Data Protection, Security, and Privacy

How Papercurve Protects Your Data

papercurve

1.0

Table of Contents

A message from the CTO

Papercurve was built on a foundation of strong data security principles. We recognize how important it is to our customers to be good stewards of your data. Data security and privacy are core to the company culture. Our commitment to you is to always be evolving, improving and innovating to ensure your data is safe with Papercurve.

Thank you.

Antonio Salumbides
Chief Technology Officer

Overview

The following whitepaper documents Papercurve's system architect. It depicts general details on common installations of our platform for clients. The content you see in this whitepaper makes some assumptions:

1. A basic default installation (as opposed to a custom installation).
2. The installation and architecture are set up for handling a medium-sized organization.

Papercurve and team can customize installations upon request if there are specific needs. If you wish to do so and are an existing client, please contact your Customer Success representative. If you are currently not a client please feel free to drop us a message info@papercurve.com

The following document was updated in August 2020. For the latest version please contact us at info@papercurve.com

Infrastructure

Overview

Papercurve creates a new installation for each client. This allows Papercurve's operations team to silo a separate instance for each client. There is no sharing of resources between individual clients and scale can be dynamically controlled via auto-scaling groups to account for usage spikes. This helps ensure availability, security, and user experience.

The Papercurve application consists of many services and systems. In the following sections, you will get a brief overview of the generic setup that comes with each Papercurve instance. The architecture is designed to ensure checks and balances along the traffic pipeline.

Our architecture includes encryption from the top down and firewalls that are updated with policies often to ensure protection from unwanted access.

In addition, all of our infrastructure is built with our partners (Amazon AWS), and we also utilize third-party services to audit and monitor performance, security, and errors in order to maintain the best level of quality and availability.

Our infrastructure's default region is AWS' Canada Central 1. Papercurve is dynamic and we can host your install to an available region of your choice or we can scope out a multi-region installation if requested.

