



Starburst Global Security

Access All of your Data - Anywhere...

- The Good: Presto allows your users to access your data from virtually any source
- The Bad: Controlling data access
- The Solution: Implement Starburst Global Security



Agenda

1

Global Security Introduction

2

System Level Security in
Starburst Presto

3

Policy Creation and Application

4

Access Policies

5

Row-Level & Masking

6

Secure Zones

7

LDAP UserSync

8

Example Company

Starburst Global Security Introduction



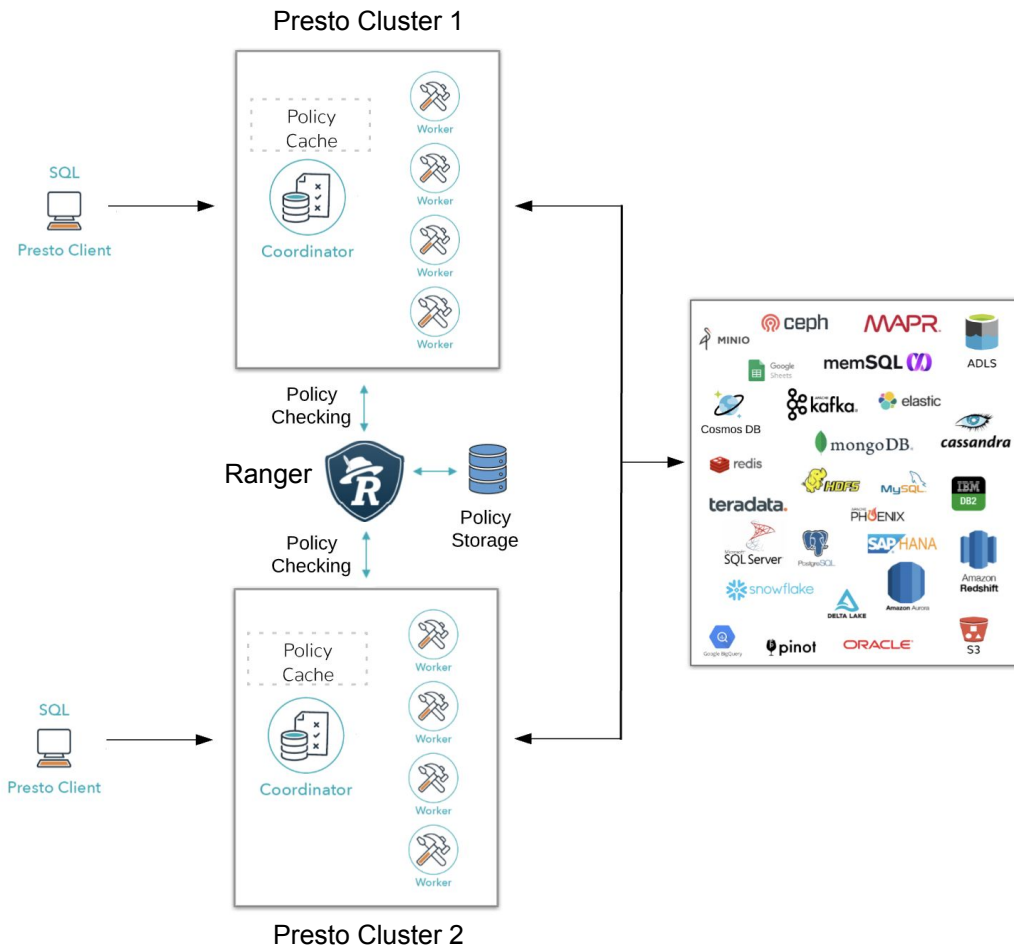
Global Security in Starburst Presto

- Flexible fine-grained access control policies across all connectors
- Column masking for any connectors
- Row-level filtering covering all connectors
- User/Group LDAP/AD synchronization
- Security Zones allow policy delegation to users and groups



