

Starburst 323e Release

Webinar - 11/19/2019



Starburst
ENTERPRISE PRESTO

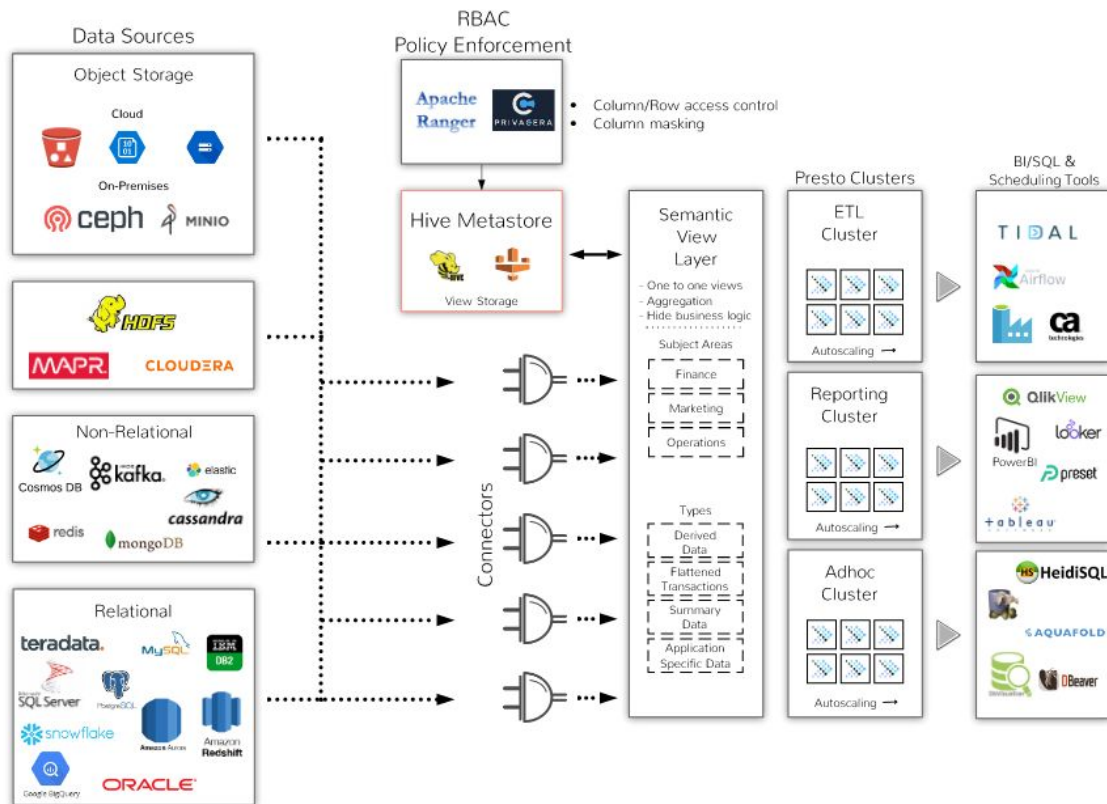
Starburst 323e Release

- Parallel Snowflake connector
- IBM DB2 connector
- MapR connector
- RedHat Certification for Openshift
- LDAP support for built-in system access control
- Starburst Secrets - Encryption for security sensitive configurations, e.g. passwords in catalogs
- Hadoop 3 support, including Hive ORC ACID support
- On-premises Coordinator HA
- Direct Parallel Teradata connector
- Power BI DirectQuery connector



Snowflake Connector - Why?

