

# HPE GreenLake for Red Hat OpenShift Container Platform

## 1. Overview

HPE GreenLake for Red Hat® OpenShift provides Hewlett Packard Enterprise and Red Hat OpenShift Container Platform (OCP) validated integrated systems wrapped in HPE GreenLake services. This solution offers an as-a-service (aaS) cloud experience for your on-premises container deployments using an HPE validated architecture and Red Hat's OpenShift container environment. In addition, HPE as your cloud experience provider and technical one-stop support specialist can now be more confident about the OpenShift experience.

The HPE GreenLake for Red Hat OpenShift cloud experience provides on-demand capacity, combining the agility and economics of the public cloud with the security and performance of on-premises IT. When you purchase this service, you can accelerate [digital transformation](#) with cloud benefits of fast deployment, scalability, and pay-per-use<sup>1</sup> economics—all within your control and your choice of on-premises or colocation environment. HPE gives you an integrated OCP built on the HPE GreenLake edge-to-cloud platform that incorporates compute, storage, and networking, to run applications and workloads on the OCP.

HPE GreenLake for Red Hat OpenShift delivers a service option with a managed infrastructure aaS. It provides a preconfigured, integrated system with your choice of cloud module instances, services for startup, advanced configuration, and ongoing support for the completed system. In addition, an account team is assigned and acts as the Customer's advocate and operational and/or technical focal point for support.

The service also includes monitoring and alerting of issues, patches, and advisories that could impact the Customer's integrated system environment. The Customer is primarily responsible for proactive monitoring, managing, and operating the integrated system with HPE support.

This data sheet and the service features in Table 1 describe the high-level components of HPE GreenLake for Red Hat OpenShift. The scope of available service options is tailored to your selected requirements and is detailed in a mutually agreed-upon and implemented statement of work (SOW) and is not available for purchase separately. Any defined terms used in this data sheet will have the meaning ascribed to them in this document or the HPE GreenLake terms.

## 2. Service benefits

The service is designed to provide the following benefits:

- A complete solution including scalable infrastructure, software, and services
- Faster time to value with industry leading HPE systems, storage, and network platforms
- Industry-leading Red Hat OCP
- Simplified IT with implementation, support, and monitoring; with operations and management options
- Convenient, pay-per-use billing based on metered usage; includes storage, container, and compute units of measure
- Cloud-like flexibility to meet business needs

<sup>1</sup> May be subject to minimums or reserve capacity may apply

### 3. Service feature highlights

HPE GreenLake for Red Hat OpenShift provides the following:

**Table 1.** Service features

Service features	Managed infrastructure aaS
Relationship management	✓
HPE GreenLake integrated system	✓
Service implementation	✓
IT service management	✓
System operations and administration	✓

✓ = Basic deliverables

### 4. Service feature specifications

#### 4.1 Relationship management

The purpose of relationship management is to govern the relationship between the Customer and HPE for smooth and swift collaboration between both parties. The team members typically increase with higher service experiences.

**Table 2.** Account team members

Team member and their role	Managed infrastructure aaS
<p><b>Managed service manager (MSM)</b>—Is the Customer's primary point of contact within HPE for the service and is responsible for coordinating service delivery. The MSM will meet the Customer monthly and remotely for the operational support meeting to review the delivery of the HPE GreenLake for Red Hat OpenShift services. The MSM will provide advice and guidance on the routine delivery of the Customer's critical IT services and running the service management processes and technology. If potential risk factors are identified through the service delivery, HPE will provide related recommendations for consideration by the scope Customer and implementation through the appropriate change management process.</p> <p>Note: The MSM will be part of the wider assigned account team supporting the overall services being provided to the Customer.</p>	✓
<p><b>Account support manager (ASM)</b>—If the managed infrastructure is selected, then the ASM will be the single point of contact overseeing the delivery of the post-implementation services. The ASM is responsible for coordinating the proactive services as described herein.</p>	✓
<p><b>Technical account manager (TAM)</b>—Will provide remote proactive services such as support planning and review and support activity reporting services.</p>	✓
<p><b>Consumption services specialist (CSS)</b>—Provides initial service setup using the HPE transition and transformation methodology (TTM) and works with the HPE account team and Customer, leading the contract change management process where capacity changes are required.</p>	✓



### 4.2 Governance

Integral to making the service experience successful is a well-structured governance process, bringing together planned and unplanned activities and making sure the service is performing as planned.

**Table 3.** Governance structure

Governance process	Managed infrastructure aaS
<p><b>Operational service plan (OSP)</b>—Working with the Customer, HPE will develop and maintain an operations handbook containing all standard operating procedures (SOPs) specific to maintaining the systems and documenting the agreed ITSM processes and escalation protocols. The document will be reviewed remotely every quarter with the Customer making required changes through the change management process.</p>	
<p><b>Account support planning (ASP)</b>—The ASM documents the expected service delivery activities during the contractual period and would cover the standard included services and any selected options.</p>	✓
<p><b>Operational service meeting (OSM)</b>—The OSM reviews service performance to target service-level objectives (SLOs) and systemic operation (trends/forecasts) regarding the service delivery, day-to-day service activities, and any incidents impacting service performance, overall quality and performance of operational change demands, operational change forecasts, and possible service improvement recommendations. The OSM is delivered remotely and covers:</p> <ul style="list-style-type: none"> <li>• Customer feedback</li> <li>• Incidents, changes, and problems review</li> <li>• Resource and capacity planning</li> <li>• Upcoming operational events</li> <li>• Operational improvement processes (progress/monitoring)</li> </ul>	✓
<p><b>Service planning and review</b></p> <p>The Customer and account team review the service features provided over the previous period and discuss trends, upcoming changes, and potential impacts. It provides an open communication forum to ensure a common understanding of business and IT requirements. During these sessions, the HPE account team may share the best practices and provide tailored advice.</p>	Choice
<p><b>ITSM reporting</b></p> <ul style="list-style-type: none"> <li>• Incident management (number of incidents, age distribution, response time per category, and resolution time per category)</li> <li>• Problem management (number of open cases, age distribution, and response time)</li> <li>• Change management (number of changes in different categories and lead time from registration to commencement of change)</li> <li>• Service request management (number of requests by category, response time by category, and completion time by category)</li> <li>• Release management (releases and planned updates)</li> <li>• Configuration management (managed devices based on CMDB)</li> <li>• Capacity management (Capacity usage against defined thresholds and plans)</li> </ul>	✓
<p><b>System operations and administration reporting</b></p> <p>HPE provides a standard set of reports, with the minimum reported measured as shown here. Specific reporting details will be communicated during the transition period:</p> <ul style="list-style-type: none"> <li>• Overview of systems based on CMDB</li> <li>• Patch and release management</li> <li>• Capacity</li> </ul>	Capacity only