Global Security Architecture

- One or more Starburst Presto clusters “point” to a Ranger instance
- For each query issued against the Presto cluster, Ranger is queried to match against existing policies
- Based on a decision tree of policies, policies are applied to the query
- Policies can be cached in Presto to improve query performance

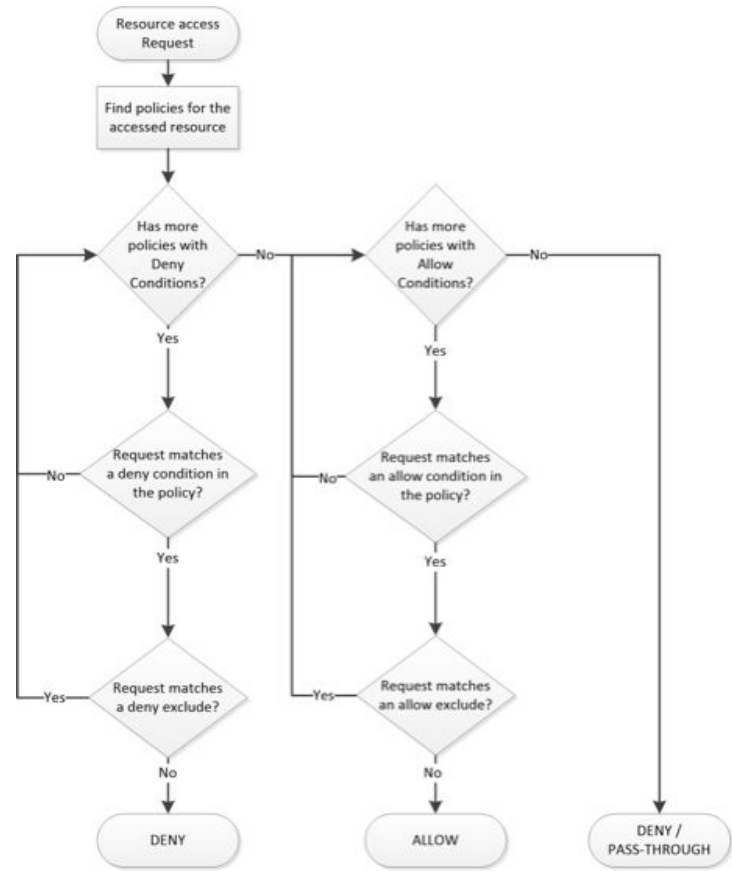


System Level Security in Starburst Presto



Policy Evaluation Flow

- Once a query is issued to Presto, the decision tree on the right is followed within Ranger
- All Deny policies are matched and enforced first
- After all deny policies have been met, allow policies are evaluated.




Apache Ranger Policy Evaluation Flow

System Level Security in Starburst Presto

- Policies applied on all connectors
- Object types are: Catalog, procedures, session properties and query
- Ability to add a validity period to each policy
- User/Group UserSync from LDAP or Active Directory