Federated Enterprise Semantic Layer

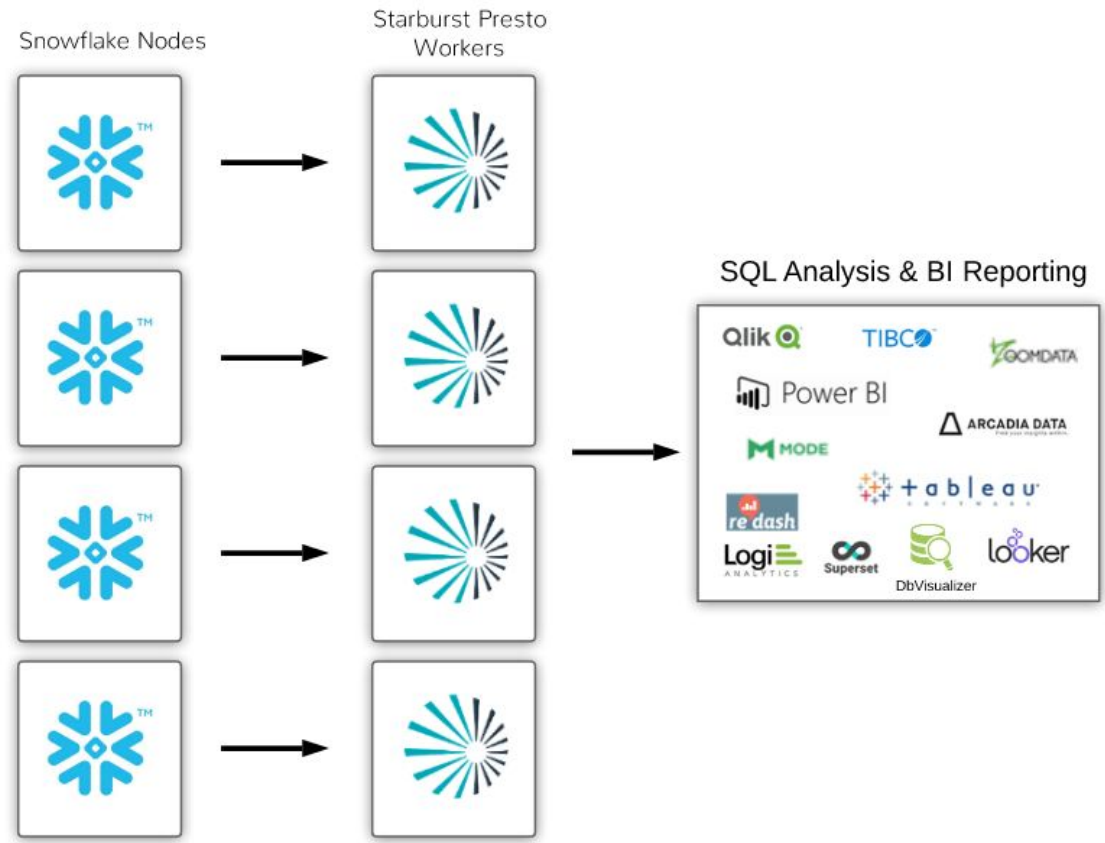


Snowflake Serial & Parallel Connectors

- Connect directly to Snowflake on any cloud
- Two options
 - JDBC - Great for when result sets are small
 - Distributed - Preferred when result sets are large
- Supports 1 Starburst Presto catalog for multiple Snowflake schemas
- Supports user impersonation
- Cost Based Optimizer uses statistics on Snowflake tables



Snowflake Distributed Option



Snowflake Connector Demo



DB2 Connector

- Supports DB2 11.5
- Any flavor of DB2 that provides a JDBC driver
- Add JDBC driver to Presto plugins directory
- DB2-to-Presto type mapping
- User impersonation
- Cost Based Optimizer uses statistics on DB2 tables



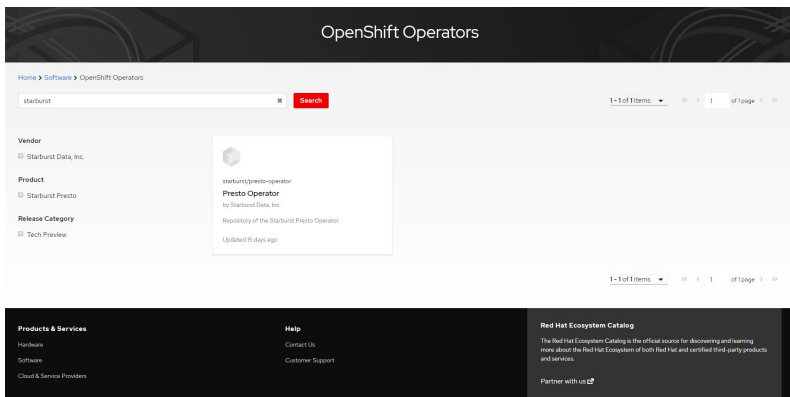
DB2 Connector



- Secure and non-secure clusters supported
- MapR client installed on each Starburst Presto node
- DB2-to-Presto type mapping
- User impersonation
- Cost Based Optimizer uses statistics on DB2 tables

