

**Service certificate**

**STACKIT**

**RabbitMQ**

**Version and start of validity**

---

Version 1.5	Valid from 2025-09-12
-------------	-----------------------

---

# Service certificate | STACKIT RabbitMQ

## Service name

STACKIT RabbitMQ

## High level service description

The STACKIT RabbitMQ Service (“RabbitMQ”) provides fully managed instances of the open-source multi-protocol messaging broker RabbitMQ. Customers can provision instances of RabbitMQ in various sizes and use them in applications. RabbitMQ instances are deployed on dedicated virtual machines to ensure isolation and high performance. Customers can choose between replica plans that offer resilient operation with a low level of downtime for use in productive environments and single instances, which we do not recommend for use in productive environments. Customers can meet increasing demand by upgrading to larger instances (service plans). Smoke tests, service guards, and the option of automated backups via Backup Manager ensure the secure operation of RabbitMQ’s service instances as well as the stored data.

## Key features

- On-demand provisioning: Easy and fast provisioning of new service instances of different sizes
- Service Instance Insulation: Deployment on dedicated virtual machines (Bad Neighborhood Protection)
- Easy capacity upgrades via service plan upgrades
- High availability: Possibility of using replicas for increased resilience and fail-safety
- Logging and monitoring: Communication of logging and monitoring information to defined end points
- Backup: Automatic creation and restoration of backups over a period of up to 14 days

## Service plans

Within the framework of the order, the customer has the option to choose from different service configurations. These mainly differ in their performance capabilities or storage capacity.

Service Plan	High availability	Number of nodes	vCPU per node	RAM (GB) per node	Disk (GB) per node
RabbitMQ-1.2.10-Replica	yes	3	1	2	10
RabbitMQ-2.4.10-Replica	yes	3	2	4	10
RabbitMQ-2.4.50-Replica	yes	3	2	4	50
RabbitMQ-4.8.10-Replica	yes	3	4	8	10
RabbitMQ-4.8.50-Replica	yes	3	4	8	50
RabbitMQ-4.8.100-Replica	yes	3	4	8	100
RabbitMQ-8.16.50-Replica	yes	3	8	16	50
RabbitMQ-8.16.100-Replica	yes	3	8	16	100
RabbitMQ-16.32.100-Replica	yes	3	16	32	100
RabbitMQ-2.4.10-Single	no	1	2	4	10
RabbitMQ-4.8.50-Single	no	1	4	8	50
RabbitMQ-8.16.100-replica-HiPerf	yes	3	8 <sup>1</sup>	16	100
RabbitMQ-4.8.100-replica-HiPerf	yes	3	4 <sup>1</sup>	8	100
RabbitMQ-16.32.100-replica-HiPerf	yes	3	16 <sup>1</sup>	32	100

<sup>1</sup> dedicated CPU for highest performance

A subsequent downgrade of the service plan initially selected by the customer is not possible.

## Metrics

- Billing per hour or part thereof per commissioned RabbitMQ instance

## SLA specifics

- A booked instance of RabbitMQ is considered available if it can be accessed at the service delivery point.
- The provisions of the STACKIT Cloud general service description apply to the availability of RabbitMQ instances.
- To achieve the highest possible availability of RabbitMQ (for example, for use in production environments), the use of replicas (3 nodes) is strongly recommended.
- When using RabbitMQ single-instance service plans, RabbitMQ may become (temporarily) unavailable, especially during maintenance work (e.g., through updates) that STACKIT performs on the underlying platform as well as on the selected service plans. Maintenance work is considered an excluded event within the meaning of section 2.4 of the STACKIT Cloud service description and is therefore not considered as downtime in the context of availability. It is strongly discouraged to use RabbitMQ with Single Instance Service Plans in production environments.
- Incorrect use of the STACKIT Cloud Portal by the customer with regard to the settings made in self-service and a resulting unavailability of RabbitMQ will not be taken into account with regard to the actually achieved availability of RabbitMQ. The customer is advised that the instance and the associated data may be lost if RabbitMQ is set up and used incorrectly.

## Backup

Instances of RabbitMQ are backed up automatically according to a schedule specified by STACKIT in line with the STACKIT Cloud general service description. This schedule cannot be changed by the customer. However, customers have the ability to create manual backups through the service dashboard and to restore automatic and manual backups.

## Additional terms

- Creating manual backups, as well as restoring backups of RabbitMQ instances, is the responsibility of the customer.

# Annex | Exportability

## (Online Register)

Data type	Description	Exportable (Yes/No)	Format	Additional notes
Customer data (database content)	Data stored by the customer in the database (if available) or within the product/service	Yes	JSON	The export is possible via the RabbitMQ API, and the data is output in JSON format.
User accounts & permissions	Information about users and their permissions	Yes	JSON	Data can be configured and exported via API. <a href="#">Further information in our docs.</a>
System metrics (instances / resources in use)	Performance data of the instance / resource in use (e.g., CPU usage, memory usage)	Yes	Graphite Metrics	It must first be configured by the customer. A Graphite endpoint must be provided by the customer. <a href="#">Further information in our docs.</a>
	Sizes and capacities <i>Capacities of the available resources / instances</i>	Yes	JSON	Data can be configured and exported via API. <a href="#">Further information in our docs.</a>
System properties (instances / resources in use)	Versions and information necessary to check compatibility	Yes	JSON	Data can be configured and exported via API. <a href="#">Further information in our docs.</a>
Product / service-related data (product properties)	Configuration data and source code <i>Configuration of IT-systems / rudimental IT, settings, customizing, IP's, VLAN, interfaces, software code, scripts</i>	No. Company confidential STACKIT.	-	-
	Log data (non personalized and personalized)	Yes	Syslog	It must first be configured by the customer. A Syslog endpoint must be provided by the customer.

<i>System-status, technical-events, etc.</i>				<a href="#">Further information in our docs.</a>
Log data (non personalized and personalized)	Yes	JSON	STACKIT Audit Log.	<a href="#">Further information in our docs.</a>
<i>Login/logout of user, user activities</i>				