LDAP/AD
User/Group
Synchronization



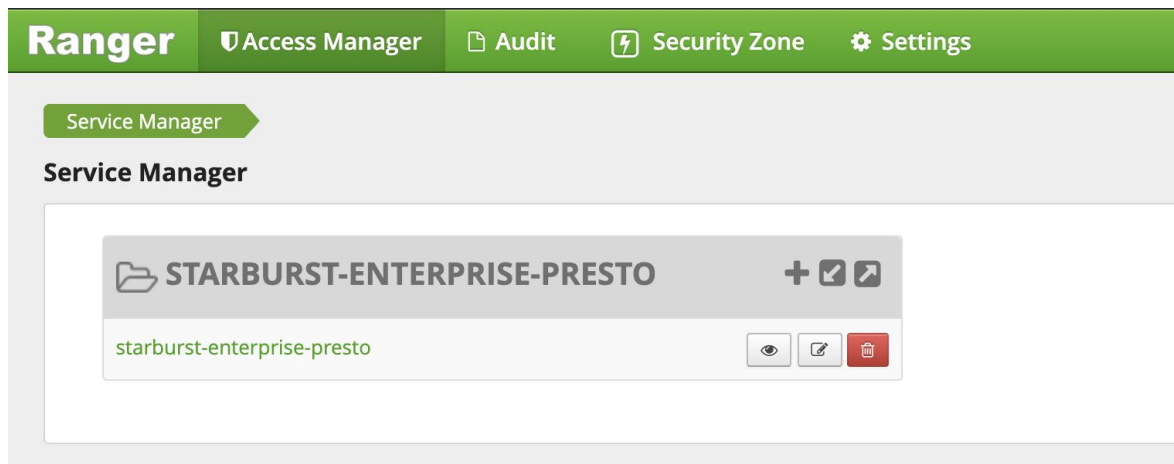
Create Policy

Policy Details :

Policy Type	Access	
Policy Name *	FinancePolicy ⓘ	enabled
Policy Label	Policy Label	
<div>catalog ▼ *</div>	<div>✕ hive ✕ sqlserver</div>	include
<div>schema ▼ *</div>	<div>✕ adlsgen2 ✕ dbo</div>	include
<div>table ▼ *</div>	<div>✕ *</div>	include
Column *	<div>✕ *</div>	include
Description	Allow Finance group to query Data Lake and Finance database.	

Ranger Interface Overview

1. Access Manager -
 - a. Resource Based Policies - Create policies for access to catalogs, schemas, procedures, session parameters, tables, column, masking and row-level filtering
 - b. Tag Based Policies - Ability to create policies based on Tags
 - c. Reports - Search and report on policies
2. Audit - Combined with an independent Solr or HDFS storage will log all access events
3. Security Zone - Delegate catalog policy control to certain groups/users
4. Settings - User and group management



Policy Creation and Application



Policy Creation

Three different types of policies

1. Access Control
2. Column Masking
3. Row Level Filtering

Access Control	Column Masking	Row Level Filtering
<ol style="list-style-type: none">1. Control access to catalogs, schemas, tables, procedures, query and session properties.2. All objects are "deny all" at first installation. Must add policies to allow access.3. Flexible "deny" then "allow" or "allow" then "deny all others"	<ol style="list-style-type: none">1. Columns from any table from any connector can be masked2. Only one column can be masked per policy3. Masking types: mask, partial mask, nullify and hash	<ol style="list-style-type: none">1. Applied at any connector row, based on condition2. Conditions are set at the user and group levels3. Filtering is done using a "where clause" Example: nation <> 9

Policy Creation - Access Policies

Policy Details :

Policy Type **Access**

Policy ID **8**

Policy Name *

Policy Label

catalog *

schema *

table *

Column *

Description

- Multiple entries or wildcards can be used for each section
- Different permission levels can be set
- Deny policies get evaluated first

add/edit permissions

☒ Select

☒ Insert

☐ Delete

☐ Update

☒ Ownership

☐ Select/Deselect All

Example: Marketing doesn't have access to the postgresql catalog, because it has not been explicitly granted

```
prestosql> select * from  
postgresql.public.customer limit 5;  
Query 20200729_191304_00249_8tbeg failed:  
Access Denied: Cannot access catalog
```

Allow Conditions :

Select Role	Select Group	Select User	Permissions
<input type="text" value="Select Roles"/>	<input type="text" value="Marketing"/>	<input type="text" value="Select Users"/>	<input checked="" type="button" value="Select"/> <input checked="" type="button" value="Insert"/> <input checked="" type="button" value="Ownership"/> <input type="button" value="✎"/>

Policy Creation - Column Masking

Policy Details :

Policy Type	Masking
Policy ID	16
Policy Name *	<input type="text" value="MarketingInterns"/>
Policy Label	<input type="text" value="Policy Label"/>
Catalog *	<input type="text" value="hive"/>
Schema *	<input type="text" value="adlsgen2"/>
Table *	<input type="text" value="customer"/>
Column *	<input type="text" value="name"/>
Description	<input type="text" value="Mask customer name for Marketing Interns"/>

