research community to track reported issues in Google services and open source tools. More information about reporting security issues can be found at <http://www.google.com/intl/en/corporate/security.html>

Incident Management

Google has an incident management process for security events that may affect the confidentiality, integrity, or availability of its systems or data. This process specifies courses of action, procedures for notification, escalation, mitigation, and documentation. Google's security incident management program is structured around the NIST guidance on handling incidents (NIST SP 800-61).

Key staff are trained in forensics and handling evidence in preparation for an event, including the use of third party and proprietary tools. Testing of incident response plans is performed for key areas, such as systems that store sensitive customer information. These tests take into consideration a variety of scenarios, including insider threats and software vulnerabilities.

To help ensure the swift resolution of security incidents, the Google Security Team is available 24x7 to all employees. When an information security incident occurs, Google's Security staff responds by logging and prioritising the incident according to its severity. Events that directly impact customers are treated with the highest priority. An individual or team is dedicated to remediating the problem and enlisting the help of product and subject experts as appropriate. Other responsibilities are deferred until the issue is resolved.

Google Security engineers conduct post-mortem investigations when necessary to determine the root cause for single events, trends spanning multiple events over time, and to develop new strategies to help prevent recurrence of similar incidents.

Network Security

Google employs multiple layers of defense to help protect the network perimeter from external attacks. Only authorised services and protocols that meet Google's security requirements are permitted to traverse the company's network. Unauthorised packets are automatically dropped.

Google's network security strategy is composed of the following elements:

- Control of the size and make-up of the network perimeter. Enforcement of network segregation using industry standard firewall and ACL technology.

- Systematic management of network firewall and ACL rules that employs change management, peer review, and automated testing.
- Restricting access to networked devices to authorised personnel.
- Routing of all traffic through custom front-end servers that help detect and stop malicious requests.
- Create internal aggregation points to enable better monitoring.
- Examination of logs for exploitation of programming errors (e.g., cross-site scripting) and generating high priority alerts if an event is found.

Operating System Security

Designed in-house from the ground up, Google's production servers are based on a stripped and hardened version of Linux that has been customised to include only the components necessary to run Google applications, such as those services required to administer the system and serve user traffic. The system is designed for Google to be able to maintain control over the entire hardware and software stack and to help provide a secure application environment.

Google's production servers are built on a standard hardened operating system (OS), and security fixes are uniformly deployed to the company's entire infrastructure. This homogeneous environment is maintained by proprietary software that continually monitors systems for binary modifications. If a modification is found that differs from the standard Google image, the system is automatically returned to its official state. These automated, self-healing mechanisms are designed to enable Google to monitor and remediate destabilising events, receive notifications about incidents, and slow down potential compromise on the network.

Using a robust change management system to provide a centralised mechanism for registering, approving, and tracking changes that impact all systems, Google minimises the risks associated with making unauthorised modifications to the standard Google OS.

Access Control

Authentication Controls

Google requires the use of a unique User ID for each employee. This account is used to identify each person's activity on Google's network, including any access to employee or customer data. This unique account is used for every system at Google. Upon hire, an employee is assigned the User ID by Human Resources and is granted a default set of privileges described below. At the end of a person's employment, policy requires that the account's access to Google's network be disabled from within the HR system.

Where passwords or passphrases are employed for authentication (e.g., login to workstations), systems enforce Google's strong password policies, including password expiration, restrictions on password reuse, and sufficient password strength.

Google makes widespread use of two-factor authentication mechanisms, such as certificates and one-time password generators.

Authorisation Controls

Access rights and levels are based on an employee's job function and role, using the concepts of least-privilege and need-to-know to match access privileges to defined responsibilities.

Google employees are only granted a limited set of default permissions to access company resources, such as email, Google's internal portal, and HR information. Requests for additional access follow a formal process that involves a request and an approval from a data or system owner, manager, or other executives, as dictated by Google's security policies. Approvals are managed by workflow tools that maintain audit records of all changes. These tools control both the modification of authorisation settings and the approval process to ensure consistent application of the approval policies.

An employee's authorisation settings are used to control access to all resources, including data and systems for Google Apps products.

Accounting

Google's policy is to log administrative access to every Google production system and all data. These logs are reviewable by Google Security staff on an as-needed basis.

