Introduction

Groove.id is a SaaS identity platform that dramatically improves security, productivity and user experience by eliminating the use of passwords at work. We're aiming to solve the problems and challenges involved in securely managing employee identities.

To do that, we must make sure that your sensitive data is secure, and protecting it is our most important responsibility. We're committed to being transparent about our security practices and helping you understand our approach.

Infrastructure

Our service is hosted on infrastructure operated by Google (Learn More).

They have a global scale technical infrastructure designed to provide security through the entire information processing lifecycle. This infrastructure provides secure deployment of services, secure storage of data with end user privacy safeguards, secure communications between services, secure and private communication with customers over the internet, and safe operation by administrators.

The security of the infrastructure is designed in progressive layers starting from the physical security of data centers, continuing on to the security of the hardware and software that underlie the infrastructure, and finally, the technical constraints and processes in place to support operational security.

Google invests heavily in securing its infrastructure with many hundreds of engineers dedicated to security and privacy.

Data Encryption

All the data that we store for you is encrypted at rest and while in transit.

For some functions, we must manage secrets on your behalf. For example, in order to implement SAML for single sign on, we must digitally sign a special assertion documents. The secret key used to sign such an assertion is an example of a secret we manage for you.

To store secrets, we use a FIPS-140-2 Level 3 compliant hardware security manager. A key in the hardware security manager is used to decrypt secrets when they are needed. An intruder would not be able to derive the secrets merely with access to our data store. Access to the hardware security manager is strictly controlled and all access to it and thus to the secrets is carefully logged.

Each customer of Groove.id has their own unique sub-key for managing secrets, so even if a secret were stolen from one customer, it would not be usable for another customer.

Finally, individual secrets are encrypted using the XSalsa20 algorithm for encryption and Poly1305 for authentication (256-bit key, 192-bit nonce).

At various points in the application, we must ensure the integrity of data that the user manages for a brief time. For example, to prevent cross-site request-forgery, we must pass a state variable through the OAuth 2.0 process. When this is needed, we digitally sign the data using a per-customer key and HMAC-SHA256.

Secure Deployment

The Groove.id source code is stored in a central code repository. Making changes to the software requires the review and approval of at least one other member of the team. Our software infrastructure is immutable, meaning that neither we nor an attacker can modify it. Rather than modify

running systems, we destroy and replace systems to deploy new versions. No virtual machine in our infrastructure runs for more than twenty four hours before being destroyed.

Because deployments are automated, it is unusual for staff to access the production environment directly. This access is extremely rare and only a very few members of the team have access. All such access is audited and recorded.

Security vulnerabilities

We have selected technologies that are resistant to the kinds of vulnerabilities that have plagued software for years. For example, we use a modern, hosted, NoSQL data store (Google Cloud Datastore) rather than an SQL database, which means we don't have to worry about SQL-injection vulnerabilities.

We have also made architectural choices that make vulnerabilities more difficult to introduce. For example, the identity and privilege level of the remote user is threaded throughout the application, all the way to the datastore, which enforces access rules in testable, auditable place. The mandatory peer-code review process serves as a backstop against intentional or accidental vulnerabilities. Finally we use automated static analysis tools that alert us to potential security problems in the code, and those checks must pass in order for code to get deployed.

We have automated tools that monitor for security vulnerabilities in the third party code we rely on and automatically propose changes to fix them.

The responsibility for pathing known vulnerabilities in the operating system layer, virtualization layer, and hardware layers falls to Google and their mature vulnerability management practice. We have configured our infrastructure so that it is automatically patched against known

vulnerabilities within twenty-four hours of a patch being released. This is a side effect of our immutable infrastructure and our policy of replacing machines every twenty four hours.

Network security

We divide our systems into separate environments for development, staging and production. Each environment is an independent domain with respect to network access control, service account credentials, and secrets. No access to the production, staging or development environments is allowed except on known protocols and ports via our front-end load balancers.

All access to our services from user devices, or between our client software and our service is protected by TLS version 1.2 or higher. Our public endpoints, (for example, *api.groove.id*) receive an A rating from Qualys SSL Labs. Network traffic on the private network that connected the various microservices that constitute Groove.id is encrypted with TLS version 1.2 or higher.

Authorizing access

To minimize the risk of data exposure, Groove.id adheres to the principle of least privilege. Employees are only authorized to access data that they reasonably must handle in order to do their job. All internal systems require our employees to authenticate with unique user accounts.

All employees have successfully completed a comprehensive criminal background check.

Data Residency

By default, all customer data is located in Iowa, United States. At your request, we can alternately host your data will be hosted in the

Netherlands.

Employee Training

All employees complete mandatory security awareness training once per year. In addition to general resistance to online threats, we teach our staff to resist social engineering attacks through our support channels. All employees are trained in protecting the identities and confidential information of our clients. Although we do not generally handle protected health information (PHI), all employees are trained to identify and report any incidental contact with it.

Authentication

We eat our own dog food. Wherever possible, we use our own instance of Groove.id to authenticate to our internal systems. In the cases where we cannot use Groove.id ourselves, such as where doing so would create a circular dependency, we require the use of hardware backed multi-factor authentication such as U2F tokens.

Logging

We collect logs from all our servers. We routinely examine logs for suspicious activity and operational issues. We scrub logs of personal data and operational secrets before archiving them.

We also generate audit records of all your activity on our system and provide you a means to receive these records if you choose.

Business Continuity

All data that we store for you are regularly backed up. We regularly simulate the backup and recovery process to make sure it works smoothly. Copies of backups are stored in multiple data centers in

different regions and are encrypted in transit and at rest.