

# Disaster Recovery

How SnapLogic will recover the platform and data in the event of a disaster; and a forward look at SnapLogic’s disaster recovery strategy.

## Introduction

The purpose of this whitepaper is to present SnapLogic’s method and means to recover from a disruptive event to ensure workloads run as expected and remain durable.


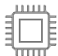



## Background

Businesses create and manage large volumes of data. As companies extend their use of artificial intelligence into daily operations, data becomes mission critical. The impact of data loss or corruption can be significant. Preparing plans to continue business in response to an event is not only good management, but essential in today’s business climate.

At SnapLogic, we are revising our view of resiliency and taking a broader approach. Beyond the standard disaster recovery and cybersecurity requirements, real resiliency should include being able to handle a full range of responses a company needs to keep its business going, whatever happens.

## SnapLogic’s Disaster Recovery Model

[ < 48 Hour Recovery Time and < 2 Hour Data Loss ]

Presentation Layer		Where users interact with the application
Application Layer		Pipeline creation and management
Data Execution Layer		Pipeline execution
Database Layer		Pipeline meta-data storage Pipeline log data storage
Snaps		Library of APIs (connectors) Used to execute pipelines

SnapLogic’s infrastructure is deployed on Amazon Web Services. When users access the application, their request is filtered through a load balancer which is a cloud networking device to help distribute large volumes of internet traffic. Requests are passed from the load balancer to a cluster of web servers which present the application(s) to the user.

The application servers, sometimes referred to as the Control Plane, facilitate and manage pipeline creation, updates, and monitoring. The data execution servers, sometimes referred to as Cloudplexes, Groundplexes, or the Data Plane, are the heart of pipeline execution.

---

Database servers store pipeline metadata and log data each time a pipeline is run.

The last component are Snaps, sometimes referred to as connectors or APIs, they are small computational units that process data.

These assets are monitored for performance 24x7x365. If performance drops below a certain threshold, an automated process notifies the appropriate team(s) to investigate and if necessary, recover or replace an infrastructure asset.

Assets are managed using a combination of automated and manual processes. The presentation, data execution, and database layers would be restored manually. The application and Snaps layers leverage fully automated processes.

In the event of a disaster, SnapLogic's end to end recovery time for the entire infrastructure is within 48 hours of being hard down (inoperable). The amount of data loss would be 2 hours or less.

The data execution servers, sometimes referred to as Cloudplexes, Groundplexes, or the Data Plane, have a shared responsibility for resiliency. Groundplex resiliency is solely the responsibility of the customer. Cloudplex resiliency is shared. SnapLogic is responsible for the resiliency of the infrastructure, once recovered, the customer is responsible for recovering / restarting their pipeline(s).

This outlines SnapLogic's disaster recovery model. If you have any questions about this white-paper or would like to discuss resiliency in more detail, please contact your SnapLogic representative who can schedule a deeper discussion.