### 4.3 HPE GreenLake integrated system

HPE GreenLake integrated system for Red Hat OpenShift is purpose-built for deploying the OCP solution based on validated configurations. The solution is modular, flexible, and factory built. It can be used to run most common needs that correspond to the OpenShift capabilities and technical use cases such as application modernization, cloud-native applications, and workload consolidation.

HPE GreenLake leverages pre-defined, tested, optimized, and validated cloud modules for OpenShift, delivering fast time to value to our Customers. The solution is designed to provide reliable solid performance and high availability while offering various configurable options such as cloud module instances.

**Table 4.** HPE GreenLake cloud module features

Cloud module	Specifications												
<b>HPE GreenLake cloud module instance:</b>	<ul style="list-style-type: none"> <li>• Compute optimized</li> </ul>												
<b>Compute</b>	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Compute instance code</th> <th style="text-align: left;">Storage type</th> <th style="text-align: left;">RAW capacity</th> <th style="text-align: left;">Memory</th> <th style="text-align: left;">CPU</th> <th style="text-align: left;">GPU</th> </tr> </thead> <tbody> <tr> <td><b>C2apd</b></td> <td>SSD</td> <td>1.6 TB</td> <td>512 GB</td> <td>2 x AMD EPYC 2.8 GHz 32 core</td> <td>NA</td> </tr> </tbody> </table>	Compute instance code	Storage type	RAW capacity	Memory	CPU	GPU	<b>C2apd</b>	SSD	1.6 TB	512 GB	2 x AMD EPYC 2.8 GHz 32 core	NA
Compute instance code	Storage type	RAW capacity	Memory	CPU	GPU								
<b>C2apd</b>	SSD	1.6 TB	512 GB	2 x AMD EPYC 2.8 GHz 32 core	NA								
<b>HPE GreenLake cloud module instance:</b>	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Storage instance code</th> <th style="text-align: left;">RAW capacity options</th> </tr> </thead> <tbody> <tr> <td><b>Bs2</b></td> <td>Choose 1: 25 TB, 50 TB, 100 TB, 150 TB</td> </tr> </tbody> </table>	Storage instance code	RAW capacity options	<b>Bs2</b>	Choose 1: 25 TB, 50 TB, 100 TB, 150 TB								
Storage instance code	RAW capacity options												
<b>Bs2</b>	Choose 1: 25 TB, 50 TB, 100 TB, 150 TB												
<b>Storage</b>													
<b>Network module</b>	✓												
<b>Integrated system management plane</b>	✓												
<b>Platform software</b>	Red Hat OpenShift OS, HPE OneView, and Aruba Fabric Composer												
<b>HPE metering</b>	<p>To measure infrastructure consumption, HPE GreenLake for Red Hat OpenShift solution uses three units of measure:</p> <ul style="list-style-type: none"> <li>• OpenShift: OCP bare-metal subscriptions                             <ul style="list-style-type: none"> <li>– Only applies to the compute instances used as Red Hat OpenShift worker nodes</li> </ul> </li> <li>• Compute: Per core—Measures the average number of cores used per day                             <ul style="list-style-type: none"> <li>– Measures the average number of cores averaging more than 3% CPU consumption per day</li> <li>– Only applies to the compute instances used as Red Hat OpenShift worker nodes</li> </ul> </li> <li>• Storage: HPE Alletra 6000 usable GiB                             <ul style="list-style-type: none"> <li>– The amount of RAW storage needed to store Customer-written data and provide data redundancy</li> <li>– Only applies to the storage instances, does not apply to compute instances</li> </ul> </li> </ul>												
<b>Summary</b>	<ul style="list-style-type: none"> <li>• Offered as a single rack solution for data center deployment</li> <li>• Minimum three worker instances ordered that will be shipped with rack / PDUs / network modules and a management plane</li> <li>• Available configurations are 3, 6, 12, or 16 worker nodes</li> <li>• Add node option up to maximum 16 nodes (increments of 3 or more)</li> </ul>												



#### 4.4 Service implementation

All service levels include the startup services

**Table 5.** Implementation features

Feature	Specifications
<b>Installation and Startup</b>	<p>Provided as part of the service. Before shipping the solution, HPE agrees to the Customer's delivery date, time, and on-site logistics.</p> <p><b>Deliverables:</b></p> <ul style="list-style-type: none"> <li>• A CSS coordinates the factory-build process per your configuration inputs provided before the build and with HPE best practices</li> <li>• The solution is comprehensively factory-tested before shipping</li> <li>• An HPE field delivery specialist installs, powers on, and connects the pre-integrated and configured solution to external peripherals and network components</li> <li>• Hardware validation is performed to help verify that there was no damage during shipping and that the solution is properly connected and functioning in the Customer's environment</li> <li>• HPE remote support and monitoring/management tools are configured, activated, and tested as appropriate</li> </ul>
<b>Service integration with Customer's IT environment</b>	<p>Using HPE for hardware integration and implementing Red Hat OpenShift allows the solution to be integrated into the Customer's IT environment. The services described in the following table are only available for inclusion under an HPE GreenLake services SOW that includes an HPE GreenLake for Red Hat OpenShift integrated system and is not available for purchase separately.</p>
<b>Feature</b>	<b>Delivery specification</b>
<b>Control plane configuration</b>	<p>If necessary, update the configuration of all physical and virtual components in the solution due to changes from what was originally provided by the Customer in the smart customer intent document (smart CID).</p>
<b>Network integration</b>	<ul style="list-style-type: none"> <li>• Use the Customer-provided data collected in the CID to appropriately configure the solution's network switches limited to the requirements of the OpenShift solution</li> <li>• Assist the Customer in validating the north-south connectivity from the solution to northern elements, such as DNS, NTP, limited to the requirements of the OpenShift solution</li> <li>• Help the Customer, when any north-south connectivity issues are identified, in diagnosing them and advising on required actions to resolve the issue</li> </ul>
<b>Red Hat OpenShift implementation</b>	<ul style="list-style-type: none"> <li>• Configure storage for use by OpenShift and the containers therein</li> </ul>
<b>Customer handoff</b>	<ul style="list-style-type: none"> <li>• Conduct an orientation session covering basic interactions with relevant solution elements and any HPE software components the Customer may need to interact with, not to exceed four hours, wherein the Customer is responsible for ensuring attendance during the session; additional HPE service time is available for purchase</li> <li>• Create and provide the Customer with an as-built document</li> </ul>