- Currently supported masking types are Mask, Partial mask, Hash and Nullify →
- Only one policy can be attached to a column at a time
- Can be applied to users and/or groups

Select Masking Option

- ☒ Mask
- ☐ Partial mask: show last 4
- ☐ Partial mask: show first 4
- ☐ Hash
- ☐ Nullify

Example




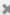
```
presto> select custkey,name from hive.adlsgen2.customer  
limit 2;
```

custkey	name
75001	xxxxxxxxxxxxxxxxxxxxxx
75002	xxxxxxxxxxxxxxxxxxxxxx

Select Group	Select User	Access Types	Select Masking Option
<input type="text" value="Marketing-Interns"/>	<input type="text" value="Select Users"/>	<input type="button" value="Select"/> <input type="button" value="Edit"/>	<input type="button" value="Mask"/> <input type="button" value="Edit"/>


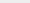
Policy Creation - Row Level Filtering

Policy Details :

Policy Type	Row Level Filter
Policy ID	10
Policy Name *	RowFilterMarketing 
Policy Label	Policy Label
Catalog *	 hive
Schema *	 adlsgen2
Table *	 nation
Description	Filter out NationKey = 24 for Marketing group.

- Each row-filter policy can only refer to one column
- Filter similar to where clause in a sql statement
- Can be applied at the user or group level

Row Filter Conditions :




Select Role	Select Group	Select User	Access Types	Row Level Filter
Select Roles	✖ Marketing	Select Users	Select 	nationkey=24 

Policy Creation - Policy Validity Period

- Provides a date/time-based period to enforce policies
- Available on access, masking and row-level filtering policies

 Add Validity Period

Policy Validity Period ×

Start Time	End Time	Time zone
2020/07/20 06:00:51 × 	2020/08/19 06:00:51 × 	US/Eastern (EDT) ×  ×

+

Cancel

Save

Reports



Reports

You can use the Reports page to help manage policies more efficiently as the number of policies increases

- Search policies by a variety of fields including groups or users
- Export results to Excel, csv or json
- Expand on different conditions:

User Access Report

Policy Name

Enter Policy Name

Policy Type

Access

Component

Select Component

Resource

Enter Resource Name

Policy Label

Select Policy Label

Zone Name

Select Zone Name

Search By

Group

Marketing

Q Search

Export

TAG

hide

Policy ID	Policy Name	Policy Labels	Resources	Policy Type	Status	Zone Name	Allow Conditions	Allow Exclude	Deny Conditions	Deny Exclude
No Policies found!										

STARBURST-ENTERPRISE-PRESTO

hide

Policy ID	Policy Name	Policy Labels	Resources	Policy Type	Status	Zone Name	Allow Conditions	Allow Exclude	Deny Conditions	Deny Exclude
8	MarketingPolicyCatalogs	--	schema:* catalog:sqlserver,mysql,sno... column:* table:*	Access	Enabled	--	+	+	+	+
17	MarketingOrdersClerk	--	schema:adisgen2 catalog:hive column:clerk table:orders	Access	Enabled	--	+	+	+	+

Allow Conditions	Allow Exclude	Deny Conditions	Deny Exclude
+	+	+	+

Allow Conditions	Groups	Users	Accesses
Marketing	--	--	select insert ownership

Security Zones



Security Zones

Security Zones allow the creation and administration of policies to groups and/or users.