Starburst Kubernetes (K8S) RedHat Certified

- Presto K8S operator fully certified on OpenShift
- Deploy Starburst K8S anywhere
- Query data in CEPH, HDFS or Cloud Storage
- Available on the RedHat Ecosystem Catalog



LDAP support for built-in system access control

- By default, all users have access to all catalogs. You can override this behavior by adding rules in Starburst's System Access Control [Plugin](#)
- The System Access Control will match an authenticated user to the access control rules for catalogs.
- The System Access Control can match groups to access control rules for catalogs. In this scenario, the authenticated user must be assigned to a group in LDAP

```
"catalogs": [  
  {  
    "user": "admin",  
    "catalog": "(mysql|system)",  
    "allow": "all"  
  },  
  {  
    "groups": ["finance", "admin"],  
    "catalog": "postgres",  
    "allow": true  
  },  
  {  
    "catalog": "hive",  
    "allow": "all"  
  },  
  {  
    "user": "alice",  
    "catalog": "postgresql",  
    "allow": "read-only"  
  },  
  {  
    "catalog": "system",  
    "allow": "none"  
  }  
]
```

Starburst Secrets



- Encrypting sensitive information in Starburst Presto configuration files
- Handled by a industry standard Java keystore
- Any configuration value can be encrypted. e.g. catalogs, configurations
- Eases concerns of having sensitive information in plain text on Presto clusters
- Keystore can be maintained outside of Presto configurations

Starburst Secrets

Demo



- Create a keystore for the Oracle catalog password
- Enable Starburst Presto to read secrets
- Replace password with secret variable in oracle.properties
- Query as normal

Starburst Secrets

1. `keytool -genseckey -alias oracle_connection_password -keyalg PBE -keystore presto-keystore.pfx -storetype PKCS12`

2. create configuration-source.properties in /etc/presto

```
[root@ip-172-31-8-140 presto]# cat configuration-source.properties
```

```
keystore.enabled=true
```

```
keystore.file-path=etc/presto-keystore.pfx
```

```
keystore.password=changeit
```

3. Verify the alias:

```
[root@ip-172-31-8-140 presto]# keytool -list -v -keystore presto-keystore.pfx
```

```
Enter keystore password:
```

```
Keystore type: PKCS12
```

```
Keystore provider: SUN
```

```
Your keystore contains 1 entry
```

```
Alias name: oracle_connection_password
```

```
Creation date: Nov 19, 2019
```

```
Entry type: SecretKeyEntry
```

4. Modify the oracle connector password:

```
[root@ip-172-31-8-140 catalog]# cat oracle.properties
```

```
connector.name=oracle
```

```
connection-url=jdbc:oracle:thin:@74.115.253.196:1521/orclpdb
```

```
connection-user=presto
```

```
connection-password=${ENV:oracle_connection_password}
```

```
oracle.impersonation.enabled=true
```

