Polymail - Security Whitepaper

Last Revision Date: February 22, 2018

Describe in detail the mechanisms used to authenticate and authorize internal code and data access.

- Developers access code via Github, and are required to use two factor authentication.
- Developer's access to the code base is managed using Github's team management features https://github.com/features

Are data assets encrypted in transit and/or at rest? If so, what is the encryption strength?

- All data assets accessed by the client are encrypted using an HTTPS connection.
- Access to data stored in our GCP and AWS networks are strictly managed using IAM roles in both providers.
- Persistent Disks in GCP are encrypted by default
    - https://cloud.google.com/compute/docs/disks/

Does your server offer forward secrecy for clients that support it?

- Yes. This is a feature of Google Cloud Load Balancer.

Where is the SSL connection between the user and your application terminated?

- At the load balancer, from there on all traffic is in Google Cloud's Virtual Private Cloud
- https://cloud.google.com/vpc/docs/vpc

How do you protect against SQL injection?

- The only SQL database we use is Cloud Spanner. Spanner's client doesn't allow direct raw SQL queries. All queries are done with parameterized input.

When manually constructing queries, how do you preserve SQL injection protection?

- We construct our queries using the library provided by Google Cloud Spanner. All arguments are parameterized input.
    - https://cloud.google.com/spanner/docs/reference/libraries

Does your application load active content, such as scripts, applets or style sheets from third party servers?

- Yes. We load third party scripts and stylesheets from Google Analytics and Intercom.

Is all content you load to the DOM protocol-relative?

- Yes, and if the user reaches our site in HTTP, we redirect to them HTTPS.
- We proxy all remote images in emails for privacy and security.

- We don't allow HTTP content to ever load due to constraint in our CSP headers.

If your application supports authentication, are authentication cookies marked with the secure attribute?

- Yes, the authentication cookies are marked with the secure attribute

How do you guard against XSS?

- We sanitize all emails before they are rendered to our client.
- As an additional measure, we have a strict Content Security Policy that blocks inline Javascript, and third-party javascript not specified in our Content Security Policy.
    - https://developer.mozilla.org/en-US/docs/Web/HTTP/CSP

Does the application set a valid and appropriate content type and character set for each page (in the Content-Type HTTP header)?

- Yes, our Content-Type and character set are set for each page.

Can you ensure the client's data will not be removed or destroyed unless action to do so initiated by client -- unless otherwise documented in the agreement?

- Yes.

Are there mechanisms in place to ensure that all changes are adequately tested in a test environment prior to implementing in production?

- Our coding process requires all code changes to go through a peer review process
- We use continuous integration (CI) to regularly run unit tests and integration tests during the development process to test new features, and detect regressions on existing features.
- New code changes are required to pass all tests in our CI pipeline before they can be merged into our master branch.
- Additionally all code must compile before it can reach production.

Who has access to client data?

- Only our backend engineering team has access to our production databases on a case by case basis. We have strict policies in place to limit and restrict access to production data.

Are vendor default passwords removed, disabled or changed prior to placing the device or system into production?

- Yes, vendor default passwords are always removed and changed before reaching production.

Does your application use test or custom accounts during the development process? If so, how are those accounts removed and/or sanitized before the application is promoted to a production environment?

- Yes. The development test data is in a separate development environment that is not shared with production data.

Do you utilize "Cloud Computing"? If so, please explain.

- Yes, all our infrastructure is in Amazon Web Services (AWS) and Google Cloud Platform (GCP)
- AWS Services that we use:
    - S3
    - DynamoDb
    - Firehose
    - SQS
    - IAM
    - SNS
    - Cloudwatch
- GCP Services that we use:
    - Kubernetes Engine
    - Compute Engine
    - Cloud Spanner
    - BigQuery
    - Stackdriver Monitoring
    - IAM

What is the name of the cloud service provider?

- Google Cloud Platform (GCP)
- Amazon Web Service (AWS)

Does your Cloud Service Provider possess any of the following: ISO 27001 Certification, SOC-2 Audit Report, SSAE 16 Audit Report, FISMA Certification?

- Yes.
    - AWS: https://aws.amazon.com/compliance/
    - GCP: https://cloud.google.com/security/compliance

How is customer data handled between original source (e.g., Google, Salesforce) and the company's database?

- All API queries are made with HTTPS so data transmission is encrypted.

Geographically speaking, where will data be stored?

- AWS: us-east
- GCP: us-central

How will data be transmitted between your organization and the client or a third party?

