



The remedies set forth within this SLA are your sole and exclusive remedies for any failure of the service.

NETWORK

INTERNAL NETWORK

The internal network includes cables, switches, routers, and firewalls within our network perimeter, under our exclusive control, or the control of our immediate upstream network partners.

DOWNTIME

Internal network downtime exists when two Enscale servers (nodes) provided by Layershift at the same hosting location (Enscale region) cannot communicate with each other for a consecutive period in excess of fifteen (15) minutes due to a fault within the internal network. It does not cover any other reason for communication difficulties such as incorrect configuration on either server or firewall restrictions.

GUARANTEES

100% uptime for internal network connectivity.

REMEDIES

Service credits as a percentage of the total fees incurred **by the Enscale servers with internal network downtime** in the calendar month prior to the relevant outage:

Uptime	Credits from monthly fees
<100%	10%
<99.5%	25%
<95%	50%
<90%	100%



INTERNET / EXTERNAL NETWORK

The Internet (External Network) includes third party networks between the internal network and end user Internet connections. It excludes end user Internet connections and their respective provider networks, as these are the exclusive remit of the respective end user connectivity provider (ISP) and cannot be influenced or circumvented by any hosting provider.

DOWNTIME

External network downtime exists when an Enscale server (node) provided by Layershift cannot send and receive data to and from the Internet for a consecutive period in excess of fifteen (15) minutes due to a fault within the external network. For practical purposes, an Enscale server (node) is **not** considered to be experiencing external network downtime providing that it can send and receive data to and from at least any two (2) major transit networks as defined by Layershift.

GUARANTEES

99.99% uptime in a calendar month

REMEDIES

Service credits as a percentage of the total fees incurred **by the Enscale server(s) with Internet downtime** in the calendar month prior to the relevant outage:

Uptime	Credits from monthly fees
Over 99.99%	-
<99.99%	10%
<99.5%	25%
<95%	50%
<90%	100%



CRITICAL INFRASTRUCTURE

Critical infrastructure includes power and HVAC, including UPS equipment and cabling, but excludes server hardware, software, and power supply units (PSUs).

GUARANTEES

100% uptime for all critical infrastructure.

DOWNTIME

Critical infrastructure downtime exists when your service is offline due to power or heat problems, and is measured from the first of (a) when infrastructure monitors alert our engineers, or (b) a technical support ticket is created to report a service problem with a critical infrastructure failure as the root cause.

REMEDIES

Service credits as a percentage of the total fees incurred **by the Enscale server(s) with critical infrastructure downtime** in the calendar month prior to the relevant outage:

Uptime	Credits from monthly fees
<100%	10%
<99.5%	25%
<95%	50%
<90%	100%



PLATFORM UPTIME

Platform includes the physical hardware and any composite software layers such as virtualization and automation technology responsible for providing the Enscale service.

GUARANTEES

Downtime caused by platform defects will be corrected within a *maximum* time period from us identifying the specific fault for services:

- ten (10) minutes.

DOWNTIME

Platform downtime exists when one or more Enscale servers (nodes) within your environment is not running and cannot be started due to hardware or underlying software layer problems, and is measured from the first of (a) when infrastructure monitors alert our engineers, or (b) a technical support ticket is created to report a service problem with a critical platform failure as the root cause.

REMEDIES

Service credits as a percentage of the total fees incurred **by the Enscale server(s) with platform downtime** in the calendar month prior to the relevant outage:

Time to repair	Credits from monthly fees
Within 10 minutes	-
<1 hour	10%
<4 hours	15%
<12 hours	25%
Over 12 hours	50%



SUPPORT AND PLATFORM MANAGEMENT

Support and Platform Management refers to the technical assistance and administration conducted by us to help you with your Enscale environments. It forms a critical part of our service to you and we know it makes the difference between your services working exactly as you need, or malfunctioning and causing hours of headaches for you and your clients.

SCOPE

Support and Platform Management included FREE with Enscale is comprised of the following in relation to your service provided by us:

- Troubleshooting and resolution of issues affecting the operation of pre-installed software stacks (e.g. Apache, Nginx, Tomcat, MySQL)
- Upgrade, patching, configuration, and optimization of pre-installed software
- Expert advice for tuning/configuring pre-installed software stacks for optimum performance and stability (e.g. tuning MySQL configuration)
- Troubleshooting and resolution of all platform and related issues

For the avoidance of doubt, this leaves the following as customer responsibilities:

- Configuration tasks within the Enscale dashboard interface (e.g. deploying a new server)
- Website and application software development / installation / debugging
- Operation of your own personal computer (we will try to assist with this where reasonably possible, but we cannot support problems with your personal computer).

GUARANTEES

Support and Platform Management is provided on a 24x7x365 basis.