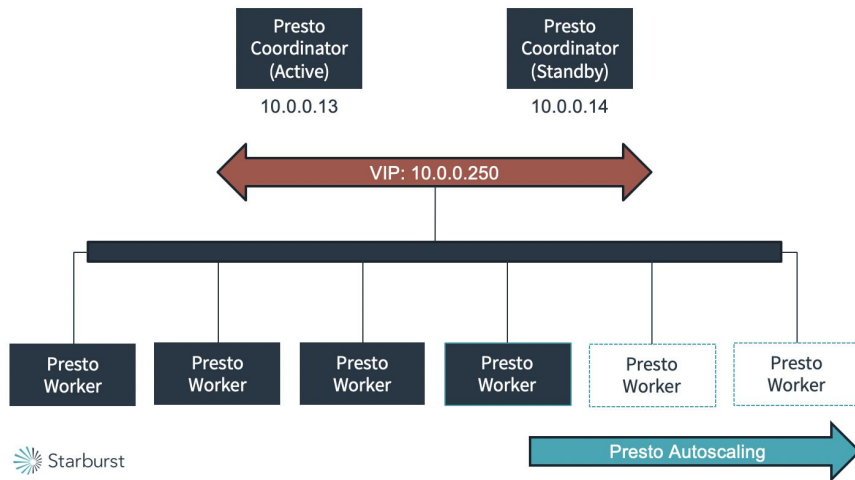
Starburst Hadoop 3.1 Support

- Starburst Presto is certified to work with Hadoop 3.1
- Including:
 - Read support for Hive transactional tables (ORC ACID)
 - Read and write for HDFS locations utilizing Erasure Coding
 - Support for Hive materialized views
 - Compatibility with the newer Apache Ranger version on HDP 3.1
 - Compatibility with the newer Hive Metastore version in 3.1
 - Compatibility with new Hive bucketing support in 3.1

Starburst On-Premises Coordinator High Availability

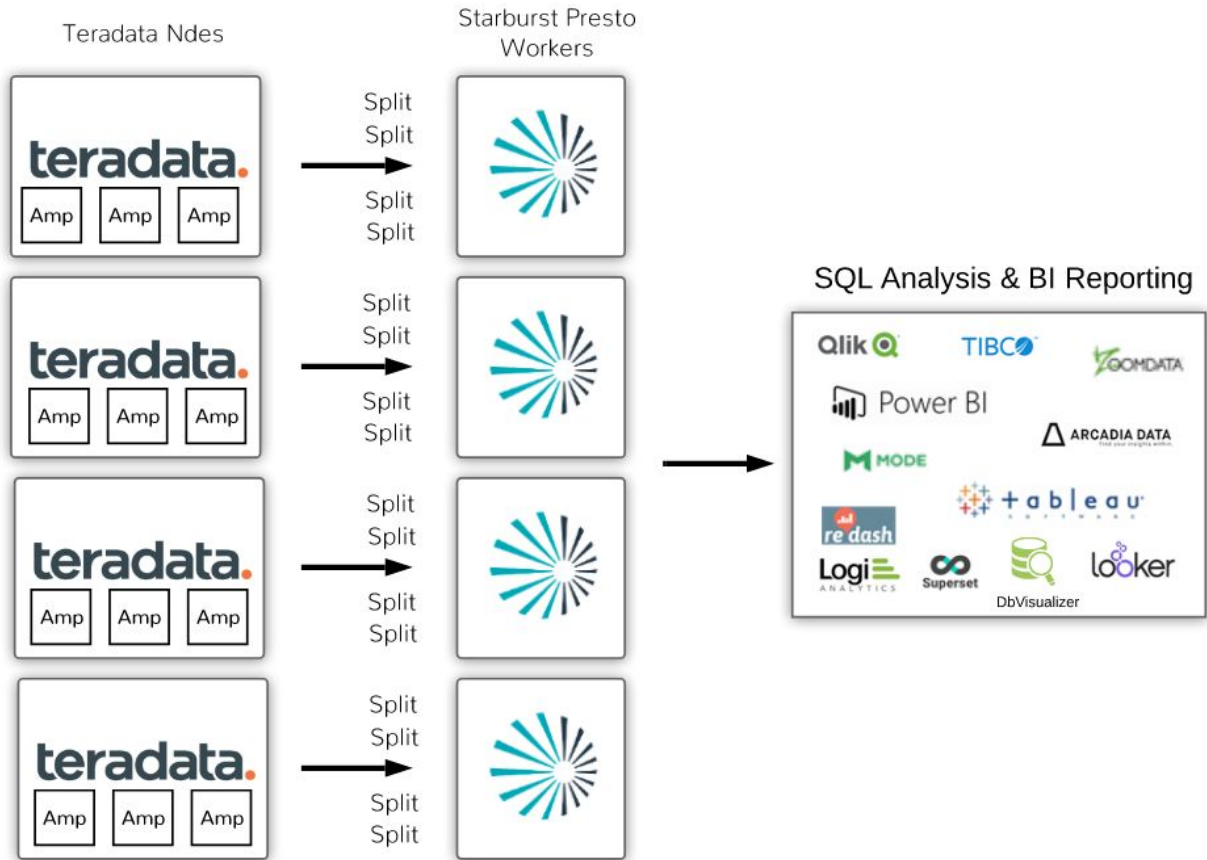


- Available for on-premises deployments
- Ensures 24/7 uptime of critical Starburst Presto deployments
- Uses a “keepalived” process that constantly monitors health
- Included in Starburst Enterprise Edition

