

COMPLIANCE

# Service Level Agreement

About uptime and help. Explains availability targets, incident handling, response expectations, and what remedy applies if targets are missed.

Version	1.1
ID	b20321d2d0e5
Prepared by	Safespring
Approver	Safespring
Date	2026-04-15
Classification	Public

During the Term of the agreement under which Safespring has agreed to provide its services to Customer (the “Agreement”), each Covered Service will provide a Monthly Uptime Percentage to Customer as set forth in the tables below (the “Service Level Objective” or “SLO”).

If Safespring does not meet the SLO, and if Customer meets its obligations under this SLA, Customer will be eligible to receive the Financial Credits described below. This SLA states Customer’s sole and exclusive remedy for any failure by Safespring to meet the SLO.

## 1. Service Level Objectives

Monthly Uptime Percentage targets are set per Covered Service. Each service is measured independently on a calendar month basis.

Covered Service	Monthly Uptime %
<b>Compute (IaaS), Network, Block Storage</b>	<b>99.99%</b>
<b>Kubernetes (control plane)</b>	<b>99.95%</b>
<b>Object Storage (S3)</b>	<b>99.95%</b>
<b>Backup Service</b>	<b>99.9%</b>

### Kubernetes SLO Scope

Kubernetes SLO applies to the control plane only. Worker node availability is governed by the underlying Compute SLO.

## 2. Definitions

The following definitions apply to the SLA:

### Covered Service

- **COMPUTE** Safespring Compute service including virtual machines, local storage, and block storage volumes.
- **NETWORK** The network connectivity between customer instances within a Safespring site and to external endpoints.

- **BLOCK STORAGE** Persistent block storage volumes attached to Compute instances.
- **KUBERNETES** The Kubernetes API and control plane provided by Safespring's managed Kubernetes offering.
- **OBJECT STORAGE** Safespring's S3-compatible object storage service.
- **BACKUP** Safespring's cloud backup service.

## Downtime

Downtime is defined per Covered Service as follows:

- **COMPUTE** A virtual machine becomes unreachable, crashes, or loses access to its attached block storage, affecting instances across two or more physical hosts. A single host failure is excluded (see Section 5).
- **NETWORK** Complete loss of network connectivity at the site level affecting all customers.
- **BLOCK STORAGE** Inability to perform read or write I/O operations on provisioned volumes.
- **KUBERNETES** The Kubernetes API is unreachable or unable to process requests for all clusters. Failure of individual worker nodes or pods is not considered Downtime and is covered under the Compute SLO.
- **OBJECT STORAGE** The S3 API is unreachable or unable to process read, write, or list requests.
- **BACKUP** The backup service is unreachable and clients are unable to start or complete backup or restore operations.

## Downtime Period

"Downtime Period" means a period of five or more consecutive minutes of Downtime. Intermittent Downtime for a period of less than five minutes will not be counted towards any Downtime Periods.

## Monthly Uptime Percentage

"Monthly Uptime Percentage" means total number of minutes in a calendar month, minus the number of minutes of Downtime suffered from all Downtime Periods in that month, divided by the total number of minutes in the month.

## Scheduled Downtime

"Scheduled Downtime" means Downtime resulting from Safespring performing maintenance during a pre-communicated Maintenance Window. Scheduled Downtime is excluded from the Monthly Uptime Percentage calculation.

## Maintenance Window

A period of time during which Safespring performs planned maintenance on a Covered Service. Maintenance Windows shall be announced at least 5 working days in advance. One maintenance window per service per month is standard. In exceptional cases involving external security threats, additional maintenance windows may be announced with as much notice as possible.

## 3. Financial Credits

Financial Credits are determined on a calendar month basis per Covered Service.

- **MUP** Monthly Uptime Percentage
- **PMB** Percentage of monthly bill for the respective Covered Service affected which did not meet SLO that will be credited to future monthly bills of Customer.

### Compute, Network, and Block Storage

MUP	PMB
99.00% to < 99.99%	10%
95.00% to < 99.00%	25%
< 95.00%	50%

### Kubernetes and Object Storage

MUP	PMB
99.00% to < 99.95%	10%
95.00% to < 99.00%	25%
< 95.00%	50%

### Backup

MUP	PMB
99.00% to < 99.90%	10%

MUP	PMB
95.00% to < 99.00%	25%
< 95.00%	50%

## Customer Must Request Financial Credit

In order to receive any of the Financial Credits described above, Customer must notify Safespring within thirty days from the time Customer becomes eligible to receive a Financial Credit. Customer must also provide log files or monitoring data showing service unavailability and the dates and times those errors occurred.

## Maximum Financial Credit

The aggregate maximum number of Financial Credits to be issued by Safespring to Customer for any and all Downtime Periods that occur in a single billing month will not exceed 50% of the amount due by Customer for the affected Covered Service for the applicable month. Financial Credits will be made in the form of a monetary credit applied to future use of the Service and will be applied within 60 days after the Financial Credit was requested.

## 4. SLA Exclusions

This SLA does not apply to any:

1. Scheduled Downtime or pre-communicated Maintenance Windows;
2. Hardware failure affecting individual compute hosts;
3. Features designated as beta;
4. Errors caused by factors outside of Safespring's reasonable control;
5. Errors that resulted from Customer's software, hardware, or third-party software or hardware;
6. Errors that resulted from abuses or other behaviors that violate the Agreement or Acceptable Use Policy;
7. Errors caused by quotas or resource limits listed in the self-service portal.

## 5. Measurement and Monitoring

Availability is measured at the connection point (delivery point) using Safespring's infrastructure monitoring systems at the datacenter.

Customer is responsible for monitoring their own instances and ensuring services are running after an interruption.

Customers are encouraged to subscribe to the status page at <https://status.safespring.com> for real-time incident updates.