

Stackable Data Platform – Service Definition

What is Stackable Data Platform

Stackable Data Platform is an open-source Big Data platform. Everything in Stackable is 100% free and open source; users are free to use it however they wish and at any scale without obligation to pay to use the software. All the source code for Stackable and the software components it uses are publicly available.

Stackable is made up of several popular open-source, data-oriented software packages designed to be deployed together to form a basis for data movement, processing, analysis, and presentation. These software packages include:

- Apache Airflow
- Apache Druid
- Apache Hadoop HDFS
- Apache HBase
- Apache Hive
- Apache Kafka
- Apache NiFi
- Apache Spark
- Apache Superset
- Trino
- Apache ZooKeeper

Stackable provides Kubernetes operators to control and manage these software packages. It is suitable for deployment in the cloud, on premise or to both as a hybrid deployment. Essentially, Stackable can be hosted anywhere, on premise or in any cloud.

Stackable Data Platform is built with scalability, security, and fault tolerance in mind. Services can be deployed in a highly available fashion to ensure continuity of service in the event of a fault. Security is also a prime concern for Stackable, and features including authentication, authorisation and encryption are available throughout the platform.

Professional services and support are available from Stackable for customers that require help with deployment, architecture, design, and operation of their data platform. Flexible support hours are offered, with office hours and 24x7 support plans available.