Systems Development and Maintenance

It is Google's policy to consider the security properties and implications of applications, systems, and services used or provided by Google throughout the entire project lifecycle.

Google's "Applications, Systems, and Services Security Policy" calls for teams and individuals to implement appropriate security measures in applications, systems, and services being developed, commensurate with identified security risks and concerns. The policy states that Google maintains a security team chartered with providing security-related guidance and risk-assessment.

Google employs a variety of measures to ensure that the software products and services Google offers to its users meet high standards of software security. This section outlines Google's current approach to software security; it may adapt and evolve in the future.

Security Consulting and Review

With regards to the design, development, deployment and operation of applications and services, the Google Security Team provides the following primary categories of consulting services to Google's Product and Engineering Teams:

- Security Design Reviews – design-level evaluations of a project's security risks and corresponding mitigating controls, as well as their appropriateness and efficacy.
- Implementation Security Reviews – implementation-level evaluation of code artifacts to assess their robustness against relevant security threats.
- Security Consulting – ongoing consultation on security risks associated with a given project and possible solutions to security concerns, often in the form of an exploration of the design space early in project life cycles.

Google recognises that many classes of security concerns arise at the product design level and therefore must be taken into consideration and addressed in the design phase of a product or service. Ensuring that such considerations are taken into account is the primary purpose of the Security Design Review. As such, the Security Design Review has the following objectives:

- Provide a high-level evaluation of the security risks associated with the project, based on an exploration of relevant threats.
- Equip the project's decision makers with the information necessary to make informed risk management decisions and integrate consideration of security into project objectives.
- Provide guidance on the choice and correct implementation of planned security controls, e.g., authentication protocols or encryption.
- Help ensure that the development team is adequately educated with regard to relevant classes of vulnerabilities, attack patterns, and appropriate mitigation strategies.

In cases where projects involve innovative features or technologies, it is the Security Team's responsibility to research and explore security threats, potential attack patterns, and technology-specific vulnerability classes related to such features and technologies.

Where appropriate, Google contracts with third party security consulting firms to complement the Google Security Team's skill set and to obtain independent third party review to validate in-house security reviews.

Security in the Context of Google's Software Lifecycle

Security is at the core of our design and development process. Google's Engineering organisation does not require Product Development teams to follow a specific software development process; rather, teams choose and implement processes that fit the project's needs. As such, a variety of software development processes are in use at Google, from Agile Software Development methodologies to more traditional, phased processes.

Google's security review processes are adapted to work within the chosen framework. That this can be done successfully hinges on Google's quality-driven engineering culture and a few requirements defined by Engineering management for project development processes:

- Peer-reviewed design documentation
- Adherence to coding style guidelines
- Peer code review
- Multi-layered security testing

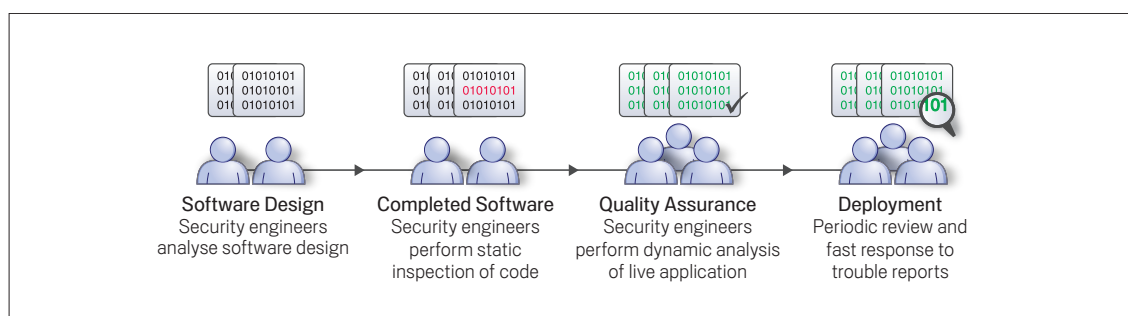


Figure 3: Google's system development and maintenance strategy

The above mandates embody Google's software engineering culture, where key objectives include software quality, robustness, and maintainability. While the primary goal of these mandates is to foster the creation of software artifacts that excel in all aspects of software quality, the Google Security Team's experience also suggests that they represent significant and scalable drivers toward reducing the incidence of security flaws and defects in software design:

- The existence of adequately detailed design documentation is a prerequisite of the security design review process, since in early project stages it is generally the only available artifact on which to base security evaluations.
- Many, if not most, classes of implementation-level security vulnerabilities are fundamentally no different from low-risk, common functional defects. Most implementation-level vulnerabilities are caused by fairly straightforward oversights on the developer's part.
- Given developers and code reviewers who are educated with respect to applicable vulnerability patterns and their avoidance, a peer review-based development culture that emphasises the creation of high-quality code is a very significant and scalable driver towards a secure code base.