For the managed infrastructure aaS experience, additional implementation activities and processes are added as described here:

HPE will follow a phased approach to provide a smooth transition to the HPE GreenLake service. In addition to the Customer responsibilities herein, HPE and the Customer will mutually agree upon the necessary activities, along with owners and timelines, to enable HPE to provide the services. These activities will be managed as a discrete project (hereinafter referred to as the implementation project) consisting of the phases and related workstreams described in the following sections, which must be completed before ongoing service delivery commences. The timescales and activities for each phase, including the mutually agreed acceptance criteria for such phases will be documented in an implementation project plan. The implementation period will vary according to the systems in scope and will be defined in the implementation project plan. Before starting the implementation project, the Customer must meet the following prerequisites:

- The Customer must grant HPE exclusive access to its systems
- Further specific assumptions relating to the implementation project are provided in the [Primary Customer responsibilities](#) section.

The implementation project will be split into the following phases. These phases, including any applicable acceptance processes, are more fully described as follows:

- Initiation phase
- Joint verification phase
- Design and implementation phase
- Acceptance into operations phase
- SLO operations and closure phase
- Final acceptance phase

**Implementation project phase acceptance:** Mutually agreed objective acceptance criteria would be defined and documented within the implementation project plan for the three acceptance milestones for the implementation project. These occur in the following phases:

1. At the end of the joint verification phase, acceptance regarding the readiness of the systems for service delivery by HPE will occur. The acceptance process and objective acceptance criteria validate that the systems meet the requirements in the [Primary Customer responsibilities](#) section. The objective criteria on which acceptance of the joint verification phase is based is called infrastructure acceptance criteria.
2. At the end of the acceptance into the operations phase, the objective acceptance into operations criteria for accepting the systems into operations will be based on tests performed on the tools and processes implemented within the individual workstreams.
3. At the end of the SLO operations and closure phase, final acceptance will occur. The parties will mutually agree upon objective acceptance criteria to validate that all processes and tools are in place and operational for HPE to manage the systems. The criteria on which final acceptance is based is called project acceptance criteria.

**Initiation phase:** In the initiation phase, HPE will work with the Customer to define the specific scope, activities, and timelines of the implementation project, which will be documented in the implementation project plan. Project teams and a steering committee will be agreed upon and assembled, followed by an initial kick-off meeting to commence detailed planning.

After completion of the initiation phase, the following phases will occur, except for the connectivity workstream:

**Joint verification phase:** In this phase, HPE will check and validate all assumptions in consultation with the Customer. All systems and related processes will also be assessed and verified based on the following aspects:

- Verification of the systems—quantities, model, type, and CMDB information
- Verification that the system configurations are suitable for meeting SLO targets
- Verification that the systems are accessible, up to date, have required security features, and are documented
- Verification that the systems have the required support/maintenance contracts
- Verification that the environment in which the systems are located conforms to HPE requirements, including temperature, humidity, UPS, security, access rules, and so on
- Validation of the high-level service design (HLSD)

The agreed acceptance process will be performed to verify/check if the systems meet the defined infrastructure acceptance criteria. Each deviation will be documented, and the impact will be defined and listed in a deviations document, with any actions, owners, and timelines identified. If any changes are required because of this activity, they will be addressed through the change management process as detailed in the HPE GreenLake SOW ASP.



**Design and implementation phase:** Specific workstreams will be initiated during this phase of the implementation project as described in the following list, with each having its tasks and timeline documented within the plan. Each workstream will deliver their input to develop the detailed-level service design (DLSD) document, which will be consolidated and documented during this phase with all required information for each workstream as described here.

In the following paragraphs, the workstreams within this phase are described separately.

**Connectivity workstream:** Working with the Customer, HPE will:

- Confirm and detail the connectivity to be implemented to enable the service delivery, including transport protocol and security measures.
- Test and implement the connectivity. Connectivity is known to be on the critical path. This workstream will commence as soon as possible after project initiation.

**Monitoring tooling workstream:** The configuration of the service monitoring tools will be defined and documented within the DLSD. After the Customer and HPE review, the tooling will be implemented according to the DLSD, including event management principles such as event categorization and prioritization. Connectivity must be in place before monitoring can be implemented.

**Service management workstream:** Service management processes and procedures will be documented in the daily agreed procedures (DAP) along with the Customer. These will include interface documents, change type definitions, and specific service-level reporting. Service management tooling will be implemented to support the agreed processes.

**Knowledge transfer workstream:** HPE will work with the Customer to gather necessary information on existing operational documentation, tasks, and activities. The Customer is responsible for providing all information requested by HPE to ensure an effective and efficient knowledge transfer process. It includes deployment, failover, recovery, and security procedures. During the transfer, processes and activities will be tested by HPE with any deviations noted, and required actions documented along with owners and timelines. Information gathered will be documented in the DLSD document.