General Architecture

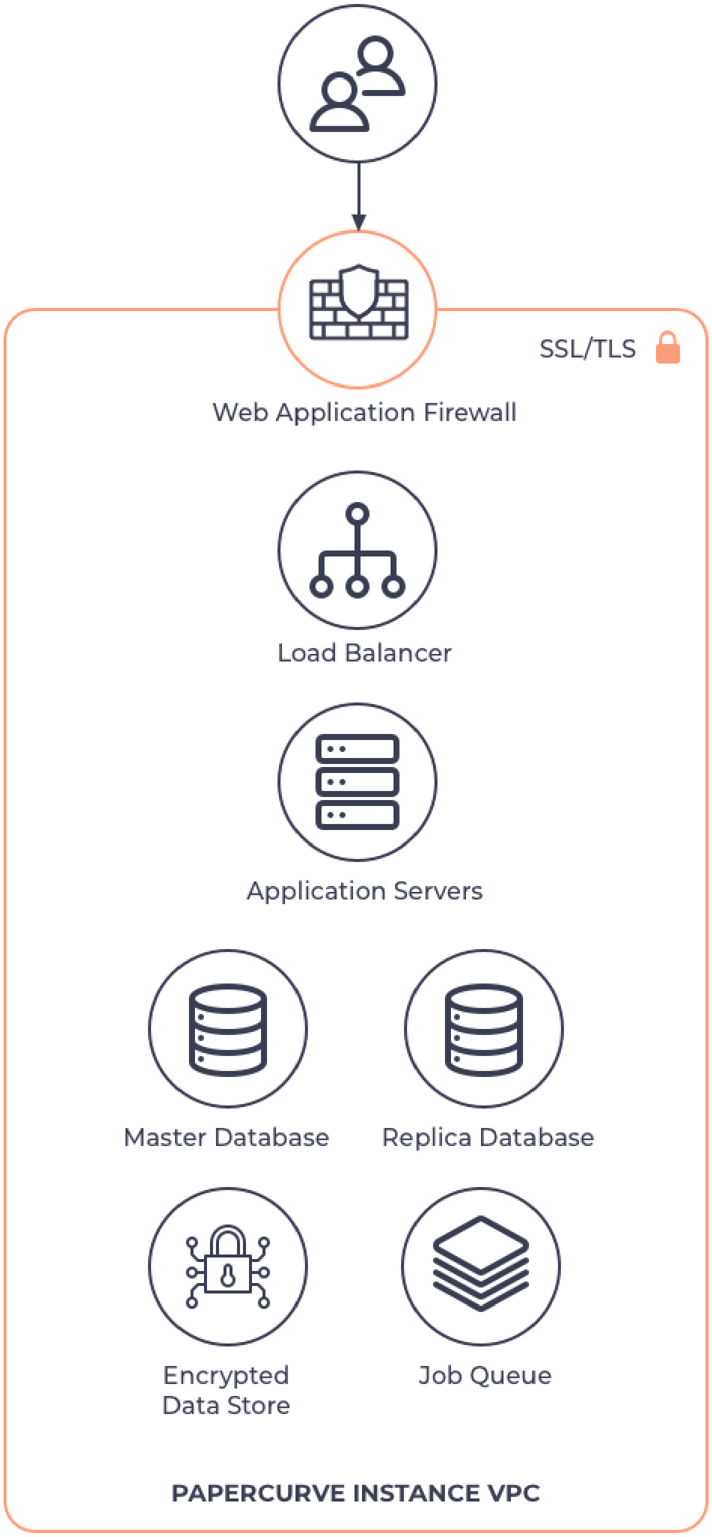

Gateway

Our gateway system consists of a web application firewall, encryption with SSL/TLS, a load balancer, and security routing rules. Before any request even being able to get through, our architecture ensures the traffic is safe with our web application firewall, it's encrypted with SHA-256 RSA encryption, and it has to pass security routing rules that deny any connection from continuing from untrusted sources.

Web Application Servers

At its core, our web application servers hold the majority of business logic that makes the Papercurve platform. Our application servers only hold the code that manages traffic and requests. It does not store files or customer data. This ensures that when traffic spikes occur we can manually or automatically scale our web servers to handle increased or decreased requests/responses.

The web application server is the glue that brings the other parts of the system together. User requests are received and data is either stored or retrieved from the database or the encrypted data store. In addition, we utilize a job queue to manage long-running tasks such as report generation, notifications, etc.

We partner with Amazon's AWS to utilize its Elastic Compute Cloud (EC2) to service our web application.

Databases

Our database store of choice is PostgreSQL. We use Amazon's RDS services to host a separate database unique to your instance. Our databases are all encrypted at rest and access controls are strictly tailored to ensure no unauthorized access is permitted. Our databases are backed up periodically (by default every 30 minutes) and we utilize a master/replica system for enterprise installations for maximum availability.

Data Storage

Any client files or assets uploaded to the Papercurve platform are stored on an encrypted data store.

To achieve this, we utilize several AWS services and tools. The first is a key management system to store and protect the encryption keys. These keys are rotated periodically. We have customized hardened access control policies, and Simple Storage System (S3) with encryption enabled.

Each Papercurve instance includes a separate store with unique encryption keys and periodic backups are performed every 30 minutes. In the event of any issues that may occur, the last backup would never be more than 30 minutes out of date.

Security

Network and Application Firewalls

Papercurve creates strict access control lists and deploys advanced rules to protect our network from attacks. Our web application firewall has been designed specifically to stop attacks. We constantly monitor and update it to ensure we are protected from new threats. As well, we box our environments per client adding the ability for us to add stricter rules.

Intrusion Detection

With the constantly evolving landscape of cyber attacks, Papercurve's web application firewall utilizes additional services to inform us of any possible signs of intrusion. If any incidents arise, a team of engineers known as Papercurve's incident response team gathers in a war room to assess the threat and begin gathering intelligence on the situation. The incident response team works closely with our customer success to ensure communication with the client is frequent and transparent.

Encryption

Papercurve ensures all communications are sent via encrypted channels. All sensitive information is stored encrypted and the distribution of materials is done through encrypted channels. In addition, our infrastructure consists of databases to store transactional persistent data and all of our databases enforce encryption at rest. These are also all enforced through our cloud provider and internal processes ensure we audit what we practice.

Access Control

At Papercurve, we strictly monitor, control, and minimize access to production systems. Members of the Papercurve Engineering team only have access to lower environments and utilize named accounts to access any production environments. In order for access to production environments, a member of the Papercurve team must go through a strict break glass process to be granted temporary access.

- Production environment access is restricted.
- When access is required, we proceed through our internal break glass access request process.
- All events are logged.

Two-factor Authentication

Papercurve enforces two-factor authentication (2FA) in order to access any resource of our infrastructure. This is enforced through our cloud platform provider and we ensure employees are trained when onboarded to understand the importance of 2FA. In addition, all Papercurve services and products provide 2FA for logins with advanced controls for administrators to revoke/reset access. Each instance can control whether 2FA is mandatory for all users, or disabled.

Availability

Papercurve partners with Amazon Web Services (AWS) for all resource and infrastructure needs. Our main resources are installed in data centers that are designated in the Canada Central region. AWS is certified SOC 2 Type II compliant and provides a high level of availability and security. On top of that Papercurve also implements 24/7 monitoring, logging, and internal processes to enhance AWS' offering.