The Google Security Team's software engineers collaborate with other engineers across Google on the development and vetting of reusable components designed and implemented to help software projects avoid certain classes of vulnerabilities. Examples include database access layers designed to be inherently robust against query-language injection vulnerabilities, or HTML templating frameworks with built-in defenses against cross-site-scripting vulnerabilities (such as the **Auto Escape** mechanism in the open-sourced **Google CTemplate** library).

Security Education

Recognising the importance of an engineering work force that is educated with respect to secure coding practices, the Google Security Team maintains an engineering outreach and education program that currently includes:

- Security training for new engineers.
- The creation and maintenance of extensive documentation on secure design and coding practices.
- Targeted, context-sensitive references to documentation and training material. For example, automated vulnerability testing tools provide engineers with references to training and background documentation related to specific bugs or classes of bugs flagged by the tool.
- Technical presentations on security-related topics.
- A security newsletter with engineering team-wide distribution that is intended to keep Google's engineering workforce abreast of new threats, attack patterns, mitigation techniques, security-related libraries and infrastructure, best practices and guidelines, etc.
- The Security Summit, a recurring Google-wide conference that brings together engineers from various teams at Google who work in security-related fields, and that offers in-depth technical presentations on security topics to Google Engineering at large.

Implementation-Level Security Testing and Review

Google employs a number of approaches to further reduce the incidence of implementation-level security vulnerabilities in its products and services:

- Implementation-level security reviews: Conducted by members of the Google Security Team, typically in later stages of product development, implementation-level security reviews aim to

validate that a software artifact has indeed been developed to be robust against relevant security threats. Such reviews typically consist of a re-evaluation of threats and countermeasures identified during security design review, targeted security reviews of security-critical code, selective code reviews to assess code quality from a security perspective, and targeted security testing.

- Automated testing for flaws in certain relevant vulnerability classes. We use both in-house developed tools and some commercially available tools for this testing.
- Security testing performed by Software Quality Engineers in the context of the project's overall software quality assessment and testing efforts.

Disaster Recovery and Business Continuity

To minimise service interruption due to hardware failure, natural disaster, or other catastrophes, Google implements a disaster recovery program at all of its data centres. This program includes multiple components to minimise the risk of any single point of failure, including the following:

- Data replication and backup: To help ensure availability in the event of a disaster, Google Apps data is replicated to multiple systems within a data centre, and also replicated to a secondary data centre.
- Google operates a geographically distributed set of data centres that is designed to maintain service continuity in the event of a disaster or other incident in a single region. High-speed connections between the data centres help ensure swift failover. Management of the data centres is also distributed to provide location-independent, around-the-clock coverage, and system administration.

In addition to the redundancy of data and regionally disparate data centres, Google also has a business continuity plan for its headquarters in Mountain View, CA. This plan accounts for major disasters, such as a seismic event or a public health crisis, and it assumes people and services may be unavailable for up to 30 days. This plan is designed to enable continued operations of our services for our customers. We conduct regular testing of our Disaster Recovery Plan.

Regulatory Compliance

Legal Information Access Process

Google follows standard legal processes in responding to third party requests for user information. Information can only be obtained by third parties through legal processes such as search warrants, court orders, subpoenas, through a statutory exemption, or through user consent. Upon receipt of a request for information disclosure, Google's Legal team reviews the request for compliance with applicable law. If the request is legally valid, it is Google's policy to notify the individual user or organisation whose information is being requested except in an emergency or where prohibited by law.

Privacy

Google maintains a strong privacy policy to help protect customer data. This policy is detailed at <http://www.google.com/a/help/intl/en/users/privacy.html> and is posted as part of every application within Google Apps. Read more about Google's privacy policies and practices at the Google Privacy Centre, located at <http://www.google.com/privacy.html>.

To put it simply, Google does not own customer data, and we believe it should stay that way.