Although we aim to provide all customers with rapid and efficient technical assistance at all times, where response times are measured in minutes, we also guarantee that the **maximum** response time to any new support request will be as follows for services with:

- High severity (e.g. your environment is down or unstable): one (1) hour maximum, **target response time - 5 minutes;**
- Normal severity (e.g. assistance tuning MySQL): two (2) hours maximum, **target response time - 15 minutes;**
- Low severity (e.g. feature request for Enscale): six (6) hours maximum, **target response time - 1 hour.**

Response times are measured from the time a ticket is created in our helpdesk, until the time an engineer logs an initial response within the ticket. This response will usually indicate findings of



the engineer's preliminary investigation into your issue.

We take responsibility for resolving a diverse range of issues of significantly varying complexity, so a reasonable time for resolving one issue is entirely unreasonable for resolving another. Therefore we are unable to guarantee support request resolution time (though other aspects of this SLA address fault repair guarantees). All of our engineers are technical experts in their field and we aim to resolve all issues as expeditiously as possible.

If we determine that multiple tickets are opened about the same or closely related issues we may merge the related tickets and reply to you only in one ticket. Any such superfluous tickets will be closed and excluded from this guarantee.

During the following holiday periods (exclusively) we operate with a reduced staffing level and may provide a reduced support service during which the maximum response time guarantee does not apply to low-priority tickets. This allows us to ensure we can maintain a high quality support service for all other issues. We will assess the severity of a ticket based on our reasonable understanding of the impact on a customer's business. All low-priority tickets will be responded to within two (2) hours of the relevant window ending, though in most cases will already have been handled during these windows.

- Between 02:00, December 25th to 08:00, December 26th (UTC)
- Between 22:00, December 31st to 08:00, January 1st (UTC)

TICKET ESCALATIONS

We offer a ticket escalation method for the majority of our services for use when an issue is having a significant impact upon your business. The ticket escalation method, if applicable to your service, is described in the automatic reply you will receive when opening a ticket. If your ticket is business critical we encourage you to escalate your ticket to ensure we can provide you with an even faster response than usual. In the vast majority of cases this is not necessary, but we have introduced this to provide added reassurance to our customers.



REMEDIES

Service credits as a percentage of the total fees incurred **by the relevant Enscale environment** in the calendar month prior to the relevant outage:

Response time outside guarantee by	Credits from monthly fees
<1 hour	2%
<4 hours	5%
Over 4 hours	10%



BACKUPS

We take regular filesystem backups of your Enscale environment as part of our service. Very few providers offer any guarantees at all when backing up your servers, even though most charge extra for it! Unlike many of our competitors we are able to restore individual files and directories without overwriting your entire environment - only targeted data is overwritten, and only upon your request.

We recommend that you follow software vendor guidelines for taking consistent copies of any frequently changing files, such as databases, to ensure our backup systems are able to reliably backup and restore consistent data.

FREQUENCY

A backup snapshot of each server (node) within your Enscale environment is taken automatically every 6 hours

RETENTION POLICY

We retain a rolling backup history, so that we can restore any one or more files from any of your recent backups upon request. The number of backup recovery snapshots held is 56 snapshots (14 days).

GUARANTEES

Backup integrity and availability are very important to us, but in rare cases operational reasons such as software or hardware malfunction, or a backup or restore process taking longer than normal to complete, may prevent us from maintaining the configured backup frequency or overall number of backup snapshots.

However, we guarantee to always have a minimum of 80% (45 backup snapshots) of the configured backup snapshots available to restore to your environment.

REMEDIES

Service credits: If less than the number of guaranteed backups are available your service **for the affected Enscale server(s)** within your environment will be provided free of charge until we reach the guaranteed level of backup restore points.



SERVICE CREDIT CONDITIONS

- i. The total value of service credits issued for any service in any one (1) month cannot exceed 100% of the total fees incurred through the use of the Enscale platform in the previous calendar month.
- ii. Requests for service credits must be submitted in writing to our billing department within 5 business days of the related incident.
- iii. You must open a support case during the failure in question, including detailed information about the problem experienced and steps to reproduce the error symptoms. In most circumstances this will be considered as the start time for the downtime.
- iv. Service credits will not be aggregated if multiple failures overlap; you will receive credit as applicable only to the service component which we identify as the root cause.
- v. Service credits are non-transferable and may only be applied to future service delivery on the account they're issued to, and cannot be exchanged for cash or other forms of payment. Any outstanding or previously accrued service credits will be forfeited upon cancellation of the Customer's account.
- vi. All downtime measurements will be exclusively at our discretion, and our sole decision in respect of service credit entitlement and amount will be final and binding.
- vii. Free trial accounts, accounts that are past-due, have been in payment arrears two (2) or more times in the previous twelve (12) months, or have outstanding AUP violations, waive any right to claim under this SLA.
- viii. Excludes periods of scheduled and emergency maintenance, issues arising due to acts or omissions of the Customer, or the occurrence of a Force Majeure event.
- ix. Where multiple servers are load balanced in a cluster configuration, this SLA shall only apply in the event the entire cluster experiences downtime as defined within each SLA section, because a load balanced cluster is intentionally designed to be fault tolerant and allow one or more servers to be offline without service impact.