

plusserver Kubernetes Engine (PSKE)

Service description

Version: EN – v1.0.3
Current as of 13. July 2022

Preamble

The following describes essential services which plusserver makes available to its customers within the framework of the "plusserver Kubernetes Engine" (hereinafter "PSKE"). Special conditions, rights and / or restrictions apply to certain services. These are also described in the following description of the services concerned or reference is made there to such special conditions, rights and/or restrictions. The concrete services to be provided to the customer depend on the scope of services contractually agreed between plusserver and the customer in the individual case.

Plusserver continuously updates the offered services and adapts them in particular to the state of the art. Accordingly, plusserver will update this service and product description from time to time. Changes to this service description during the term of an existing contract will take effect in accordance with the agreements made with the customer.

1. General description

"plusserver Kubernetes Engine" uses Kubernetes as an orchestration layer for container environments. Kubernetes is a scalable open source system for the automatic provision and scaling of container applications (services). Official Kubernetes distributions are supported and used by plusserver as a basis.

The basic topology consists of a redundant management system that provides the corresponding Kubernetes functions. The IT systems required for management are provided centrally on the cloud infrastructure of plusserver or another cloud provider supported by the "plusserver Kubernetes Engine" (e.g. AWS) and managed by plusserver. The size and number of customer-owned Kubernetes environments can be freely defined.

The Kubernetes clusters (workers) are provided on the basis of the compute resources of plusserver or another cloud provider supported by the "plusserver Kubernetes Engine" (e.g. AWS), which serve as the basis for the compute resources required within the framework of Kubernetes (in particular processor performance, storage, network and security resources).

Access is via a management system (e.g. Gardener), which can be operated for management either via a graphical interface or via API. The administration of Kubernetes clusters created via this system is done exclusively via API.

Plusserver does not assume any support or responsibility for the applications, data, workloads or clusters provided by the client.

2. Essential performance features

- Administrative access to the Kubernetes cluster via GUI, CLI or API
- Use of the Kubernetes cluster based on the compute resources of plusserver or the respective cloud provider (e.g. AWS)
- Persistent storage
- Usage-based billing

3. Technical details / performance features

Resources based on the compute resources of plusserver or another cloud provider supported by the "plusserver Kubernetes Engine" (e.g. AWS) are provided for the client and used for Kubernetes. The client gets access to the self-service environment, which in turn uses orchestration software (e.g. Gardener) to manage its Kubernetes clusters.

Management system: The components of the management system are provided redundantly and offer the basic Kubernetes services required.

Workers based on the compute resources of external supported cloud providers: The client is free to choose the infrastructure resources required for the workers from other cloud providers supported by the "plusserver Kubernetes Engine".

Workers based on plusserver compute resources: The workers are provisioned by the client via the management server and host the client workloads.

a. Processor, working memory

A virtual server consists of the resources listed here with the respective minimum and maximum values.

Resource	Minimum	Maximum
Processor	2 vCore	16 vCores
Working memory	4 GByte	64 Gbyte

Processor power and main memory are considered to be actually used if they are allocated in switched-on instances.

A VCPU has at least 2 gigahertz clock frequency.

b. Volumes

Hard disks (volumes) are provided with the following scope of services:

Value	Datastore
IOPS/TB	4.000
Latency	< 3 ms

For the use of the hard disks (volumes), the storage period applies to the size of all configured hard disks regardless of the operating state of the associated virtual machine.

c. Persistent Storage

Persistent storage capacities are also connected via the resources.

d. IP addresses

Customers can consume additional public IPv4 addresses or reserve fixed subnets for a fee.

The number of IP addresses ordered by the client is provided by plusserver at the router in the client's project and billed according to consumption. Floating IPs are considered to be actually used if they are assigned to a port of an instance, router or load balancer.

Unless otherwise agreed, the configuration of the internal network behind the router is carried out by the customer. This applies both to the segmentation of the internal networks and to the addressing. The configuration of further private networks is freely possible up to the specified maximum.

The number of configurable project-internal subnets is limited to 100, a larger number requires a separate release.

e. Loadbalancer

Load balancers can optionally be used within a project to distribute the load of network connections.

The configuration of the load balancers provided is carried out by the customer, unless otherwise agreed. This applies both to the creation of members of the distribution, monitoring and addressing.

f. Internal traffic

Data traffic exclusively behind the router, i.e. within the customer's project, is included and does not have to be remunerated separately by the customer.

g. Connection to the internet

A client project can be connected to the Internet via the router. For this purpose, a connection of the router is created via the pluscloud open infrastructure to the central uplink switches of plusserver. This connection can be used via the floating IP addresses assigned to the router.

The traffic routed to the internet via the routers is measured on an aggregated basis over the month as "traffic outgoing".

4. Necessary cooperation of the client

For the use of "plusserver Kubernetes Engine", it is necessary that the client provides the following cooperation services:

- The customer is responsible for rolling out the configuration and the resources contained in a workload and for operating the workload. Plusserver assumes no responsibility for this.
- The customer is responsible for regularly creating a backup of his own workload.

5. Software used / applicable licence conditions

- The licence conditions of any software used shall, insofar as relevant for the user, be contractually bound to the user.
- Relevant licence terms and conditions of use are shown separately
- Software components of the product (e.g. Gardener and Kubernetes) are open source software and are made available under the Apache 2.0 licence.

6. Remuneration model /tariffs/breakdown type

a. Accounting Kubernetes Cluster

Billing takes place dynamically for each Kubernetes cluster created by the customer and monthly in arrears. The contractually agreed fees are charged per Kubernetes cluster.

b. Billing of workers on the basis of plusserver compute resources

Worker nodes require resources that are obtained from plusserver on the basis of compute resources.

- The billing of resources for provisioned workers on the basis of the compute resources of plusserver takes place in principle according to actual use (on demand) and monthly in retrospect. The contractually agreed fees are charged per resource.
- The use of vCPU and RAM resources is recorded to the minute; billing is rounded to the full hour. All resources that are assigned to switched-on instants are recorded.
- The billing of resources for hard disks / volumes is done according to the monthly average used. All used storage areas are recorded regardless of the operating status of an associated VM.
- The billing of data traffic to destinations on the Internet outside the plusserver network ("outgoing traffic") is calculated according to gigabytes, depending on use. Data traffic from any source to the client's project is free of charge.
- Optionally, minimum purchases (commitments) can be contractually agreed for the resources mentioned.
- The provision and billing of load balancing and routing instances is carried out dynamically according to the maximum value used in parallel within a billing period.

c. Commitments

If a minimum purchase (commitment) has been agreed for individual services, the customer is obliged to pay for these services in the amount of this minimum purchase quantity, even if the actual use was less than the agreed period in each case. Parts of the Commitment not used in a month shall expire. If the customer has used services in excess of an agreed commitment in a billing period, this excess use will be billed at the agreed rates.

d. Billing of workers on the basis of external compute resources of the cloud providers supported by the "plusserver Kubernetes Engine"

If the integration of external cloud resources of another cloud provider supported by the "plusserver Kubernetes Engine" (e.g. AWS) is possible, the client can generally choose between the following options:

- The billing of the compute resources of another cloud provider supported by the "plusserver Kubernetes Engine" is carried out via plusserver. This requires an additional contractual agreement between plusserver and the client regarding the services of the respective cloud provider.
- The billing of the compute resources of another cloud provider supported by the "plusserver Kubernetes Engine" takes place directly via a contract between the respective cloud provider and the client.

e. Other plusserver services

Separate remuneration models, tariffs and/or billing methods may apply for other plusserver services. Details are contained in the relevant service descriptions or in the service overview relevant for the customer.