Google's Approach to IT Security

A Google White Paper

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.

Introduction

Google technologies that use cloud computing (including Gmail, Google Calendar, Google Docs, Google App Engine, Google Cloud Storage among others) provide famiiar, easy to use products and services for business and personal/consumer settings. These services enable users to access their data from Internet-capable devices. This common cloud computing environment allows CPU, memory and storage resources to be shared and utilized by many users while also offering security benefits.

Google provides these cloud services in a manner drawn from its experience with operating its own business, as well as its core services like Google Search. Security is a design component of each of Google's cloud computing elements, such as compartmentalization, server assignment, data storage, and processing.

This paper will explain the ways Google creates a platform for offering its cloud products, covering topics like information security, physical security and operational security.

The policies, procedures and technologies 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.

Overview

Google's security strategy provides controls at multiple levels of data storage, access, and transfer. The strategy includes the following ten components:

- Google corporate security policies
- Organizational security
- Data asset management
- Access control
- Personnel security
- Physical and environmental security
- Infrastructure security
- Systems and software development and maintenance
- Disaster recovery and business continuity

Google Corporate Security Policies

Google's commitment to security is outlined in both Google's Code of Conduct: **http://investor.google.com/corporate/code-of-conduct.html** and Google's Security Philosophy: **http://www.google.com/intl/en/about/corporate/company/security.html**.

These policies cover a wide array of security related topics ranging from general policies that every employee must comply with such as account, data, and physical security, along with more specialized policies covering internal applications and systems that employees are required to follow.

These security policies are periodically reviewed and updated. Employees are also required to receive regular security training on security topics such as the safe use of the Internet, working from remote locations safely, and how to label and handle sensitive data. Additional training is routinely given on policy topics of interest, including in areas of emerging technology, such as the safe use of mobile devices and social technologies.

Organizational Security

Google's security organization is broken down into several teams that focus on information security, global security auditing, and compliance, as well as physical security for protection of Google's hardware infrastructure. These teams work together to address Google's overall global computing environment.

Information Security Team

Google employs a full-time Information Security Team that is composed of over 250 experts in information, application, and network security. This team is responsible for maintaining the company's perimeter and internal defense systems, developing processes for secure development and security review, and building customized 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 multi-phase process
- Conducts security design and implementation-level reviews
- Provides ongoing consultation on security risks associated with a given project
- Monitors for suspicious activity on Google's networks, systems and applications, and follows formal incident response processes to recognize, analyze, and remedi ate information security threats
- Drives compliance with established policies through security evaluations and internal audits
- Develops and delivers training for employees on complying with Google security policy, including in the areas of data security and secure development
- Engages outside security experts to conduct periodic security assessments of Google's infrastructure and applications
- Runs a vulnerability management program to help discover problem areas on Google's networks, and participates in remediating known issues within expected time-lines

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

- Publishing new techniques for secure programming to remain current with security trends and issues
- Working with software vendors and maintainers to identify and remediate vulner abilities in third party open and closed source software
- 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 organizing, open source projects such as RatProxy, a web application security audit tool (**http://code.google.com/p/ratproxy/**)
- Building training curricula for top universities
- Running and participating in academic conferences
- Managing Google's Vulnerability Rewards Program (**http://www.google.com/about/corporate/company/rewardprogram.html**)

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

Global Internal Audit and Global Compliance Team

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 reviewing and auditing adherence to said compliance requirements, such as Sarbanes-Oxley and Payment Card Industry standards (PCI).

Physical Security Team

Google maintains a global team of staff, headquartered in the United States, dedicated to the physical security of Google's office and data center facilities. Google's security officers are qualified with training to protect high security enterprises with mission-critical infrastructures.

Data Asset Management

Google's data assets - comprising customer and end-user assets as well as corporate data assets - are managed under security policies and procedures. In addition to specific controls on how data is handled, all Google personnel handling data assets are also required to comply with the procedures and guidelines defined by the security policies.

