

Heata Compute Service Level Agreement (SLA)

Last Updated: March 2026

Key components

There are two components to the *Heata Network*; cloud-based *Access Nodes*¹, and *Heata Worker Nodes* which are distributed, one per home, providing free hot water as a byproduct of their processing.

A cluster of *Heata Worker Nodes* undertakes the processing, whilst the *Access Node* provides access to the cluster of *Heata Worker Nodes* via a Wireguard Virtual Private Network (VPN).

A *service* is made up of an *Access Node* and a cluster of *Worker Nodes*, the number of nodes in the cluster is determined by the *Client* during set-up.

Where necessary the *Heata Network* can also call upon *Cloud Burst Nodes* to manage peaks in demand and to guarantee service.

¹ Also referred to as *Head Nodes* or *Control Planes* in some service descriptions

SERVICE LEVEL GUARANTEES:

Maintenance

Heata will:

Ensure that *Scheduled Maintenance* of the *Heata Network* occurs between 12:00 am and 08:00 am *Local Time* for the affected region(s);

and

Notify you or your *Reseller Partner(s)* (as applicable) of any maintenance window for the affected region(s) at least 1 business day in advance.

Access Nodes are hosted within [GCP](#), please check this link for more detail on Maintenance schedules that apply to *Access Nodes*.

Uptime and Committed Capacity

Access Nodes have a monthly *Uptime* of >99.5%.

Worker Nodes have a *Committed Capacity* of >99.5%, calculated across all *Jobs* in the *Applicable Period*, weighted by *Requested Capacity*. The percentage refers to the amount of time per *Job* that a Client had their full *Requested Capacity* committed to work on the Job.

Exclusions

Scheduled or Emergency Maintenance Windows:

Outages leading to *Downtime* or *Reduced Capacity* do not include times designated, and notified in advance, to be *Scheduled Maintenance*, or *Emergency Maintenance* in accordance with the Terms of this SLA.

Third party platforms

Any *Downtime* or *Reduced Capacity* resulting from the failure of third-party platform providers beyond Heata's reasonable control, provided Heata is within its own *Uptime* commitments.

Force Majeure:

Any *Downtime* or *Reduced Capacity* resulting from a Force Majeure event as defined in this Agreement

Calculating Uptime and Committed Capacity

Access Nodes - Uptime Percentage Calculation
The <i>Uptime Percentage</i> for each <i>Access Node</i> is calculated by the percentage of Minutes in the <i>Applicable Period</i> in which an <i>Access Node</i> had <i>Downtime</i> .
$\frac{\text{Total Minutes in an Applicable Period} - \text{Minutes of Downtime in an Applicable Period}}{\text{Total Minutes in an Applicable Period}} \times 100$
<p>Eg. in this scenario the <i>Uptime Percentage</i> would be 99.5%</p> <p>Example time period: 30 days. Total Minutes in Applicable Period: 43,200 Minutes of Downtime in Applicable Period: 216</p> <p>43,200 minutes – 216 minutes = 42,984 minutes (42,984 / 43,200) x 100 = 99.5%</p>

Worker Node Cluster - Committed Capacity Percentage Calculation
This is calculated by comparing the <i>Committed Capacity</i> with the <i>Requested Capacity</i> . <i>Reduced Capacity</i> occurs when the number of functional cores or nodes available to the client drops below the client's <i>Requested Capacity</i> during a <i>Job</i> . <i>Committed Capacity %</i> is calculated across all <i>Jobs</i> in the <i>Applicable Period</i> , using <i>Minutes Under Capacity</i> accumulated across those <i>Jobs</i> .
$\frac{(\text{Job Length in Minutes} \times \text{Requested Capacity}) - \text{Minutes Under Capacity}}{\text{Job Length in Minutes} \times \text{Requested Capacity}} \times 100$
<p>Eg. in this scenario the <i>Committed Capacity</i> would be 96%.</p> <p>Job time: 100 minutes Requested Capacity: 250 cores Total Core minutes if at full capacity: 25,000</p> <p>Committed capacity: 200 cores operational for 100% of the time. 50 cores were unavailable for 20 minutes</p> <p>25,000 core minutes - 1000 core minutes = 24,000 core minutes (24,000 / 25,000) x 100 = 96%</p>

SERVICE CREDIT FOR NOT MEETING SLA

Client Access Node(s) Monthly Uptime %	Service Credit
< 99.5%	10%
< 99.0%	25%
< 95.0%	50%

Committed Capacity as a % of Requested Capacity	Service Credit
< 99.5%	10%
< 95.0%	25%
< 90.0%	50%

Requesting Service Credit

To receive any Service Credit, you must notify us within 60 days from the date of an outage that qualifies for credit. Notification must be done via email, which must include reasonable evidence of loss of connectivity, including timestamps, error messages, and request IDs where available. If these are not provided, you will forfeit your rights to receiving Service Credit.

