INSTRUCTURE

Disaster Recovery Plan & Procedures

Engineering, Security, and Operations

January 2021

In this guide, you'll learn:

How we recover from disasters and quickly restore our services to minimize loss and disruption to both our customers and our operations.

Table of Contents

Disaster Recovery Overview 3

Introduction 3

Disaster Recovery Procedures 5

Key Organizational Resources 6

Communication Strategy 7

Disaster Resilience 8

Backup and Recovery Practices 10

Disaster Recovery Plan Testing 11

Disaster Recovery Overview

INTRODUCTION

No business wants a disaster, whether it's the catastrophic loss of a datacenter or a crazy panda running around the office pulling out cables. But if or when the time comes, having a robust disaster recovery plan in place allows us to restore our services as quickly as possible and minimize loss or disruption to both our customers and our internal operations.

This document is an overview of the disaster recovery plan and procedures Instructure has established to recover from disasters affecting its production operations. We describe how our Software as a Service (SaaS) offering has been architected to recover from disaster scenarios, the steps we will take if a disaster is declared, our policies, communication strategies and customer notification procedures, and several example scenarios and impact assessments.

KEY TERMS AND ASSUMPTIONS

In the Software as a Service (SaaS) space, there are some key terms in relation to Disaster Recovery.

1) In the context of a disaster recovery scenario, two terms are commonly used to describe how a recovery process may be affected: **Recovery Time Objective (RTO)** and **Recovery Point Objective (RPO)**. The RTO represents how long it will take to restore access to data, and the RPO how much data is at risk of being lost. For example, if it takes 8 hours for a service to be recovered, the RTO is 8 hours. If the last 4 hours of data will potentially be lost due to a disaster, the RPO is 4 hours.

2) While '**Disaster Recovery**' and '**High Availability**' are shared concepts in relation to business continuity, they impact disaster recovery planning differently. Disaster Recovery essentially infers there will be some form of downtime involved, measured in hours or days. High Availability, however, is about ensuring ongoing continuity of operations in a disaster recovery scenario, especially through the design of architectural redundancies such as automated failover of components.

Our services are architected to achieve both exceptionally low RPO and RTO in the most common scenarios and High Availability for our customers due to the distributed and resilient nature of our infrastructure. For the vast majority of failure scenarios, the need to failover to another cloud region is obviated and the impacts to our services will be minimal.

The primary assumption of our disaster recovery plan is that it only addresses events that would affect an entire datacenter or our architecture as a whole. Failures of individual components will be recovered through robust architectural redundancies and failover mechanisms.

DISASTER RECOVERY IN A SAAS WORLD

Instructure's educational software (and associated data) is hosted in the cloud by Instructure and delivered over the internet through the world's most trusted public cloud provider, Amazon Web Services (AWS). This Software as a Service (SaaS) delivery model means that our customers don't have to worry about maintaining server hardware or software, patches, service packs, or, in the context of this document, disaster recovery.

Not only do we maintain our own robust disaster recovery plans and procedures, but we also benefit from using AWS, an Infrastructure as a Service (IaaS) world-leader that bakes redundancy into its services by providing numerous regions, availability zones, and data centers that allow us to recovery quickly in the event of an unforeseen disaster.

Given the nature of the SaaS delivery model, Instructure is responsible for providing disaster recovery in relation to our software and associated data. Naturally, best practice also dictates that our customers develop and maintain their own disaster recovery plans and procedures.

DEFINITION OF A DISASTER

A disaster is defined as any disruptive event that has potentially long-term adverse effects on Instructure's services. In general, potential disaster events will be addressed with the highest priority at all levels at Instructure. Such events can be intentional or unintentional, as follows:

- **Natural disasters**: Tornado, earthquake, hurricane, fire, landslide, flood, electrical storm, and tsunami.
- **Supply systems**: Utility failures such as severed gas or water lines, communication line failures, electrical power outages/surges, and energy shortage.
- **Human-made/political**: Terrorism, theft, disgruntled worker, arson, labor strike, sabotage, riots, war, vandalism, virus, and hacker attacks.

DISASTER RECOVERY PROCEDURES

DISASTER MONITORING PHASE

Instructure monitors the performance of our services around-the-clock using external performance monitoring tools and internal, open- and closed-source monitoring tools. These tools are configured to send real-time alerts to our personnel when certain events occur that would warrant investigation into a potential looming disaster scenario.

ACTIVATION PHASE

All potential disasters are escalated immediately to both the Executive Leadership Team and the Senior Director of Production Engineering (or a designated officer) who are responsible for assessing the event and confirming the disaster. Once confirmed, the Incident Officer is authorized to declare a disaster and begin activation of the Disaster Recovery Team (DRT). Because disasters can vary in terms of severity and disruption, and can also happen with or without notice, the DRT will assess and analyze the impact of the disaster and act quickly to mitigate any further damage.