Information Access

Google has 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 consumer and business customer data (as well as Google's own data) is distributed among a shared infrastructure composed of Google's many homogeneous machines and located across Google's data centers.

Google services store user data in a variety of distributed storage technologies for unstructured and structured data, such as Google File System (GFS), and distributed file systems evolved from GFS, such as BigTable.

The layers of the Google application and storage stack require that requests coming from other components are authenticated and authorized. Service-to-service authentication is based on a security protocol that relies on authentication infrastructure built into the Google production platform 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.

For example, a Google web application front-end might receive an end-user authenticated external request to display user data. The front-end in turn makes a remote procedure call to an application back-end to process the request. This remote procedure call is authenticated by the back-end, and will only be processed if the caller is authenticated as an authorized front-end application. If authorized, the application back-end will make a remote procedure call to a storage layer to retrieve the requested data. The storage layer again authenticates and authorizes the request, and will only process the request if the requester (the service back-end) is authenticated as authorized to access to the data store in question.

Access by production application administrative engineers to production environments is similarly controlled. A centralized 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 short-lived personal x509 certificates; issuance of personal certificates is in turn guarded by two-factor authentication.

Rather than segregating each customer's data onto a single machine or set of machines, Google consumer and business customer data (as well as Google's own data) is distributed among a shared infrastructure composed of Google's many homogeneous machines and located across Google's many data centers.

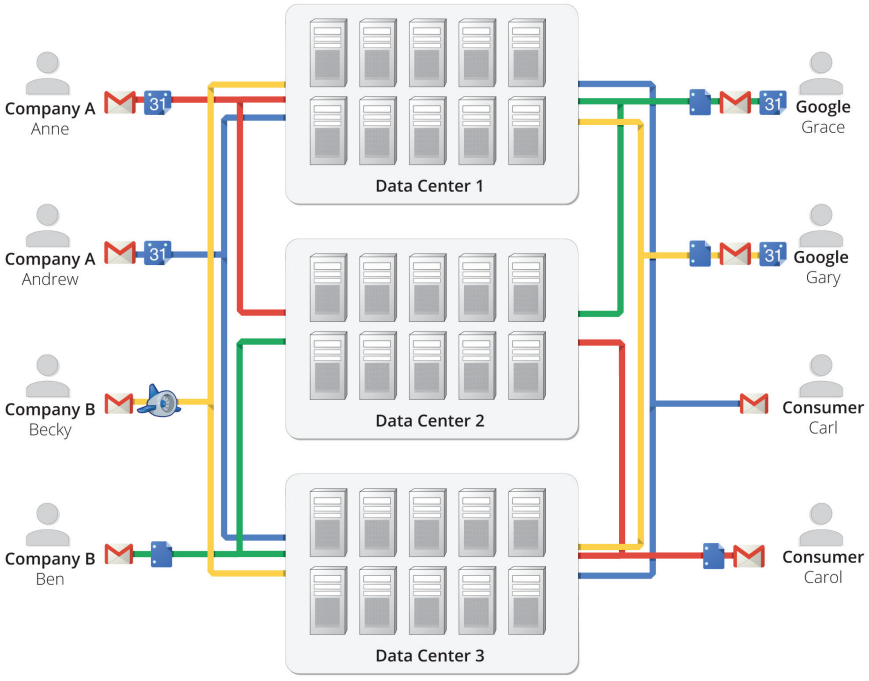

Figure 1: Google's Multi-tenant, distributed environment

Administrative access to the production environment for debugging and maintenance purposes is based on secure shell (SSH) connections. SSH connections into the production environment are authenticated using short-lived public-key certificates that are issued to individual administrative users; issuance of such certificates is in turn authenticated via two-factor authentication. All connections to the production environment are forced by network-level controls to pass through security proxies; these proxies provide centralized auditing of connections into the production environment, and allow for control over production access (e.g., in response to an incident such as suspected compromise of an administrative user's account). For both scenarios, group memberships that grant access to production services or accounts are established on an as-needed basis.