**Acceptance into operations phase:** Following completion of the design and implementation phase, the acceptance into the operations phase will commence by performing several integral tests under the agreed objective project acceptance criteria for completion of this phase. These tests will be performed by simulation of real-life cases with Customer involvement to validate the end-to-end result of the implementation (simulation sessions). If acceptance does not occur, each deviation will be described, and the impact will be defined and listed in a deviation document. The Customer and HPE will review the deviations and mutually agree upon the next steps to correct any deviations within a reasonable time frame.

Once completed, SLO operations and the closure phase will commence.

**SLO operations and closure phase:** During this phase, HPE will deliver the service based on the SLOs; however, deviations from the SLO may occur. Any discrepancies will be investigated and addressed, with any adjustments made to the service as needed and in agreement between HPE and the Customer. It is intended to be a pre-handover period, with the preceding Customer's operations team still in place to assist where necessary and will typically last for two months. Detailed timelines for this phase will be established during the initiation phase and documented within the implementation project plan.

**Final acceptance phase:** After the successful completion of the SLO operations and closure phase, as determined by HPE, final acceptance testing will take place using the agreed project objective acceptance criteria. Once acceptance occurs, the service implementation project will be completed, and full operations will commence with HPE responsible for providing the service as set forth herein.

## Service implementation output

At a high level, the implementation project described in the previous section will provide the following tasks or deliverables:

- Mutually agreed upon:
  - Project plan and planning
  - Connectivity requirements
  - Objective infrastructure acceptance criteria



- Objective project acceptance criteria
- Infrastructure objective acceptance criteria
- Process interfacing documents
- List of standard changes
- Governance document
- Service-level reports (based upon standard reporting)
- DLSD document
- Successful completion of the following:
  - Kick-off meeting
  - Verification tasks completed
  - Connectivity implemented
  - Monitoring platform implemented
  - Service management tooling configured
  - Knowledge transfer
  - HPE and Customer simulation session
  - HPE Service Desk operations
  - HPE and Customer governance in place
  - Acceptance testing completion and sign-off at identified phases

#### 4.5 Service additions

The following optional extra services can enhance your experience.

##### **HPE GreenLake with colocation—Equinix, CyrusOne, or Digital Realty options**

As an optional service, HPE GreenLake with colocation combines the benefits of a cloud experience while letting Customers control the IT with advantages of colocation, which relieves Customers from the CAPEX and the burden of running a data center on their premises. HPE GreenLake with colocation helps deliver the OpenShift services at any of the colocation partner's location under a single agreement. For more information, see the HPE GreenLake with colocation data sheets:

- [HPE GreenLake with colocation—Equinix data sheet](#)
- [HPE GreenLake with colocation – CyrusOne data sheet](#)
- [HPE GreenLake with colocation – Digital Realty Trust](#)

##### **HPE Education Services**

- A one-year HPE Digital Learner Silver Subscription service can optionally be included with your HPE GreenLake for Red Hat OpenShift purchase. It includes training for one named user seat on all HPE technologies eLearning plus additional select IT eLearning. See [HPE.com/education](https://www.hpe.com/education) for more information.

##### **HPE Cloud Consulting Services**

- As optional services, HPE offers several consulting services that focus on meeting specific Customer objectives. These services are offered as either blocks of consulting hours to meet a limited set of Customer needs, or as a scoped SOW offering to meet the needs of a broader Customer project.
- HU6S0A1: HPE Cloud Consulting 5-day Offsite Service
- HU6S1A1: HPE Cloud Consulting 1-day Offsite Service
- H8Q71A1: HPE Cloud Consulting 5-Day Onsite Service
- H8Q72A1: HPE Cloud SOW Consulting Service





## 5. Service details

### 5.1 Service management levels

There are choices in the degree of service management required to be delivered by HPE, as shown in Table 6.

**Table 6.** Service options

Option	Managed infrastructure aaS
<b>Red Hat OpenShift OS management</b>	
Proactive product support	✓
Advanced configuration and integration	
Installation and startup	✓
<b>Hardware management</b>	
Firmware implementation	✓
Firmware advice	✓
Incident resolution	✓
Proactive product support	✓
Hardware failure monitoring	✓
HPE GreenLake Central onboarding	✓
Advanced configuration and integration	✓
Installation and startup	✓
<b>Control plane</b>	
Installation to integration, operations, and maintenance	✓
<b>Common services</b>	
Consumption analytics	✓
Consumption billing	✓
Hardware capacity planning	✓
HPE technical account management	✓
HPE Digital Learner Silver Subscription service includes training on all HPE technologies plus IT technologies for one learner, one year	✓



### 5.2 Managed infrastructure aaS service-level options

The service levels here will apply upon completion of the final acceptance.

Regardless of the Customer coverage window, incidents with covered hardware or software can be reported to HPE via telephone, web portal, chat, or as an automated equipment reporting event. The Customer can use HPE electronic remote support solution 24 hours a day, 7 days a week.

**Table 7.** Service levels

Service-level options	Service feature	Coverage window	Feature description
<b>Critical</b>	Enhanced phone response	Remote response 24x7; service is available 24 hours per day, 7 days per week, including HPE holidays	15-minute call back for priority 1 incidents, 1 hour for priority 2 and 3; where available direct phone access to product specialists without the need for a call back (all priorities)
	6-hour hardware call-to-repair <sup>2</sup>	On-site response 24x7; service is available 24 hours per day, 7 days per week, including HPE holidays	For priority 1 and 2 incidents, HPE returns the coverage hardware to operating condition within 6 hours <sup>3</sup>
	Outage management	Remote response 24x7; service is available 24 hours per day, 7 days per week, including HPE holidays	Available for severity 1 business impacting situations, HPE provides priority access to incident recovery specialists to expedite return to service
<b>Essential</b>	Enhanced phone response	Remote response 24x7; service is available 24 hours per day, 7 days per week, including HPE holidays	15-minute call back for priority 1 incidents, 1 hour for priority 2 and 3; where available direct phone access to product specialists without the need for a call back (all priorities)
	24x7 on-site coverage	On-site response 24x7; service is available 24 hours per day, 7 days per week, including HPE holidays	4-hour on-site response <sup>4</sup> for covered hardware

All service levels provide 24x7 access to online self-service and self-solve capabilities. You also get 24x7 incident logging, for supported devices, along with 24x7 HPE InfoSight analytics and automated incident submission.

