



CANVAS DISASTER RECOVERY PLAN AND PROCEDURES

Instructure Security, Engineering, and Operations

INSTRUCTURE, INC.

6330 South 3000 East
Salt Lake City, Utah 84121

Table of Contents

Disaster Recovery Plan and Procedures.....	1
Policy and Practices	1
Definition of Disaster	1
Declaration of Disaster	1
Key Organizational Resources	2
Disaster Recovery Team	2
Notification	2
Notifying Staff.....	2
Notifying Clients and Business Partners	2
Testing	3
Disaster Recovery Solution	3
Current Operating Infrastructure.....	3
Objectives	4
Backup and Recovery Practices	4
Sample Disaster Scenarios.....	5
Complete Loss of a Master Database.....	5
Simultaneous Complete Loss of Master and Slave Databases	6
Database Destruction by Hacker	6
Complete Loss of Primary Hosting Facility	7

DISASTER RECOVERY PLAN AND PROCEDURES

This document describes the plan and procedures that Instructure has established to recover from disasters affecting its production operations. We describe how the Canvas LMS Software as a Service (SaaS) offering has been architected to recover from disaster scenarios, the steps to be taken when disasters are declared, the policies regarding notification of partners during disasters, and several example scenarios and how they affect the service. Our disaster recovery procedures address events which would affect an entire facility. Failures of individual components are recovered through architectural redundancies and fail-over mechanisms.

POLICY AND PRACTICES

DEFINITION OF DISASTER

A disaster is defined as any disruptive event that has potentially long-term adverse effects on the Instructure service. 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, vandalism, virus, and hacker attacks.

DECLARATION OF DISASTER

All potential disasters will be escalated immediately to a designated officer who is authorized to declare a disaster. The incident officer will be responsible for assessing the event and confirming the disaster. Once the disaster is declared, the incident officer will be responsible for directing recovery efforts and notifications.

KEY ORGANIZATIONAL RESOURCES

DISASTER RECOVERY TEAM

The Disaster Recovery Team (DRT) is made up of key engineers and operations employees. 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

NOTIFICATION

There are several parties that must be notified at various stages during disaster events.

NOTIFYING STAFF

The incident officer is responsible for making sure the DRT and any other necessary staff are notified of a disaster event and mobilized. Notification of staff will generally happen via cell phone.

NOTIFYING CLIENTS AND BUSINESS PARTNERS

Clients and business partners will be notified at various stages of disaster recovery using email and our official status page. If these methods are unavailable, notification will happen via alternative means (cell phone, etc.) as provided by each client institution. The stages of notification are:

- **Disaster Declaration:** When a disaster is declared the notification will include a description of the event, the effect to the service, and any potential impact to data.
- **Completion of Recovery:** Once recovery is complete and the service is available, the notification will include general information about 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, this notification will include an estimate of how long the degradation will continue.

TESTING

A Disaster Recovery Plan is only useful insofar as it is tested regularly. The incident officer is responsible for ensuring that the plan is tested in its entirety at least annually and in part whenever major components are changed.

DISASTER RECOVERY SOLUTION

CURRENT OPERATING INFRASTRUCTURE

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

LOAD BALANCERS

All web traffic to .instructure.com is served by two load balancers in an active/passive configuration. The load balancers are responsible for directing traffic to the next tier.

APP SERVERS

App servers process incoming requests from clients 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 website performance, Canvas aggressively caches data in a caching layer. The data stored here 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 Master and a Slave database, located in geographically separate sites. The data from each Master is replicated asynchronously in near real-time to its corresponding Slave. Each Master 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.

OBJECTIVES

In the context of a disaster recovery scenario, there are two terms which are commonly used to describe how the data may be affected: Recovery Time Objective (RTO) and Recovery Point Objective (RPO). The RTO is how long it will take to make access to the data available again, and the RPO is how much of the most recent data will be preserved. For example, if it takes 12 hours for a service to recover, but on a failure up to 24 hours of data may be lost, the RTO is 12 hours and the RPO is 24 hours.

The Canvas platform has been architected to achieve an exceptionally low RPO and RTO in the common case due to the distributed and resilient nature of its infrastructure. For the vast majority of failure scenarios, the need to “failover” to another cloud region is obviated. In the event of a catastrophe, which would necessitate the need to move hosting regions, it would in all likelihood require multiple days for Instructure to restore service to an acceptable level.

BACKUP AND RECOVERY PRACTICES

Instructor, student, course, assignment data from database

Backup	<p>Data is replicated asynchronously in near real-time to remote site (monitored, etc.).</p> <p>Nightly backups of every database are stored at a remote site.</p>
Recovery	<p>When slave is caught up:</p> <ul style="list-style-type: none"> • Promote slave database to master, following replication docs • Provision new database using provisioning tools • Establish new database as new slave, following replication docs • When slave is > 24 hours behind • Copy last nightly backup to slave database • Load slave with nightly backup • Provision new database using provisioning tools • Establish new database as new slave, following replication docs

Static assets from courses and assignments such as documents and other content files

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

Web applications

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

SAMPLE DISASTER SCENARIOS

Following are several different possible disaster scenarios and their RPO/RTO, services affected, and recovery overview. Note that these are intended only to convey magnitude of impact and recovery efforts required under different situations.

COMPLETE LOSS OF A MASTER DATABASE

Services Affected	Most accounts hosted on the affected database
Recovery Overview	<p>When the slave database is up-to-date (common case): The slave is promoted to be the new master according to the steps described above</p> <p>When the slave database is inconsistent: The slave is populated with the latest nightly snapshot and brought online as the new master</p>
RPO	5 minutes (consistent slave, common case), 24 hours (inconsistent slave)

RTO	1 hour (consistent slave, common case), 6 hours (inconsistent slave)
Likelihood	Once a year

SIMULTANEOUS COMPLETE LOSS OF MASTER AND SLAVE DATABASES

Services Affected	Most accounts hosted on the affected database.
Recovery Overview	<p>New master and slave databases are brought online in separate locations</p> <p>The master database is populated with data from the offsite backup</p> <p>App servers pointed to new master database</p> <p>Replication re-established with the new slave database</p>
RPO	24 hours
RTO	6 hours
Likelihood	Once every 20 years (the master and slave databases are hosted in geographically separate locations, which makes simultaneous failure very unlikely)

DATABASE DESTRUCTION BY HACKER

Services Affected	Most accounts hosted on the affected database.
Recovery Overview	<p>The master database is restored from the most recent complete backup</p> <p>Replication is re-established with the slave database</p>
RPO	24 hours
RTO	6 hours
Likelihood	Once every 10 years

COMPLETE LOSS OF PRIMARY HOSTING FACILITY

Services Affected	LMS for most accounts
Recovery Overview	<p>New load balancers and app servers are brought up in the secondary site with the slave database</p> <p>The old slave database is promoted to master database.</p> <p>A new database slave is brought up at a third site and replication re-established</p> <p>DNS is pointed to the new load balancers at the recovery site and services are restored</p>
RPO	4 hours
RTO	Commercially Reasonable
Likelihood	Extremely Unlikely