

STACKIT GmbH & Co. KG
Stiftsbergstraße 1 | 74172 Neckarsulm

Service Certificate – STACKIT Intake

Service name

STACKIT Intake

High level service description

The STACKIT Intake service provides a fully managed platform for ingesting data streams in the STACKIT data platform. The STACKIT Intake service thus enables the use of data for downstream processes such as preparation, aggregation, and analysis. Customers can instantiate intake engines with different maximum throughput capacities (“intake runners”). Within these intake runners, customers can then define intake pipelines (or simply “intakes”) that can store streams of any JSON messages in Apache Iceberg tables from the Dremio Iceberg REST catalog for the data platform, in line with their reserved throughput capacity. The data is ingested via the Apache Kafka protocol.

Key Features

- **Fully managed service:** STACKIT takes on full responsibility for providing, operating, and maintaining the intake infrastructure.
- **Reliable delivery and buffering:** 24-hour buffering ensures the intake is resilient to temporary downtimes in downstream systems (such as Dremio maintenance work).
- **Idempotency:** prevents duplicated data storage in the destination table, even in the event of repetition or errors.
- **Protocol compatibility:** Apache Kafka protocol for easy integration of existing data producers.
- **Lakehouse-ready:** direct, seamless storage of data streams in Apache Iceberg tables within the STACKIT data platform, with only a minimal delay.
- **Flexible schema management:** intake of any JSON messages. The STACKIT Intake service automatically derives a suitable table schema and supports additive schema evolution.
- **Error handling:** provision of a dead letter queue (DLQ) for messages that cannot be delivered due to format or schema problems.
- **Easy to manage:** purchase and configuration of intake runners and intakes through the STACKIT cloud portal, the STACKIT CLI, and Terraform.
- **Changes to capacity:** customers can increase the throughput capacity of the intake runners.
- **Billing model:** simple billing through STACKIT based on reserved maximum throughput capacity.

STACKIT GmbH & Co. KG | Stiftsbergstraße 1 | 74172 Neckarsulm

Kommanditgesellschaft | Sitz: Neckarsulm | Registergericht: Stuttgart, HRA 741347 | USt-IdNr. DE368560082
Commerzbank AG | DE55 6004 0071 0521 9928 00 | SWIFT/BIC COBADEFFXXX

Vertretungsberechtigte Komplementärin: STACKIT Beteiligungs-GmbH | Sitz: Neckarsulm
Registergericht: Stuttgart, HRB 795936 | Geschäftsführer: Christian Müller, Robin Hermann, Bernd Wagner

Service Plans

The service does not scale up or down automatically; instead, the customer sets the throughput capacity while instantiating the intake runner. Customers choose the two parameters that together define the maximum hourly throughput capacity of the runner across all the included intakes:

- The maximum number of messages per hour
- The maximum message size in KiB

These parameters determine both the upper limit of the hourly throughput (in GiB/h) and the size of the internal 24-hour message buffer.

Metric

STACKIT bills to the precise hour based on the reserved maximum hourly throughput capacity, rounded up to the next full GiB/h.

SLA specifics

For the availability of intake runners, the provisions in the latest version of the general STACKIT Cloud service description apply as a basic principle. This service description is available at: <https://stackit.com/en/gtc/service-description>.

- An intake runner is deemed to be **unavailable** for one minute
 - (i) if its Kafka protocol endpoint at the service delivery point cannot be accessed over the **entire** service minute
 - or
 - (ii) if **all** queries in the service minute return only Kafka protocol error codes 2, 8, 9, 15, 19, 20, 56, or 72.
- Otherwise, the intake runner is deemed to be **available** for the service minute.
- In particular, an intake runner is deemed to be **available** for a service minute if no query is made through the Kafka protocol endpoint in this service minute.

Backup

For the purpose of data recovery, intake runners are automatically backed up according to the general STACKIT Cloud service description (in its latest version, which is available at: <https://stackit.com/en/gtc/service-description>) based on a schedule specified by STACKIT. This schedule cannot be changed by the customer.

The backups cover only the configuration data of the intake runners, together with the configured intakes and intake users. JSON messages sent to the intakes are not backed up, because STACKIT Intake only buffers them and does not permanently store them.

Customers can implement their own backup concepts with the aid of the STACKIT Intake API, the STACKIT SDK, and the STACKIT CLI.

Additional terms

- Customers are solely responsible for managing the passwords for the technical users used to submit the JSON messages (intake users). Customers cannot read them out later via API.
- STACKIT Intake enforces a minimum password length and complexity for intake users. Further-reaching policies (such as password rotation) can be implemented by the customer under their own responsibility.
- Customers are responsible for keeping the Dremio personal access tokens (PATs) in the intake configurations up to date. Expiring PATs are not automatically extended by STACKIT Intake.

Annex: Exportability (Online Register)

Data Type	Description	Exportable (Yes/No)	Format	Additional notes
Customer data (Messages)	<i>Records customers send to Intakes to ingest into the Dremio Iceberg catalog of the data platform</i>	No. Company confidential.	N/A	Records are buffered for at most 24h and then deleted.
Intake Users & Permissions	User / password credentials with permission to send messages to an Intake	No. Company confidential.	N/A	Intake users can be viewed / accessed from the portal and API. Passwords cannot be retrieved for security reasons
Intake Runners and Intakes	Configuration of Intake Runners and Intakes	No. Company confidential.	N/A	Intakes and Intake Runners can be viewed / accessed from the portal and API
System metrics (instances / resources in use)	Performance metrics of instances and used resources (e.g., CPU load, memory usage, message throughput)	No. Company confidential.	N/A	Currently, customers do not have access to such metrics; these are only available internally
Product / service-related data (product properties)	<i>Configuration data and source code</i> <i>Configuration of IT-Systems/rudimental IT, Settings, Customizing, IP's, VLAN, Interfaces, Software Code, Scripts</i>	No. Company confidential.	N/A	Only internally available

	<i>Log Data</i>	No. Company confidential	N/A	Intake customers currently do not have access to log data on Intake usage.
--	-----------------	--------------------------------	-----	--

Version and start of validity

Version 1.0, valid from 04.02.2026