<sup>2</sup> Hardware call-to-repair time begins when the initial incident has been received and acknowledged by HPE or at the start time for work scheduled in agreement with the Customer, as specified in the hardware call-to-repair section. Hardware call-to-repair time ends with HPE's determination that the hardware is repaired, or when HPE has determined that no on-site intervention is required. For hardware cases originating from software incidents, call-to-repair time begins when HPE has made the determination that the cause is attributable to the covered HPE hardware.

<sup>3</sup> Hardware call-to-repair time begins when the initial incident has been received and acknowledged by HPE or at the start time for work scheduled in agreement with the Customer, as specified in the hardware call-to-repair section. Hardware call-to-repair time ends with HPE's determination that the hardware is repaired, or when HPE has determined that no on-site intervention is required. For hardware cases originating from software incidents, call-to-repair time begins when HPE has made the determination that the cause is attributable to the covered HPE hardware.

<sup>4</sup> On-site response time begins when the initial support incident has been received, acknowledged, and confirmed to be a hardware issue by HPE. The on-site response time ends when the HPE authorized representative arrives at your site, or when the reported event is closed with the explanation that HPE has determined that no on-site intervention is required.



### 5.3 Managed infrastructure aaS service-level objectives

The following SLO will apply upon completion of final acceptance.

**Table 8.** SLOs

Event management		Service level	
<b>Hours of coverage</b>		24x7	
<b>Incident management</b>			
Response matrix	Description	Incident response time <sup>5</sup> objective	Incident resolution time <sup>6</sup> objective
Priority 1	Critical	90% < 15 min	90% < 4 hours
Priority 2	High	90% < 1 hour	90% < 8 hours
Priority 3	Medium	90% < 4 hours	90% < 16 hours
Priority 4	Low	90% < 8 hours	90% < 24 hours
<b>Problem management</b>		<b>SLO</b>	
Hours of coverage		Standard workday	
Time to initiate problem investigation		5 standard workdays from problem registration	
<b>Change management</b>		<b>SLO</b>	
Hours of coverage		Standard business hours	
Hours of coverage (emergency change)		24x7	
Commence normal change planning		5–7 business days	
Commence standard change planning		2 business days	
Commence emergency change follow-up and documentation update		Within 24 hours	
<b>Service request management</b>			
Priority	Service window	Response time objective	Completion time objective
Priority 1/2	N/A; priority 1 or 2 situations should be logged using the incident management process.		
Priority 3	Monday to Friday 08:00–17:00	4 hours	24 hours
Priority 4	Monday to Friday 08:00–17:00	8 hours	48 hours

<sup>5</sup> Incident response time is a target and is measured as the time elapsed from when an incident is first raised/opened in the respective ticket flow tool to when there is a response/acknowledgement from HPE registered in the same tool. Response time begins when HPE creates a support case in the record system and actively starts to troubleshoot and remediate an incident.

<sup>6</sup> Incident resolution time is a target and is measured as the time taken for HPE to resolve an incident. It excludes the SLO exclusions set forth here, as well as any time attributable to Customer or any third parties with whom HPE engages to resolve the incident, or time attributed to address any hardware- or software-related incidents. It is regardless of the vendor, dependent upon the underlying service agreement, and related service level with the vendor. Resolution may depend on and include the implementation of resilience measures or configurations as verified by HPE in the joint verification stage of the implementation project.



**Table 8.** SLOs (continued)

<b>Event management</b>		<b>Service level</b>
<b>Hours of coverage</b>		24x7
<b>Service request management</b>		
<b>Release management</b>	<b>SLO</b>	
Hours of coverage	Standard workday	
Frequency of operating system and system software review	Quarterly or as documented in the ASP	
<b>Configuration management</b>	<b>SLO</b>	
Hours of coverage	Standard working day	
Update CMDB following CI change	90% < 1 standard working day	
Update documentation	90% < 2 standard working days	
<b>Capacity management</b>	<b>SLO</b>	
Hours of coverage	Standard working day	

**SLO exclusions**

The following are excluded from SLOs:

- Delays in the Customer approval process
- Incidents due to Customer’s applications, hardware, software, services, or facilities
- Events due to Customer WAN/LAN-related issues
- Force majeure at Customer or HPE site
- Any act or omission on the part of Customer, its contractors or vendors, or any other entity over which Customer exercises control or has the right to exercise control, including Customer’s failure to perform its obligations or responsibilities
- Factors outside HPE’s reasonable control
- Planned outages and scheduled maintenance
- Interruptions or incidents not reported by Customer or where no ticket was opened
- Incidents occurring before the 31st day from the effective date of the change order for any systems added to the services following the change management process
- Incidents occurring at the operating system level and above (for IaaS managed infrastructure)

**5.4 IT service management**

ITSM refers to the implementation and management of quality IT services that service management tools enable, and are based on ITSM norms and best practices. ITSM is not provided with managed infrastructure aaS. Each management topic in this section includes both a description and a table that shows the specific HPE and Customer activities that will be performed.

**Event management:** HPE will provide event management services that will receive, respond to, categorize, and log events generated by the systems. According to its nature, events that require additional action will be logged as incidents or request for change (RFC). The following table shows the specific HPE and Customer activities that will be performed.

Legend for tables in the document: **R = Responsible, A = Accountable, C = Consulted, I = Informed**



**Table 9.** Event management

Activity	Customer	HPE
Design service monitoring	C	A/R
Monitoring setup	C	A/R
Monitoring and events identification	I	A/R
Alarm identification and management	I	A/R

**Incident management:** HPE will implement an incident management process that will respond to and resolve Incidents related to the systems. The process will manage an incident throughout its lifecycle through closure by HPE, including incident registration, categorization, prioritization, investigation and diagnosis, incident resolution, and closure. The process will encompass communications and dialogue with the Customer throughout the life of the incident to resolve the incident in an agreed manner.