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 authorized individuals using a process approved by the Google Security Team.

Then, another authorized 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 physically destroyed. Each facility is audited on a weekly basis to monitor compliance with the disk erase policy.

Access Control

In order to secure Google's vast data assets, Google employs a number of authentication and authorization controls that are designed to protect against unauthorized access.

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, their account's access to Google's network is disabled from within the HR system.

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

Google makes widespread use of two-factor (2-step) authentication mechanisms, such as certificates and one-time password generators. Two-factor authentication is required for all access to production environments and resources through Google's Single Sign On system. Third party applications using Google Apps for Business can also use two-factor authentication.

Authorization 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 their email, and Google's internal portal. Employees are granted access to certain additional resources based on their specific job function. 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 authorization settings and the approval process to ensure consistent application of the approval policies.

An employee's authorization settings are used to control access to all resources, including data and systems for Google's cloud technologies and 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.

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 emphasized 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 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

Google has policies, procedures, and infrastructure to handle both physical security of its data centers as well as the environment from which the data centers operate.

Physical Security Controls

Google's data centers 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 center include the following: 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.

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. Additional security controls such as thermal imaging cameras, perimeter fences, and biometrics may be used on a risk basis.

Access to all data center facilities is restricted to authorized Google employees, approved visitors, and approved third parties whose job it is to operate the data center. Google maintains a visitor access policy and procedures stating that data center 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 center facilities. Google audits who has access to its data centers on a quarterly basis.

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

Environmental Controls

Google employs a set of controls to support its operating environment.

Power

To support Google's continuous, 24x7 operations, Google data center electrical

Google's computing clusters are architected with resiliency and redundancy in mind, helping minimize single points of failure and the impact of common equipment failures and environmental risks.

power systems include redundant systems. A primary and alternate power source, each with equal capacity, is provided for every critical component in the data center. 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 alternate power supply 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 center 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

Automated fire detection and suppression equipment helps prevent damage to computing hardware. The fire detection systems utilize heat, smoke, and water sensors located in the data center 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 centers. Data center 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 centers can be found at **http://www.google.com/corporate/green/datacenters/summit.html**.

Infrastructure Security

Google security policies provide a series of threat prevention and infrastructure management procedures.

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 address malware risks.

This strategy begins with manual and automated scanners that analyze Google's search index for websites that may be vehicles for malware or phishing. More information about this process is available at **http://goo.gl/eAcef**. 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, operate for Google employees as well. Secondly, Google makes use of anti-virus software and proprietary techniques in Gmail, on servers, and on workstations to address malware.

Monitoring

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

At multiple 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 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. Network analysis is supplemented by automated analysis of system logs.

Vulnerability Management

Google employs a team that has the responsibility to manage vulnerabilities in a timely manner. The Google Security Team scans for security threats using commercial and in-house-developed tools, 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, prioritized according to severity, and assigned an owner. The vulnerability management team tracks such issues and follows up until they can verify that the vulnerability has 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. Under Google's Vulnerability Reward Program (**http://www.google.com/about/company/rewardprogram.html**), security researches receive rewards for the submission of valid reports of security vulnerabilities in Google services. 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 and procedures for notification, escalation, mitigation, and documentation.

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 identified areas, such as systems that store sensitive customer information. These tests take into consideration a variety of scenarios, including insider threats and software vulnerabilities.

The Google Security Team is available 24x7 to all employees. When an information security incident occurs, Google's Security staff responds by logging and prioritizing the incident according to its severity. Events that directly impact customers are treated with the highest priority. An individual or team is assigned to remediating the problem and enlisting the help of product and subject experts as appropriate.

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 authorized services and protocols that meet Google's security requirements are permitted to traverse the company's network. Unauthorized

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.
- Management of network firewall and ACL rules that employs change management, peer review, and automated testing.
- Restricting access to networked devices to authorized personnel.
- Routing of all external traffic through custom front-end servers that help detect and stop malicious requests.
- Create internal aggregation points to support 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.

