CANVAS DISASTER RECOVERY PLAN AND PROCEDURES

Instructure Security, Engineering, and Operations

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:

- x   r   r : Tornado, earthquake, hurricane, fire, landslide, flood, electrical storm, and tsunami.

- x   m : Utility failures such as severed gas or water lines, communication line failures, electrical power outages/surges, and energy shortage.

- x   m -m   /   : 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:

- x Establish communication between the individuals necessary to execute recovery
- x Determine steps necessary to recover completely from the disaster
- x Execute the recovery steps
- x Verify that recovery is complete
- x 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

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

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.

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

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

Web applications

B Web application source code is stored in versioned source control and backed up to multiple locations

There is no state stored on the application server

COMPLETE LOSS OF PRIMARY HOSTING FACILITY

r **Aff**	LMS for most accounts
r **r**	New load balancers and app servers are brought up in the secondary site with the slave database The old slave database is promoted to master database. A new database slave 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
	4 hours
	Commercially Reasonable
⊤	Extremely Unlikely