**Table 10.** Incident management

Activity	Customer	HPE
Verify log, categorize, and prioritize incident	R	A/R
Conduct initial diagnosis	I	A/R
Investigate and diagnose incident	I	A/R
Resolve incident and recover service	I	A/R
Close incident	C	A/R
Manage major incidents	C	A/R
Handle infrastructure-related security incidents	C	A/R

Manage security incidents will only include security on the integrated system (IaaS managed infrastructure). OS level and above security is not included in managed infrastructure.

**Problem management:** HPE will implement a problem management process to address repeated or priority 1 incidents, identified by HPE within the systems. The process will consist of problem identification, registration, root cause analysis, potential workaround, corrective action, and reporting of identified problems.

**Table 11.** Problem management

Activity	Customer	HPE
Identify problem record	R	A/R
Classify and verify problem	I	A/R
Investigate and diagnose problem	C	A/R
Resolve problem	I	A/R
Close problem record	C	A/R

Recommendations resulting from problem management activities will be logged as requests for change and handled according to the change management process.



**Change management:** HPE will implement a change management process relating to changes to be made to the systems. The process will coordinate HPE’s activities in relation to the implementation of those changes. The Customer is expected to maintain ownership of change evaluation and authorization, including change advisory board (CAB) facilitation. If changes to the systems impact the scope of the services described in this data sheet, the changes will be subject to the change management process as detailed in the account support plan.

**Table 12.** Change management

Activity	Customer	HPE
Initiate RFC	R	A/R
Enable RFC assessment / initial authorization	I	A/R
Change planning	I	A/R
Change approval	R	A/R
Change schedule	C/I	A/R
Build/test	I/R	A/R
Implement	I	A/R
Close	I	A/R

**Service request management:** HPE will implement a service request management process to provide for the implementation of standard predefined changes that do not need to follow the full change management process or to process common requests that may occur.

**Table 13.** Service request management

Activity	Customer	HPE
Initiate service request	R	A
Define generic service request	C	A/R
Approve request	R	A
Fulfill request	I	A/R
Close service request	I	A/R

**Release management:** HPE will monitor software releases and their applicability to the systems, contingent upon the agreements between the Customer and the relevant software vendors, and according to the operating system patch analysis and management services as documented in the account support plan.

Release management applies to the following types of software if included in the systems:

- **System software:** Refers to technical software implemented to run and/or support the infrastructure platform. Examples are (storage) system firmware, and so on

Release management services are dependent on all products installed in the systems being covered within an active maintenance contract that provides access to the system and operating system documentation and updates.



**Table 14.** Release management

Activity	Customer	HPE
Identify release	I	A/R
Validate release	I	A/R
Approve release	I	A/R
Deployment	C	A/R

**Configuration management:** HPE will create and maintain a CMDB that will be used to deliver services and facilitate ITSM activities. The CMDB will contain relevant configuration information relating to the systems. CMDB information will be available to the Customer, and standard reports can be generated for Customer use upon request, or as defined during the implementation project. However, the CMDB should not be considered as a replacement for a more comprehensive CMDB encompassing the whole of the Customer's IT environment.

**Table 15.** Configuration management

Activity	Customer	HPE
Provide scope, criteria, requirements, and updates for CMDB structure design	I	A/R
Upload configuration items in CMDB from HPE	I	A/R
Provide relationship information	C	A/R
Baseline CMDB from HPE	I	A/R
Give CMDB updates (new records or records updates)	R	A/R
Define audit scope	A/R	C
Perform audit and provide status	C	A/R
Plan improvements	C	A/R

**Capacity management:** Based on the Customer's capacity requirements and key performance indicators (KPIs), HPE will maintain a capacity plan relating to the systems. It will be based on Customer analysis of their capacity requirements concerning business demand.



**Table 16.** Capacity management

Activity	Customer	HPE
<b>Manage business capacity input</b>	A/R	I
<b>Handle Customer services capacity input</b>	A/R	C
<b>Define the scope</b>	A/R	C
<b>Document, maintain, and communicate the overall capacity plan</b>	A/R	I/C
<b>Document and communicate initial infrastructure (in scope) capacity plan: Utilization and performance measures, thresholds, tools</b>	A/R	I/C
<b>Maintain infrastructure (in scope) capacity plan during operations</b>	I/C	A/R
<b>Recommend configuration updates or related changes intended to help improve utilization and performance</b>	C	A/R
<b>Approve recommendations and trigger respective actions</b>	I/C	A/R

**5.5 Metering tools**

- HPE will provide metering tools to measure the systems’ usage for billing.
- The HPE metering tools collect the data needed to determine usage (for example, server on/off or virtual machine memory consumption, physical core utilization percent, or written GB.)
- All metering happens inside the Customer’s data center firewall, and an encrypted communication (email or API) with usage data is sent daily to HPE.
- Metering of the systems will begin when HPE’s installation of the metering tools is complete.
- HPE will be responsible for ongoing administration, upgrades, and maintenance of the metering tools. This will include the identification and installation of metering tool patches.
- Each calendar quarter, with three business days advance notice, HPE may inspect, audit, and test the metering scripts, all usage data, and systems to validate proper usage data.
- Inspections may be immediate in case of an apparent metering script failure, a discrepancy in usage data, or suspected manipulation.

**5.6 System operations and administration**

HPE will remotely monitor and manage the infrastructure, performing certain operational and administrative activities. These activities include:

- Health status and availability monitoring
- Checking of critical infrastructure log files
- Creation and maintenance of standard operating procedures (SOPs)
- Performing regular scheduled maintenance procedures according to agreed SOPs
- Capacity reporting
- Capacity additions
- Firmware management
- System/device configuration changes

Operational and administrative activities are designed to be delivered remotely. However, to meet certain Customer operational requirements, some activities may be provided by HPE resources located on-site within the Customer’s premises. Any on-site activities will be identified in the HPE GreenLake SOW.





## 6. System operations and administration activities

### 6.1 Managed infrastructure aaS

The following table outlines the specific activities that HPE will perform to manage the Red Hat OpenShift solution.

