Red Hat | RED HAT PARTNER CONNECT

PRODUCT BRIEF

Starburst

FAST FACTS

Company:
Starburst

Contact:
info@starburstdata.com

Website:
www.starburstdata.com

Company Description

Starburst is the Presto company. Founded by the creators and leading developers of the open source Presto project, Starburst Presto Enterprise gives analysts the freedom to analyze diverse data sets wherever their location, without compromising on performance. From data virtualization across data silos on-premises to high-performance data lake analytics in the cloud, Starburst allows you to query anything, anywhere.

Product Profile

Presto is the fastest growing open source SQL query engine, used by small and Fortune 500 companies alike. It is suitable for querying large scale data sets data across modern cloud-based and on-premise data management environments. Starburst Presto Enterprise is a fully supported, production-tested and expanded distribution of the open source Presto SQL query engine. It includes additional connectors for commercial database systems, query optimization, management tools and other improvements. Presto users can take advantage of performance improvements of 10x and more. Users of other tools migrating to Presto enjoy even larger improvements in performance and reduction of cost for queries. The simplified management of Presto clusters, the query speed and the combination of data sources, make new, formerly impossible use cases now possible.

Red Hat TECHNOLOGY PARTNER

Red Hat CERTIFIED TECHNOLOGY

Overview

Starburst Presto Enterprise is the enterprise distribution of Presto, the high performance, distributed SQL engine for querying any data source.

Executive Summary

Running Starburst Presto Enterprise on Red Hat's OpenShift, allows users to run multiple Presto clusters, scale up and down dynamically, and optimize for query speed and cost as desired. Connecting to numerous data sources and querying massive data sets has never been easier or faster.

Using the Starburst Presto Enterprise Kubernetes Operator with OpenShift eases the burden and complexity of configuring, deploying, managing, and monitoring Presto clusters and the separate workers as containerized applications.

Statement from Partner

"Using Starburst Presto Enterprise on the production-ready, dynamic Kubernetes infra-structure provided by Red Hat's OpenShift Container Application Platform allows owners of Presto deployments to focus on the data sources they provide to their users.", says Matthew Fuller, Starburst VP of Engineering, and contin-ues, *"They can create Presto clusters with ease, and maintain and scale them as desired. This allows them to expose the tremendous power of querying many data sources of massive scale to their users. The data analysts can use their tools of choice and leverage their knowledge of SQL to provide new, valuable insights to their business customers."*

Statement from Red Hat Connect

"Developers and architects looking to build new applications in, and for the cloud, or migrate existing applications to a cloud-based infrastruc-ture, partner with Red Hat to develop and deliver more supportable solutions sooner," said Mike Werner, Sr. Director, Global Technology Part-ners, Red Hat. *"Red Hat certification assures a supportable platform for all types of customer deployment models. Red Hat is thrilled to work with software partners like Starburst, resulting in the world's largest open, and commercially supportable application ecosystem.*

Starburst

Product Benefits

Running Starburst Presto Enterprise on Red Hat's OpenShift Container Application Platform enables you to provide the benefits of Presto to your users with a minimal management overhead for your infrastructure:

- Deploy and configure Presto on a Kubernetes cluster with ease, using the Starburst Presto Enterprise Operator and Containers on Red Hat's OpenShift Container Application Platform.
- Achieve high availability of the Presto clusters for your mission critical queries, reports and analyses.
- Leverage autoscaling for efficient use of Presto cluster resources, handling usage spikes, and deal with increased user demand quickly and easily.
- Perform fast, interactive queries on any data, wherever the data is stored. Users can query and combine data from popular data sources such as Ceph, S3, Azure Blob Storage, Google Cloud Storage, HDFS, MySQL, Microsoft SQL Server, PostgreSQL, Cassandra, MongoDB, Kafka, and Teradata, Snowflake, and many others.
- Keep your data protected with security protections in the whole pipeline, pass through credential support or even dedicated, private Presto clusters for specific users.
- Query numerous open source and commercial data sources at massive scale, right there in the cloud or other deployment, enabling new use cases and analysis patterns.

Use Cases

- Efficiently query data natively from your Data Lake in distributed storage such as Ceph, Cloud Storage, or HDFS.
- Efficiently query data natively from relational data sources such as Teradata, MySQL, Oracle, PostgreSQL, and SQL Server.
- Access non-relational data sources such as Apache Kafka, Elasticsearch, MongoDB, and Apache Cassandra with the simplicity and power of SQL
- Execute federated queries across the data sources without having to move or copy data to a single system.
- Allow analysts to use their favorite business intelligence tools such as Tableau, Power BI, Qlik, and Apache Superset and query data from distributed storage, relational, and non-relational data sources.

Starburst Enterprise Presto is a Red Hat certified application and is available for customer download from the Red Hat certified container registry.

https://access.redhat.com/containers/#/registry.connect.redhat.com/starburst/presto-operator