- All data is transferred using SSL (HTTPS)

Please describe the usage of Transport Layer Security (TLS) for any network communication in your environment. Please provide version of TLS used. Is SSL v3 disabled?

- We use the Google Cloud Load Balancer's for TLS/SSL termination.
- Yes, SSL v3 is disabled, and not supported by Google Cloud Load Balancer

Have you deployed HTTP Strict Transport Security (HSTS) on your server?

- No.

How do you construct your Session IDs?

- We use a cryptographically secure pseudo random number generator provided by the Go standard library. https://golang.org/pkg/crypto/rand/
- We store a signed token as a cookie to indicate that the user is successfully logged in.

Does your application protect all state-changing actions against XSRF?

- All our api endpoints supports the CORS standard, and deny requests from unknown origins.
- We use POST, PATCH, and DELETE HTTP methods for any operation that changes data.
    - Our cookies have SameSite=Lax which prevents the cookie from being used in POST requests sent from a different domain.

Do sessions automatically time out after a specified period of inactivity?

- Web sessions expire between 1 to 2 days after a 30min period of inactivity, or immediately after the latter.
- Web session with *Remember Me* selected during login expire after 4 months
- iOS and macOS sessions expire after 1 year
    - Session tokens are stored securely on macOS Keychain to limit access from 3rd party apps.
- All tokens expire after 6mo of inactivity.

Does your application offer a "log out" button that not only terminates a session but invalidates the session ID?

- Yes.

Does your application offer a "Log out of all devices" button that terminates active sessions and invalidates the session ID?

- Yes.

Do you implement Horizontal Access Control such that users cannot access other user's data through URL hacking and such?

- Yes.

Do you implement Vertical Access Control such that users cannot access other user's data who are of higher privilege than them?

- Yes

Is there a formalized process in place to track, mitigate or accept risks identified through audits, risk assessments, monitoring, testing, etc.?

- Product Planning Process considers user feature requests and bugs discovered post-release.
- Product development process cadence is traditional agile process of one-week sprints, where sprint planning considers product roadmap and outstanding bugs and prioritizes based on importance, severity and level of effort.
- Bugs and feature requests (including tech debt retirement) are logged in JIRA for consideration, execution, and then historical reference.
- When a high severity issue is identified that requires attention and resolution outside of the normal sprint process, the team has an escalation process in place.

Please describe the usage of Secure Shell (SSH) for any network communication in your environment. Please provide the version of SSH used.

- We use SSH to login into the publicly accessible servers. We use ssh with certificates only, we have disabled ssh with user and password.

Describe your Cryptographic key management process including storage devices, key fragmentation, key officers, etc.

- We use Amazon IAM and Google Cloud IAM to manage and store our keys.

How often are inactive userID(s) deleted or disabled after they are no longer needed?

- We have an automatic cleanup process to remove deleted users.

Security requirements for your encryption keys:

- Our system uses IAM symetrics keys that are 256-bits
- People are asked to generate 4096 bit key pairs to ssh into the infrastructure

How are keys managed?

- For server side keys, we use Amazon IAM and Google Cloud IAM

Are clear text login passwords to Internet accessible systems are prohibited by your organization?

- Yes

Wireless network connections used to transmit confidential information are protected by technology stronger than WEP (e.g., WPA, IPSEC, SSL/TLS, etc.)

- Yes. The security of our wireless network connection is WPA2 Personal
- All confidential communications are done using SSL/TLS

Password storing, password policy, password reset, etc

- We try to use either Google SSO or Github SSO for all services. When that is not possible, we use 1Password to manage and rotate passwords for different services.

Are unique user IDs used for access to any system collecting, storing or processing data?

- Yes, user IDs are unique and enforced by our database's primary key.

Do you have a process to revoke customer data access from employees?

- Access to our databases in AWS and GCP are managed by IAM roles. This access can be revoked at anytime by one of our company administrators.

What is your disaster recovery and business continuity policy?

- Our services and databases are spread across multiple availability zones, with multiple replicas.
- All our services are managed via Kubernetes templates and Google Cloud Deployment templates that can be used to spin up additional zones quickly.

What are your internal procedures in the event of a data spill?

- Our internal procedures for data spills follow the Report, Analyze & Assess, Clean & Rectify, Document & Learn stepwise process.
- First, engineering identifies the owner(s) of the spilled data, and reports the level of the spill. Second, engineering analyzes the nature of the spilled data and potential access to determine impact. Based on this analysis, the engineering organization moves to remedy the spill using the relevant sanitization technologies. Lastly a post-mortem is performed, and all prior findings are documented to then be embedded in ongoing security training and education program.