Transport Layer Security

Google provides many services that make use of the Hypertext Transfer Protocol Secure (HTTPS) for more secure browser connections. Services such as Gmail, Google Search, and Google+ support HTTPS by default for users who are signed into their Google Accounts. Information sent via HTTPS is encrypted from the time it leaves Google until it is received by the recipient's computer.

Operating System Security

Based on a proprietary design, Google's production servers are based on a version of Linux that has been customized 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 support a secure application environment.

Google's production servers are built on a standard 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 destabilizing events, receive notifications about incidents, and slow down potential compromise on the network. Using a change management system to provide a centralized mechanism for registering, approving, and tracking changes that impact all systems, Google reduces the risks associated with making unauthorized modifications to the standard Google OS.

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.

This section outlines Google's current approach to software security; it may adapt and evolve in the future.

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.

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 recognizes that many classes of security concerns arise at the product design level and therefore should be taken into consideration and addressed in the design phase of a product or service. 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 a key component of our design and development process. Google's Engineering organization 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. Engineering management has defined requirements for project development processes:

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

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 can reduce the incidence of security flaws and defects in software design and implementation:

- 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 classes of implementation-level security vulnerabilities are fundamentally no different from low-risk, common functional defects. Many 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 emphasizes the creation of high-quality code supports 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 [**http://google-ctemplate.googlecode.com/svn/trunk/doc/auto_escape.html**] in the open-source Google CTemplate library [**http://code.google.com/p/ctemplate/?redir=1**]).

Security Education

Recognizing the importance of an engineering workforce 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.
- In-depth training in application security for select engineers, with the goal of fostering the development of resident security experts on development project teams.
- The creation and maintenance of 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, 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 intended to reduce the incidence of implementation-level security vulnerabilities in its products and services:

- Implementation-level security reviews, which are conducted by members of the Google Security Team typically in later stages of product development, aim to validate that a software artifact has protection 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.

To minimize service interruption due to hardware failure, natural disaster, or other catastrophe, Google implements a disaster recovery program at all of its data centers. This program includes multiple components to minimize the risk of any single point of failure.

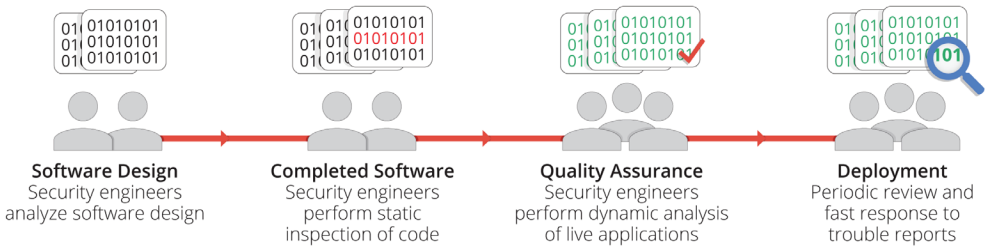

Figure 2: Google's System development and maintenance strategy

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

- Data replication and backup: Google application data is replicated to multiple systems within a data center, and in some cases also replicated to multiple data centers.
- Google operates a geographically distributed set of data centers 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 centers help to support swift failover. Management of the data centers 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 centers, 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.

Google conducts regular testing of its Disaster Recovery Plan. For example, During such tests, a disaster in a geographic location or region is simulated by taking IT systems and business and operational processes in that location off-line, and allowing such systems and processes to transfer to fail-over locations designated by the Disaster Recovery Plan. During the course of the test, it is verified that business and operations functions can operate at the fail-over location, and hidden/unknown dependencies on the off-line location are identified and logged for later remediation.

Summary

As described above, Google employs a multi-layered security strategy consisting of the ten core components illustrated in this paper that support a platform that is used by millions of organizations, including Google, to run their businesses on Google cloud technologies and products.