Maximum Service Credit

The maximum Service Credit provided by us for any outages will not exceed 50% of the amount due from you for that billing month. Credit will be provided in the form of monetary credit which can be used for any of our services and will be issued within 30 days of the successful Service Credit request. The credit will expire 12 months after the date of issue. Service Credits are your sole and exclusive remedy for any failure by Heata to meet the Service Levels.

SERVICE DELIVERY

Cloud burst facility

To guarantee delivery of workloads we may use cloud instances (*Cloud Burst Nodes*) in addition to our distributed *Worker Nodes*. This will be recorded and we can provide you with this information.

UNDERUTILISATION OF ALLOCATED CAPACITY

In situations where a *Client* has *Allocated Capacity* but is not actively using it, Heata reserves the right to utilise this capacity to maintain efficient utilisation of the network and ensure reasonable hot water delivery to the household(s). This will not interfere with the Client's use of the *Allocated Capacity*, or it may trigger a *Minutes Under Capacity* claim. Any such utilisation will be fully sandboxed and will not provide any access to Client data, environments, or networks, and if the Client submits workloads, these will take priority.

SUPPORT TIERS

We offer three levels of support to ensure you have the right coverage for your team.

Feature	Standard	Priority	Enterprise
Monthly Cost	Included	Request a Quote	Request a Quote
Primary Channels	Email, Ticketing	Email, Ticketing, Slack	Email, Ticketing, Slack, Phone
Technical Review	Self service	Annual Review	Quarterly Review
SLA Credit	No	Yes	Yes

*Response times**

	Standard	Priority	Enterprise
P1 Critical	8 hours	2 hours	30 mins
P2 Significant Issue	12 hours	4 hours	2 hours
P3 General technical issue	24 hours	12 hours	8 hours
P4 Feature request	48 hours	24 hours	12 hours

**Response times apply during Business Days between 9am and 5pm for Standard and Priority tiers, and 24/7/365 for Enterprise tiers*

DEFINITIONS

General

Billing Month:

The monthly period for which Heata issues an invoice to the Client.

Business day:

Monday through Friday, excluding UK Bank Holidays.

Client:

A customer buying compute from Heata.

Cloud Burst:

This is when Heata uses compute capacity based in the cloud to guarantee delivery of workloads to meet the SLA or to deal with overcapacity on the network.

Cluster:

A group of *Heata Worker Nodes* assigned to an individual plan.

Compute Service:

Compute Service means the provision of Access Nodes and Worker Nodes required to submit and execute Jobs, including access via VPN

Distribution Framework:

The specific software architecture, APIs, or job-scheduling tools (e.g., Ray, Slurm, Kubernetes, etc) used by the Client to submit and manage workloads.

Force Majeure:

Any act, event, or condition beyond the reasonable control of Heata, including but not limited to: natural disasters, acts of war or terrorism, civil unrest, government action, nationwide strikes, or widespread failures of the public power grid or internet infrastructure.

Access Node:

The cloud-based virtual machine that provides access to the *Heata Worker Nodes* assigned to a *Client*. A *Client* could have multiple *Access Nodes* assigned.

Heata Compute:

Compute services provided by the *Heata Network*.

The Heata Network:

The Heata distributed data centre, comprising *Heata Worker Nodes*.

Local Time:

Local Time means the local time in the region where the affected service is deployed.

Material Activity:

Any action essential to the Client's use of the Service, including but not limited to: the submission of new jobs, the retrieval of processed data, or the ability to monitor the status of active workloads.

Reseller Partner(s):

Companies who resell Heata compute capacity or plans to third parties.

Service Credit:

A credit (GBP), calculated as a percentage of the fees paid for the Applicable Period, which Heata may apply to the Client's future invoices as the sole remedy for a failure to meet Uptime or Capacity commitments.