Google adheres to the following principles regarding customer data:

- Google will not share data with others except as noted in the Google **Privacy Policy**.
- Google provides capabilities for customers to **take data with them** if they choose to use external services in conjunction with Google Apps or stop using Google services altogether.

User content is only scanned or indexed in the following cases to provide customers with a high-quality service:

- Some user data, such as email messages and documents, are scanned and indexed so users within a customer's domain can search for information in their own Google Apps accounts.
- Email is scanned so Google can perform spam filtering and virus detection.
- Email is scanned so Google can display contextually relevant advertising in some circumstances.
- Except when users choose to publish information publicly, Google Apps data is not part of the general google.com index.

Scanning and indexing procedures are automated and involve no human interaction. Google may also take down any content that violates the **Terms of Service** for Google Apps products.

Safe Harbor

Google adheres to the United States Safe Harbor Privacy Principles of Notice, Choice, Onward Transfer, Security, Data Integrity, Access and Enforcement, and is registered with the **U.S. Department of Commerce's Safe Harbor Program**.

SAS 70

Google has obtained a SAS 70 Type II attestation and will continue to seek similar attestation for the Google Apps messaging and collaboration products as well as for our security and compliance products, powered by Postini. A SAS 70 audit is an independent assessment by an outside audit firm that validates the subject company's adherence to its defined controls and confirms that these controls are operating effectively. When complete, the audit firm provides a report that details the company's compliance with these controls.

Security Feature Customisations

In addition to the various security controls described above that Google put in place to help protect the security and privacy of customer data, Google Apps also provides several additional security options that can be utilised by a customer's domain administrators. We are always working to give customers more choices when managing the security controls for their domain.

Single Sign-On (SSO)

Google Apps offers the Single Sign-On (SSO) service to customers with Premier, Education, and Partner Editions. Google Apps has a SAML-based SSO API that administrators can integrate into their LDAP, or other SSO system. This feature allows administrators to utilise the authentication mechanism of their choice, such as certificates, hardware tokens, biometrics, and other options.

Password Length and Strength

Administrators can set password length requirements for their domain users and view password strength indicators that help identify passwords that meet the length requirement but may still not be strong enough.

The password strength indicators can assess password strength in real-time and help administrators spot passwords that may become less secure over time based on emerging patterns of attacks.

Administrator-based Single Sign-Out

Administrators can reset a user's sign-in cookies to help prevent unauthorised access to their account. This will log out that user from all current web browser sessions and require new authentication the next time that user tries to access Google Apps.

Combined with the existing ability for administrators to reset user passwords, this feature to reset users' sign-in cookies improves security in the cloud in case of device theft or loss.

Secure Browser Connections (HTTPS)

Google Apps Premier and Education Editions offer domain administrators the ability to force all users in their domain to use Hypertext Transfer Protocol Secure (HTTPS) for services such as Gmail, Docs, Calendar, Sites, etc. Information sent via HTTPS is encrypted from the time it leaves Google until it is received by the recipients' computer.

Policy-enforced Secure Mail Transfer (TLS for SMTP)

With policy-enforced Transfer Layer Security (TLS) for Simple Mail Transfer Protocol (SMTP), administrators can set up policies designed for securely sending and receiving mail between specific domains. For example, an administrator could specify that all external mail sent by their accounting team members to their bank must be secured with TLS — or deferred if TLS is not possible. Similarly, an administrator could mandate a secure TLS connection between their domain and their outside legal counsel, auditors, or any other partners with whom employees may trade sensitive communications.

Archive Search

Google understands that archival services can assist customers in their compliance with various industry specific needs. By implementing Google Message Discovery, powered by Postini, customers can create a centralised and search-capable email repository for their organisation allowing for searching across the archive to locate and export email. The product can save and index all messages based on customer-defined retention policies. Customers can identify relevant messages, retain, search, and export the data to share as needed with outside vendors.

Conclusion

Google is committed to keeping the information stored on its computer systems safe and secure. Each of the ten components of Google's multi-layered security strategy is endorsed and defended throughout the organisation. Google Apps provides controls at each level of data storage, access, and transfer. Millions of organisations, including Google, run their businesses on Google Apps, and Google invests in that trust every day. With Google Apps, users can rest assured that Google values the privacy, confidentiality, integrity, and availability of their data.