Once a disaster has been officially declared, the Incident Officer is responsible for directing the DRT recovery efforts and ongoing notifications.

EXECUTION PHASE

Recovery operations commence once the disaster has been declared, the disaster recovery plan activated, the relevant staff notified, and the Disaster Recovery Team (DRT) prepped to perform the recovery activities as outlined in *Backup and Recovery Practices, Performing Recovery*.

KEY ORGANIZATIONAL RESOURCES

INCIDENT OFFICER

Jon Fletcher, Senior Director of Production Engineering

DISASTER RECOVERY TEAM

The Disaster Recovery Team (DRT) is made up of key engineers and operations employees across all areas of our business. The responsibilities of the DRT include:

- Establish communication between the individuals necessary to execute recovery
- Determine steps necessary to recover completely from the disaster
- Execute the recovery steps
- Verify that recovery is complete
- Inform the incident officer of completion

COMMUNICATION STRATEGY

NOTIFYING INTERNAL STAKEHOLDERS

The Incident Officer is responsible for making sure the DRT and any other necessary staff are notified of an emergency or disaster and mobilized.

The DRT (and other key operational staff) have a scheduled on-call roster and are contactable 24x7 in an emergency or disaster. We use a paging platform that specializes in SaaS incident response which allows us to page key staff to commence activation at a moment's notice.

NOTIFYING CUSTOMERS

- **Disaster Declaration**: Impacted customers and business partners will be notified immediately if a disaster is declared. The notification will include a description of the event, the effect to the service, and any potential impact to data.
- **Updates throughout Execution Phase**: Impacted customers and business partners will be kept up to date throughout the disaster recovery process via phone, messaging, and/or email. We will also post official status updates on status.instructure.com.
- **Completion of Recovery:** Once recovery is complete and services have resumed, our customer notifications will include general information about the steps taken to recovery, and any data that may have been impacted. If the recovery is partial and the service is still in a degraded state, notifications will include an estimate of how long the degradation will continue.

If the primary contact(s) for disaster recovery (nominated by the customer) is unavailable, we will notify the alternative contact (also nominated by the customer). If, for any reason, we are unable to contact the customer's primary and alternative contacts, we will endeavor to make contact with other representatives of the customer's organization.

DISASTER RESILIENCE

OPERATING INFRASTRUCTURE

Instructure's services are based on a multi-tier cloud-based architecture. Each component is redundant with active monitoring for failure detection and automated failover. The different tiers are:

Load Balancers

All web traffic to our services is served by load balancers in active/passive configurations. The load balancers are responsible for directing traffic to the next tier.

App Servers

App servers process incoming client requests from the load balancers. App servers implement all the business logic, but do not persist any important data. Asynchronous jobs also run on the app servers. The number of app servers varies based on demand but will always be at least two in active/active configurations.

Caching

To improve performance, Instructure's software aggressively caches data in a caching layer. The data stored in the caching layer is strictly a performance cache. Any data loss resulting from the loss of any of these servers would be limited to a small number of page view statistics that may not have been flushed to persistent storage. The number of cache servers is variable, and the cache data will be partitioned among all servers.

Databases

Most structured data—courses, user information, and assignments, for example—is stored in a database. This data is sharded between instances based on account and on demand. Each shard has a primary and a secondary database, located in geographically separate sites. The data from each primary is replicated asynchronously in near real-time to its corresponding secondary. Each primary is also backed up completely every 24 hours, and the backup is stored in a third geographically separate site. The infrastructure also includes an internal database proxy layer for the relational databases that enables the Operations Team to perform maintenance on the relational database servers with minimal downtime.

Third-Party Object Store

Content—such as documents, PDFs, audio, and video—is stored in a third-party scalable object store.

DATA CENTERS

Data centers are built in clusters in various global regions where we operate. All data centers are online and continually serving our customers; no data center is "cold."

In the case of failure, automated processes move customer data traffic away from the affected area. Our core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites. N in this context simply refers to the amount of capacity needed to run a service at full load. N+1 indicates an additional, duplicate layer has been added to support primary service failure and therefore provide failover and redundancy at equivalent capacity.

As the world leader in Infrastructure as a Service (IaaS), Amazon Web Services (AWS) provides us with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region.

Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by region).

In addition to utilizing discrete uninterruptible power supply (UPS) and onsite backup generators, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to the AWS Global Backbone, a carrier-class backbone built to standards of the largest ISPs in the world (known as Tier 1 transit providers).

DATA SOVEREIGNTY

We architect our AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. For location-dependent privacy and compliance with data sovereignty requirements, such as the EU Data Privacy Directive, data is not replicated between regions. However, in the unlikely event of a disaster that affects a customer's entire region, all services and data can be relocated to numerous active regions within the AWS infrastructure that Instructure uses.

