Stripe Payments Company
Business Continuity Plan

January 2018

Table of Contents

1. Overview

Stripe Payments Company ("SPC" or the "Company") has established a Business Continuity Plan ("BCP") to provide its employees ("Employees") and executive management with guidance on how the organization guards against future disasters that could endanger its long-term health or the accomplishment of its primary mission.

SPC's BCP takes into account disasters that can occur on multiple geographic levels—local, regional, and national disasters like fires, earthquakes, or pandemic illness. SPC's BCP takes into consideration any potential disasters that would require recovery; it includes everything from technological viruses to terrorist attacks. The ultimate goal is to help expedite the recovery of SPC's critical functions and manpower following these types of disasters.

2. Scope

This BCP applies only to SPC employees. All Stripes have company-issued computer equipment which is taken home at the end of the business day. Access to SPC work product is only accessible through a VPN network; SPC's infrastructure recovery will likely rely on cloud-hosted software-as-a-service (SaaS) offerings such as Github.com which are accessible remotely and do not necessarily use servers provided by Amazon Web Services.

3. Measurement

The SPC executive team will verify compliance to this policy through various methods, including but not limited to, periodic walkthroughs, video monitoring, business tool reports, internal and external audits, and feedback to the policy owner. The plan, at a minimum, should be reviewed and updated on an annual basis. Testing against the BCP should also be, at a minimum, conducted on an annual basis.

4. Emergency

If the emergency situation appears to affect the main data repositories or if access to the data is prohibited, senior leadership such as the CISO or VP of Engineering will closely monitor the event, notifying SPC personnel on the BCP team as required for assisting in

damage assessment. Once access to the situation is permitted, an assessment of the damage is made to determine the estimated length of the outage. If the estimated outage is anticipated to be more than 24 hours, the CTO will notify all customers and company personnel. If the outage is estimated to be longer than 24 hours, then CTO will notify the Chairman of the BCP team for contingency planning.

The BCP team will consist of a representative from each of the following teams:

- Executive management
- Legal
- People Operations
- Finance
- Communications
- Engineering

5. Backup

In the initial stage of the backup phase, the goal is to resume processing critical applications. Processing will resume either at the main data repositories or at the designated emergency repositories, depending on the results of the assessment of damage to the Company's servers, buildings, or any other type of infrastructure. In the backup phase, the initial servers must support critical applications for up to 2 weeks. During this time, the BCP team assumes that systems will be operating in a degraded mode, but will attempt to ensure full functionality by 4 weeks. However, if the damaged area requires a longer period of maintenance, then further actions determined by the CTO will be taken.

SPC's primary infrastructure is located on virtual servers via a virtual private cloud on Amazon Web Services. Amazon Web Services' Disaster Recovery Whitepaper informs a part of SPC's disaster recovery infrastructure. Pieces of the whitepaper cover issues such as backup servers, restoration procedures, standby solutions, and multi-region deployments.

6. Notifications

In the case of a disaster, the CEO will have ultimate responsibility to notify employees, vendors, customers, and partners appropriately. The CEO, along with the communications team, legal team, and executive team, will follow pre-established external media policy.

7. Recovery

The time required for recovery of the functional area and the eventual restoration of normal processing depends on the damage caused by the disaster. The time frame for recovery can vary from several days to several months. In either case, the recovery process begins immediately after the disaster. There is no parallel recovery that takes place at a designated hot site outside the Company's standard AWS region. The failure of an AWS region would take the Company completely offline until the infrastructure could be provisioned in a backup region a process that would conservatively take 24 hours to complete. The primary goal is to restore normal operations as soon as possible.

8. Contingency Location

In the event that SPC's physical office is unavailable, People Operations will direct employees to work from home until the situation has been resolved. The following protocol should be observed:

- Landlord to liaise/contact SPC's Office Manager ("OM") regarding the issue
- The OM will email SPC employees or visiting employees to advise them of the situation (i.e. not to attend work physically), through the email list 'sfoffice'; as a backup
- The OM will also post a message in the communal internal chat tool, Slack, to advise employees and key management in the US of the situation;
- The OM will keep up to date of the situation till such time as the building is a physical state where employees can return to work, and when the landlord advises as such, the OM will communicate this to all relevant SF-based employees or visiting employees; communication will take the form of an email to the 'sfoffice' list, and a post in #announce channel in Slack;

9. Critical Functions for Recovery

There are no physical servers in San Francisco for user facing operations and all SPC work product, including information pertaining to US clients, is designed and intended to be accessed through a VPN and not to be downloaded to devices/personal storage devices.

1. SPC makes heavy use of Amazon Web Services (AWS)—and use of AWS helps with disaster recovery by balancing between multiple data centers in the same geographic region. We do not actively utilise multiple geographic regions. They also improve redundancy against individual server failure and loss of network connectivity in our most commonly used configurations.
2. AWS Regions cover multiple separate and distinct data centers. The loss of a data center is expected to cause the correlated loss of multiple servers. To protect against this, all servers are expected to have redundant copies in multiple data centers, with alerting if this expectation is violated for critical stateful services, such as databases.

9.1. Single data center loss

While SPC can continue to operate at reduced capacity after loss of a data center, long-term recovery requires provisioning new servers in an unimpacted data center. The time varies per server type but for stateless servers recovery is expected in less than one hour. For stateful servers, like databases, the time to complete recovery and redundancy is longer at around 72 hours. *The loss of a single data center is therefore not expected to cause significant uservisible downtime.*

9.2. Multiple data center loss

The loss of all data centers of an AWS Region causes all SPC stateless and stateful servers to be inaccessible and SPC to be inaccessible. To protect against this business-critical data loss, data is backed up to a secondary Region (set of data centers) on a nightly basis. Recovery requires provisioning new servers within an unimpacted Region and reloading of this backup data. The time to recovery for the loss of a Region is approximately two days. The loss of all data centers of a Region is therefore likely to cause significant downtime. *This has never happened in the nine-year history of AWS and therefore the probability of loss is sufficiently small that the risk is acceptable.*

9.3. Complete AWS loss

Should AWS as a whole become unavailable, SPC would be offline and unable to service users. Recovery from the loss of AWS requires a new hosting relationship with a similar service like Microsoft's Azure or Google's Cloud Platform. We do not have any backups that are not hosted within AWS. We would not be able to recover from a total, sudden, and permanent AWS outage. *The likelihood of this eventuality is so small however that this is an acceptable risk.*

10. Summary

There are limited situations where information pertaining to US clients would be permanently inaccessible. The major impact to business in the US (aside from those matters addressed in the DRP) is in the event that a physical calamity, such as a fire or a building collapse destroys our physical office space.

11. Laws, Rules, Regulations and Other Sources

- Amazon Web Services' Disaster Recovery Whitepaper

12. Document Information

Owner	Kevin Riggle
Approver	Peiter "Mudge" Zatko
Legal Entity Applicability	Stripe Payments Company (SPC)
Effective Date	2018-01-31
Annual Review Date	2018-01-31
Team	security@stripe.com

13. Revision History

Version	Date	Status	Approval
1.1	2018-01-31	Approved	Peiter "Mudge" Zatko