**Table 17.** System operations and administration activities

Activity	Customer	HPE
Monitor health status of a Red Hat OpenShift cluster (head nodes, worker nodes, filesystems, and performance metrics)	A/R	
Supervise running services and changed state of services	A/R	
Check log files on critical events/incidents	A/R	
Respond to infrastructure security events/incidents	I	A/R
Add or remove the nodes for a Red Hat OpenShift head and worker nodes	A/R	
Start or stop a cluster	A/R	
Manage cluster settings	A/R	
Perform node maintenance (restart Red Hat OpenShift cluster nodes or offline for maintenance)	A/R	
Update Red Hat OpenShift clusters	A/R	
Update software-defined network infrastructure for Red Hat OpenShift	A/R	
Implement a Red Hat OpenShift stretched cluster (future)	A/R	
Manage Red Hat OpenShift stretched cluster operations (future)	A/R	
Create volumes, create stretch volumes, protect volumes, expand volumes, delete volumes, replace drives	A/R	
Provide incident management and coordination with break-fix support teams	I	A/R
Address performance incidents within the scope of HPE responsibilities	I	A/R
Monitor health alerts of Red Hat OpenShift platform	A/R	
Supervise CPU, memory, network, disk usage, and performance of Red Hat OpenShift platform	A/R	
Monitor uptime of Red Hat OpenShift platform	A/R	
Respond to Red Hat operating system and OpenShift security events/incidents	A/R	
Troubleshoot head node and worker nodes in OpenShift	A/R	
Manage the hardware resources available to cluster nodes, such as networking, memory, and storage	I	A/R
Upgrade Red Hat OpenShift version	A/R	



## 7. Service prerequisites

- Customer must maintain active support contracts through HPE, for all operating systems and software products within the systems being supported during the agreement/SOW term.
- All operating systems to be maintained are under manufacturer general availability support.

## 8. Primary Customer responsibilities

In accordance with the [HPE GreenLake Terms data sheet](#), the following are responsibilities of the Customer in relation to this service

### General:

- Designate a senior-level single point of contact (SPOC) who will be authorized to act as the primary contact in dealing with HPE and other internal stakeholders, including:
  - Assigning the necessary stakeholders during project transition and duration
  - Identifying a focal point to work collaboratively with the HPE account team
  - Responsibility for all Customer aspects of the services covered by the overall HPE GreenLake SOW
  - Authorized to make decisions relative to the data sheet / SOW, including identification and assignment of Customer resources
  - Certified to approve changes to the overall HPE GreenLake SOW
  - Sign and return the mutually agreed HPE GreenLake SOW for which this data sheet is appended, along with a purchase order for the Services to HPE
- Provide HPE with a contact list that contains at least one primary and one backup contact who will be the Customer's point of contact during all operational support coverage hours as described in the service levels section of this data sheet
- Give network details needed for factory integration
- Offer the current network architecture, standards, and detailed design documentation
- Give HPE access to network services such as NTP, DNS, default gateways, routes, and remediation, if needed
- Provide Active Directory credentials for commissioning and integration of the solution, if needed
- Maintain a separate backup system
- Is responsible for the security of their proprietary and confidential information, as well as properly sanitizing or removing data from products that may be replaced and returned to HPE as part of the repair process to ensure the safeguarding of Customer data; For more information on Customer responsibilities, including those outlined in [HPE Media Sanitization Policy and Media Handling Policy for Healthcare customers](#)
- Allow HPE VPN connectivity to their network so that HPE can manage the Customer's on-premises HPE GreenLake environment.
- Ensure data collected by the diagnostic and metering tools are sent daily:
  - SMTP: 7-Zip encrypted file attachments of usage data; Customer must notify HPE if a change is made to the email relay, network, or proxy to help eliminate any interruption of usage data collection
  - Or
  - An encrypted API connection
- Make modifications to the existing network that are required and identified during the planning stages of this service, prior to arriving on-site to perform configuration and integration tasks



- Cover Customer affiliates through an absolute and unconditional guarantee, or ensure it is separately credit approved by HPE
- Ensure Customer's rights and obligations under the services agreement are not assignable without HPE's consent
- Implement Customer services agreement
- Make sure All HPE integrated systems in a Red Hat OpenShift cluster are included in the HPE GreenLake agreement; Customers are prohibited from moving or modifying the integrated system hardware once deployed to a cluster
- Ensure that site preparation (for example, power and cooling) has been completed at the location where the hardware will be installed
- Make certain that the access path can accommodate rack height, weight, and clearance requirements for delivery of the racked solution

**Service implementation:**

During the implementation project, the Customer will:

- Provide HPE with physical access to site locations for initial installation and testing of remote connectivity software and/or hardware
- Offer all equipment on the Customer site relevant to provision the services
- Be deemed to have accepted the systems upon loading of Customer data, or if HPE has not received written notice of problems within five business days after the completion of the implementation project
- Schedule and receive delivery of these services within 180 days from order acceptance; HPE reserves the right to reprice for services not scheduled and delivered in 180 days

**9. HPE responsibilities**

- HPE will deliver the tasks outlined in this data sheet.

**10. General provisions / Other exclusions**

- The HPE GreenLake for Red Hat OpenShift services are limited to the systems as provided under the HPE GreenLake SOW to which this data sheet is incorporated.
- This service applies to new and expansion implementations of HPE GreenLake for Red Hat OpenShift (configuration options listed in Table 1) only and does not address modification of any configurations of any equipment that is not part of the HPE GreenLake for Red Hat OpenShift solution.
- The services (or portions thereof) will be provided from locations determined by HPE, which may be outside the country of purchase.
- HPE GreenLake for Red Hat OpenShift is governed by HPE standard terms for professional services as part of these services.
- This service does not include re-racking or reconfiguration of any work performed prior to shipment, and does not include planning, design, reconfiguration, implementation, or assessment of the Customer's existing LAN, WAN, or SAN environment (However, these are available from HPE as separate services.).
- HPE will notify Customer of any location change, if applicable, during the term of the SOW. In addition to the metering tools, HPE will install certain hardware and software tools to deliver remote monitoring services. HPE owns such tools, and delivery of these services is contingent upon its installation. Customer may not use, transfer, assign, pledge, or in any way encumber or convey the tools. HPE will remove the tools upon termination or expiration of the SOW. Customer acknowledges that it is responsible (administratively and financially) for obtaining all required approvals, licenses, authorizations, consents, and permits for HPE to perform the HPE GreenLake Management Services.