BACKUP AND RECOVERY PRACTICES

Customer data is backed-up automatically both in real-time and on a 24-hour schedule to multiple geographic locations in the customer's region, ensuring the security and reliability of data in the event of a disaster or outage of any scale. Database is backed up from one live database to another, with no additional load on our systems. Static files are stored in secure, geographically redundant storage systems. Recovery backups are encrypted using the AES-GCM 256-bit algorithm and stored within a highly secured separate location.

Instructure retains full database backups (also known as "snapshots") for up to 12 months. We retain 4-hour snapshots (data backed up every four hours) for a 42-hour window, daily snapshots for a 60-day window, and monthly snapshots for a 12-month window. Object data such as files, documents, and uploaded media, etc., are recoverable in the event of a deletion or modification for a period of one year. Further detail on our backup and recovery procedures is included below:

CUSTOMER DATA FROM PRODUCTION DATABASES

Performing Backup	Data is replicated asynchronously in near real-time to remote site (monitored, etc.). Nightly backups of every database are stored at a remote site.
Performing Recovery	When secondary database is up to date (common case): • Promote secondary database to primary, following replication docs • Provision new database using provisioning tools • Establish new database as new secondary, following replication docs When secondary is > 24 hours behind (unlikely): • Copy last nightly backup to secondary database • Load secondary database with nightly backup • Provision new database using provisioning tools • Establish new database as new secondary, following replication docs

STATIC ASSETS SUCH AS DOCUMENTS AND OTHER CONTENT FILES

Performing Backup	Files are stored on a scalable, encrypted, geographically redundant storage (Amazon S3)
Performing Recovery	Recovery in case of failures is built into the scalable storage system

WEB APPLICATIONS

Performing Backup	Web application source code is stored in versioned source control and backed up to multiple locations There is no state stored on the application servers that would need to be backed up
Performing Recovery	Not applicable

DISASTER RECOVERY PLAN TESTING

A Disaster Recovery Plan is only useful insofar as it is tested regularly.

The Incident Officer is responsible for ensuring the Disaster Recovery Plan is tested in its entirety at least annually and in part whenever major components of our architecture are changed.

We frequently test our ability to restore from backup as part of our regular release cycle, as non-production sites are populated from production backups. For example, Canvas beta instance(s) are restored each week from production backup data, thus testing our ability to recover from data loss each and every week (verifiable in a client's own instance).

SAMPLE DISASTER SCENARIOS

We have outlined below several possible disaster scenarios, the services affected, recovery strategies, and the Recovery Point Objective (RPO) / Recovery Time Objective (RTO), services affected, and recovery overview. Note that these are intended only to convey magnitude of impact and recovery efforts required under different situations. Likelihood is an estimated chance of the scenario occurring but does not guarantee occurrence. Last Incident refers to the last time we encountered this disaster recovery scenario in a live environment.

Complete Loss of a Primary Database

Services Affected	Most accounts hosted on the affected database
Recovery Overview	When the secondary database is up to date (common case): The secondary database is promoted to be the new primary database according to the steps described above When the secondary database is inconsistent: The secondary database is populated with the latest nightly snapshot and brought online as the new primary database.
RPO	5 minutes (consistent secondary, common case), 24 hours (inconsistent secondary)
RTO	1 hour (consistent secondary, common case), 6 hours (inconsistent secondary)
Likelihood	Unlikely (Once every 5+ years)
Last Incident	Never

Simultaneous Complete Loss of Primary and Secondary Databases

Services Affected	Most accounts hosted on the affected database.
Recovery Overview	New primary and secondary databases are brought online in separate locations The primary database is populated with data from the remote backup App servers are pointed to the new primary database Replication re-established with the new secondary database
RPO	24 hours
RTO	6 hours
Likelihood	Rare (Once every 20 years; the primary and secondary databases are hosted in geographically separate locations, which makes simultaneous failure unlikely)
Last Incident	Never

Database Destruction by Security Breach

Services Affected	Most accounts hosted on the affected database.
Recovery Overview	The primary database is restored from the most recent complete backup Replication is re-established with the secondary database
RPO	24 hours
RTO	6 hours
Likelihood	Highly Unlikely (Once every 10+ years)
Last Incident	Never

Complete Loss of Primary Hosting Facility

Services Affected	Platform for most accounts
Recovery Overview	New load balancers and app servers are brought up in the secondary site with the secondary database The old secondary database is promoted to primary database. A new secondary database is brought up at a third site and replication re-established DNS is pointed to the new load balancers at the recovery site and services are restored
RPO	4 hours
RTO	Commercially Reasonable
Likelihood	Extremely Unlikely (Once every 100+ years)
Last Incident	Never