Virtual Private Network (VPN)

A Virtual Private Network is an encrypted connection over the Internet from a device to a network, known as a tunnel. Security checks can be conducted on endpoints to ensure that they meet certain requirements before connecting. Traffic is encrypted as it travels through the tunnel, this ensures that sensitive data is safely transmitted and that unauthorised people cannot eavesdrop on the traffic. It is commonly used by companies to enable offsite employees to access the corporate network.

Worker Node:

A single Heata server that makes up a component of the *Heata Network*, comprised of a number of cores as specified in the order form.

Jobs and tasks

Job:

A connected series (including the singular case) of *Tasks*.

Task:

An indivisible unit of a *Job*. Here is an example Task:

```
docker run -v /your/images:/imgs dpokidov/imagemagick /imgs/sample.png -resize 100x100 /imgs/resized-sample.png
```

Supported Job Submission:

Submission of a well defined *Job* to whichever version of distribution framework is being offered, according to the documentation associated with the distribution framework.

Uptime and Downtime

Allocated Capacity:

The total amount of compute resources (expressed in Worker Nodes or Cores as specified in the Order Form) that Heata commits to making available for the Client's use. This represents the maximum capacity the Client can provision for a *Job*.

Applicable Period:

The specific calendar month in which a Service Level failure occurred and for which Uptime or Committed Capacity is being calculated.

Capacity Minutes:

A unit of measurement calculated by multiplying the number of compute units (e.g., Cores or *Worker Nodes*) by the duration of the *Job* or Outage in minutes.

Committed Capacity:

The actual capacity (measured in Node/Core Minutes; total *Job* time x committed *Worker Node/Core Minutes*) deployed to perform a *Job*.



harnessing heat from
cloud compute

[Heata Compute Service Level Agreement](#)
Bit Warmer Limited (trading as 'Heata').
Registered office: Wood Farm, Portsmouth Road, Godalming, England, GU7 2JR.
Company Number: 11110428. VAT number: 294 2496 68.

Downtime:

Loss of connectivity to *Client Access Node(s)*.

Minutes Not Available:

Minutes Not Available is defined as minutes a *Client Access Node* is *Not Available for Job Submission*. From the time a *Not Available for Job Submission* condition occurs, *Minutes Not Available* will accumulate until:

- (a) a subsequent *Supported Job Submission* succeeds;
- (b) another *Not Available for Job Submission* condition results from another *Supported Job Submission* attempt, or;
- (c) 1 hour has elapsed with no *Supported Job Submission* attempted. *Minutes Not Available* will resume accumulation upon a subsequent *Not Available for Job Submission* condition.

Minutes Under Capacity:

Minutes Under Capacity is defined as minutes per *Job* where the *Committed Capacity* is lower than the *Requested Capacity*. From the time a *Reduced Capacity* condition occurs, *Minutes Under Capacity* will accumulate until the *Requested Capacity* has been replaced and the cluster is back to full capacity. *Minutes Under Capacity* does not include the standard initial spin-up or provisioning time.

Not Available for Job Submission:

Loss of connectivity to *Access Node(s)* such that *Jobs* cannot be submitted.

Requested Capacity:

The amount of compute resources (expressed in Nodes or Cores) specified by the Client for a particular *Job* or *Cluster* submission within the *Distribution Framework*. For some services, the *Allocated Capacity* and *Requested Capacity* will be the same.

Reduced Capacity:

A condition where the total available compute resources in a Client's cluster drops below the *Requested Capacity* specified for an active *Job* or *Cluster*.

Uptime:

The total accumulated *Minutes Available* in a given month for *Access Nodes*. Uptime Percentage is the percentage of *Minutes Available* in the *Applicable Period* for *Access Nodes*.

Outage:

An *Outage* is defined as:

Service Element	Definition of an Outage
Access Node	Loss of connectivity to <i>Access Node</i> for a period of time (' <i>Downtime</i> ')
Worker Node	If a <i>Worker Node</i> becomes unavailable, an outage is the period of time before that Node is automatically replaced in the cluster. (' <i>Reduced Capacity</i> ')

Maintenance

Emergency Maintenance:

Essential maintenance that is required to prevent a material impact on our service, including (but not limited to) the risk of a cyber attack, data integrity, or platform availability.

Scheduled Maintenance:

Planned maintenance which is likely to affect the services such that the *Service* is unavailable for *Clients* to perform any Material Activity.

Maintenance Window:

The period of time where maintenance is being undertaken.