- All deliverables are accepted upon delivery unless otherwise specified within the SOW.
- Customer acknowledges that the ability of HPE to provide these services is contingent upon the accuracy and completeness of information and data provided by Customer, as well as Customer's cooperation and timely performance of its obligations. In the event, any such data is found to be inaccurate or incomplete or Customer fails to perform their obligations, the parties agree to negotiate in good faith equitable changes, which may include, without limitation, changes to the charges to Customer for the services.
- Service is provided in the English language.
- Our ability to deliver this service is dependent upon the Customer's full and timely cooperation with HPE and authorized partners, as well as the accuracy and completeness of any information and data Customer provides to HPE.
- HPE and authorized partners assume that all information provided by the Customer is accurate and will collaborate with the Customer to determine acceptable estimates for any information that is not available.
- Equipment will be specifically identified in the technical configuration / bill of materials (BOM) of the SOW this data sheet is applied to.
- HPE owned equipment will reside at a Customer-controlled facility or a facility controlled by HPE. If not, waiver and consent to be obtained from third party.
- Should the Customer notify HPE of the intention to terminate more than 90 days before the end of the term, the Customer will owe HPE all monthly payments plus any other applicable fee charges through the end of the system term or until system equipment is returned, whichever is greater.
- HPE will manage the integrated system infrastructure using integrated system management plane.
- Service does not include integration of the tools with any Customer systems, for example, ticket transfer. This functionality is available to be scoped as a separate project.
- MSM activities will be delivered remotely and are provided during standard HPE business hours.
- HPE is not responsible for any service activities relating to Customer applications and workloads.
- ITSM change management activities provide implementation of changes to the systems requiring up to eight hours of labor time. Requests requiring additional effort will be treated as projects and require separate scoping and funding.
- Incidents requiring product support contracts (from HPE or other vendors) are subject to the service levels associated with those support contracts.
- At HPE's discretion, based on the hardware type or the complexity of the configuration, a site environmental requirements document that focuses on power, cooling, space requirements, and network connections may be sent to the Customer. The document must be verified so that all requirements are met before solution delivery. If required, an HPE field delivery specialist will then work with the Customer to verify that all environmental requirements have been met.

## 11. Ordering and pricing information

The price of these services is incorporated into the monthly unit price as set forth in the SOW to which this data sheet is incorporated. Service ordering and invoicing are also according to the terms defined within the ordering and pricing section of the HPE GreenLake SOW, into which this data sheet is incorporated.



## 12. Defined terms

Term	Definition
<b>Call</b>	A notification from the Customer to the HPE Service Desk regarding an incident relating to the systems
<b>Emergency change</b>	Urgent RFCs to the systems for which the normal change procedure does not meet the requirements; late submission of a change does not constitute an emergency change
<b>Event</b>	An occurrence within the systems and within the scope of the services as observed by HPE, or because of a call that has relevance to either the Customer as the user of the systems or HPE in providing these services
<b>General availability</b>	The release of a product that is available to the general public including the availability of associated vendor support products
<b>Incident</b>	An event within the systems that results in an unplanned interruption or degradation of the functionality provided by the systems or that has not yet affected the level of functionality provided by the systems (for example, failure of one disk from a mirror set)
<b>Incident resolution</b>	Action taken to resolve an incident and restore the affected functionality or to implement a workaround
<b>Incident resolution time</b>	The time taken for HPE to resolve an incident, excluding the SLO exclusions identified in the <a href="#">General provisions / Other exclusions</a> section of this document, and any time attributable to Customer or any third parties with whom HPE is required to work with to resolve the incident; incident resolution time is a target only and may depend on and include the implementation of resilience measures or configurations as verified by HPE in the joint verification phase
<b>Minor change</b>	Changes that have a low-risk factor (as defined in the change management process) and are unlikely to impact service levels
<b>Normal change</b>	RFCs that must follow the complete change management process
<b>Problem</b>	A problem is the unknown cause of one or more incidents, often identified as a result of multiple similar incidents
<b>Request for change (RFC)</b>	A request specifically and only for a change to the systems that may lead to a change in the composition or configuration; five categories of RFCs are distinguished: <ul style="list-style-type: none"> <li>• Standard change</li> <li>• Normal change</li> <li>• Minor change</li> <li>• Major change</li> <li>• Emergency change</li> </ul> An RFC is requested using the RFC form
<b>Response</b>	A response is when HPE contacts the incident initiator or actively starts to work on the incident
<b>Response time</b>	Response time is a target and is measured as the time elapsed from when an incident is first raised to when there is a response
<b>Service</b>	The HPE GreenLake services, detailed in this agreement, that HPE will perform for Customer
<b>Standard change</b>	Predefined and preapproved RFCs with manageable risk, routine tasks, clearly specified, and with standardized implementation
<b>Standard workday</b>	During standard local business working days and hours, excluding HPE holidays. Specifics will be defined during the implementation project
<b>System</b>	System refers to the infrastructure components, devices, and software products, which comprise the system provided by the related HPE GreenLake SOW
<b>System availability</b>	The service-level availability means with respect to the systems expressed as a percentage of scheduled uptime for the relevant infrastructure component (that is, availability = [scheduled uptime minus unexcused downtime] divided by scheduled uptime)
<b>Workaround</b>	Reducing or helping eliminate the impact of an incident for which a full resolution is not yet available



Data sheet

## Learn more at

[HPE.com/us/en/GreenLake/governance-and-management.html](https://hpe.com/us/en/GreenLake/governance-and-management.html)

Make the right purchase decision.  
Contact our presales specialists.



Chat now (sales)



Call now

Visit **HPE GreenLake**



Get updates

  
**Hewlett Packard  
Enterprise**

© Copyright 2023 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

AMD is a trademark of Advanced Micro Devices, Inc. Active Directory is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. All third-party marks are property of their respective owners.

a50008215ENW, V1