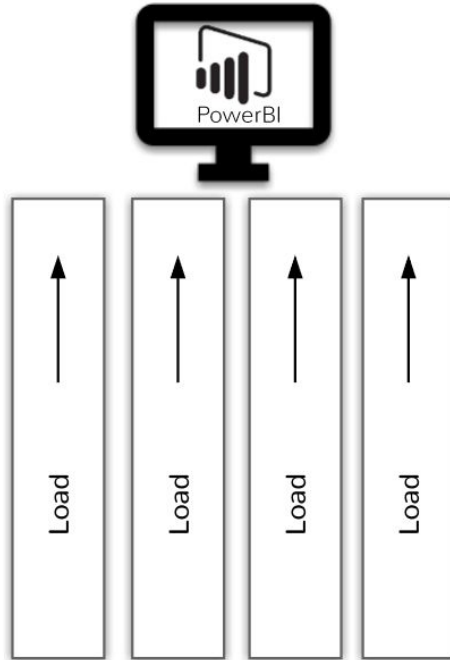


- Connects at the Worker to AMP level in Teradata
- Teradata Table Operator UDF and custom transmitter components
- Direct access to storage on Teradata via UDF
- Presto controller and many receivers for parallel run of many queries
- High performance for SELECT queries
- Already in successful deployment at large media customer

Teradata Parallel Direct Connector



Power BI - Import Function

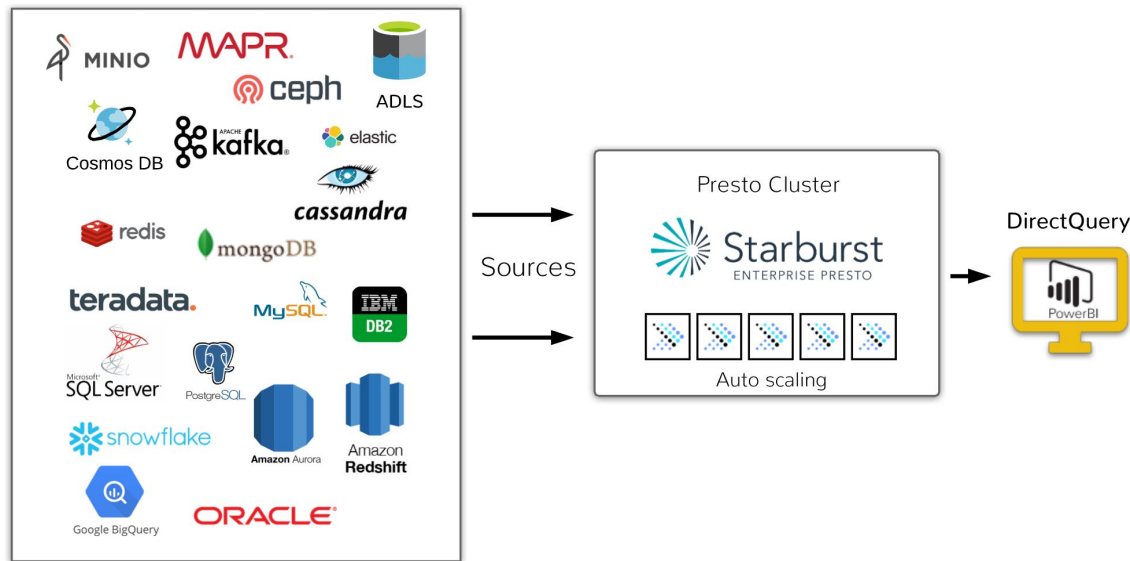


- Data is loaded into the Power BI desktop
- Connectors are limited
- Not recommended for data > 1GB
- Limited to the resources at the desktop



Limited Sources

Power BI - DirectQuery Connector



- Starburst Presto handles data federation and querying
- No data limits using the power of Starburst Presto
- Develop reports & dashboards from data living anywhere