Performance Monitoring

At Papercurve, we monitor all aspects such as service load, server performance, resource access on our products/platforms. Each server in rotation is monitored for resource usage such as CPU thresholds, memory usage, job/task handling, and service load. Our databases are monitored and tuned for optimal performance, availability, locking, and resource usage.

We utilize Amazon Web Services' CloudWatch and in addition, our site reliability engineers have internal tools such as Prometheus, Grafana, and Kibana configured to be alerted and debug/respond to any performance issues that occur.

Disaster Recovery

Papercurve exercises recovery simulations twice a year to ensure we are ready for unforeseen or unavoidable disasters. We have enacted procedures and designed infrastructure to minimize impact to our clients but are still preparing for events that can occur. In extension, we have preventative options that can lessen the impact of downtime by adding redundancy to client installations.

By default, we continually perform periodic backups of data every 30 minutes. We also have snapshots of databases occurring every 30 minutes. Depending on client requirements of mission-critical systems this interval can be increased. Papercurve is also flexible to design a hardened multi-regional custom installation of any client instance.

Security Incident Handling

In the event of a security incident, our incident response engineers are called into our war room to prioritize investigations and proceed with collecting data, sifting through logs and assessing the incident. Once the impact of the incident is determined, fixing and prevention are prioritized. At every available point, our incident response team is in constant communication with our Customer Success team to effectively communicate the issue transparently with the affected client(s).

Processing Integrity

Quality Assurance

At our core, Papercurve prioritizes Quality Assurance (QA) to built a product out customers can depend on. Our QA team tests every code change that is introduced in our products. We have multiple lower-level environments that are separate from production to conduct tests so data is never compromised or tainted. In addition, we utilize continuous integration to automate test runs prior to our QA team testing in order to catch issues early.

Unit Testing and Test-Driven Development

Testing starts at the very beginning of the code development lifecycle. Our engineers adhere to test drive development practices to ensure individual units of code are thoroughly tested. These tests ensure we keep a level of quality as our software matures and to be satisfied that we haven't made unwanted changes to existing code.

Automation testing

At Papercurve we believe that testing needs to be tackled from many different fronts. One such front is with our automated test suite. The Quality Assurance team at Papercurve creates new test cases as the application grows and automates these frontend tests using services at our partner's at Cypress.io to ensure each code commit is regulated with automatic testing.

These automated tests can be run on local machines as our engineers make changes, and as QA is vetting software. In addition, before any permanent merge commits to our codebase, these automated tests have to pass with a 100% success rate before allowing the new code to be added.

Continuous Integration

Continuous Integration (CI) is a development practice where developers integrate code into a shared repository frequently, preferably several times a day. Each integration can then be verified by an automated build and automated tests.

One of the key benefits of integrating regularly is that you can detect errors quickly and locate them more easily. As each change introduced is typically small, pinpointing the specific change that introduced a defect can be done quickly.

At Papercurve our CI framework consists of unit tests, frontend tests, and code analysis. Every code commit results in running our CI and the three units of testing need to pass specific thresholds before anything can be permanently merged to our master codebase.

Continuous Integration allows Papercurve to safely and efficiently evolve with the confidence that our changes are thoroughly tested.

Load Testing

Upon specific changes, Papercurve periodically deploys load testing to ensure that we know the type of load that our services can handle. We have specific metrics that we try and meet to ensure that there is plenty of capability for handling sustained concurrent loads on the platform. Since we do not share resources across clients we also have the benefit of ensuring that your traffic is only users of your instance.

Compliance

FDA 21 CFR Part 11 Compliance

Papercurve is FDA 21 CFR Part 11 compliant. Our service was engineered to ensure we comply with the regulation and it has been tested to ensure the individual taking action is the correct person. Each user in Papercurve is named and secured with a password. Each user has a role with specific permissions to ensure everyone has the right level of access.

When a person in Papercurve submits their approval for a piece of content, the approval meets FDA requirements for electronic signatures.

Papercurve records all key actions such as uploads, approvals, deletions and edits in a downloadable activity log. This serves as an automatic auditable record of everything that happens in the platform, who did it, and when.

Documents and videos are stored to meet document retention requirements. In an instance when content is deleted by a user, these deleted documents are archived and made available in the event of an audit.

GDPR

Papercurve takes significant steps to align with the General Data Protection Regulation (GDPR):

- Improved security procedures.
- Employee training on privacy and security best practices.
- Data collection assessment.
- Vendor risk assessment.
- Data breach incident response plans.