 **MarketingResources** 

This zone will delegate the Marketing group to manage policies on the SQL Server Marketing database

Zone Administration

Admin Users --

Admin Usergroups **Marketing**

Auditor Users **admin**

Auditor Usergroups --

Zone Tag Services

No tag based services are associated with this zone

Services

Service Name	Service Type	Resource
starburst-enterprise-presto	STARBURST-ENTERPRISE-PRESTO	catalog : sqlserver

- Multiple users/groups can be added to a zone
- Filter similar to where clause in a sql statement
- Can be applied at the user or group level

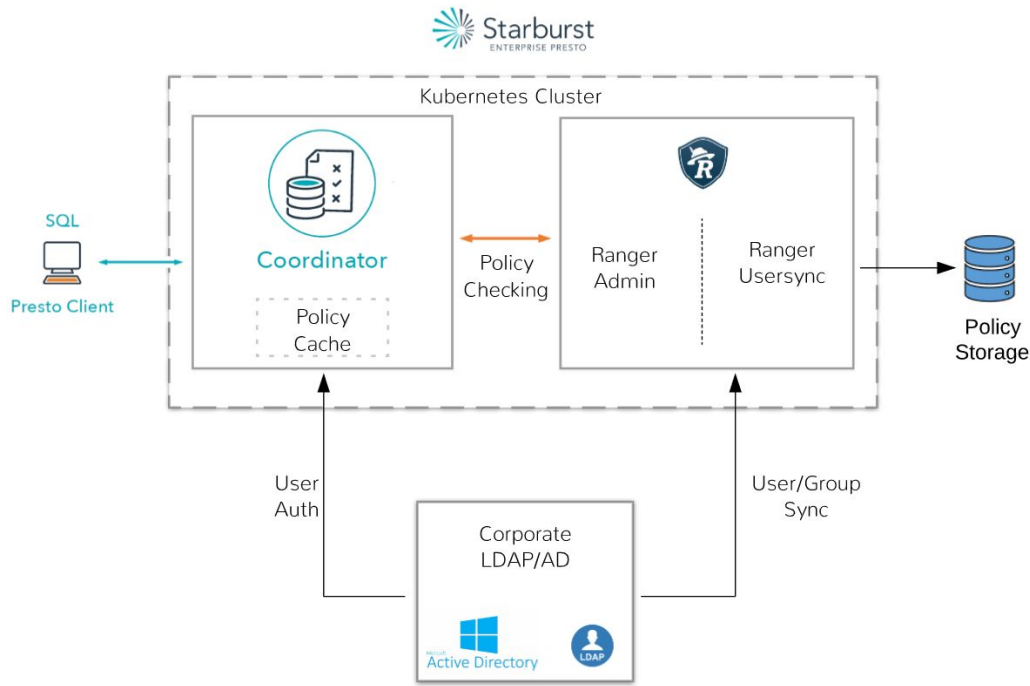
LDAP Usersync



LDAP/Active Directory Usersync

Ranger users and groups can be synchronized with a corporate LDAP or Active Directory (AD) system to reduce administration time in Ranger


- Users that are authenticated using Presto are also applied within Ranger
- Usersync runs as a container within Starburst Presto k8s service
- Users/Groups are synchronized every hour



LDAP/Active Directory Usersync - Syncing Tips

Synchronizing LDAP/AD objects can be a challenge. Determine what sets of users and groups they belong to before getting started.

- [ldapsearch](#) is a handy tool to query LDAP/AD
- Usersync runs as its own container in k8s
- Logs can be found by viewing the container logs for ranger-usersync



Common Usersync Variables: (bind user/p required)

RANGER__SYNC_LDAP_URL: ldaps://xxxx:636

RANGER__SYNC_LDAP_BIND_DN:

RANGER__SYNC_LDAP_BIND_PASSWORD:

RANGER__SYNC_LDAP_SEARCH_BASE:

RANGER__SYNC_LDAP_USER_SEARCH_BASE:

RANGER__SYNC_GROUP_SEARCH_ENABLED:

RANGER__SYNC_GROUP_USER_MAP_SYNC_ENABLED:

RANGER__SYNC_LDAP_USER_NAME_ATTRIBUTE:

Starburst Ranger Usersync Documentation:

<https://docs.starburstdata.com/latest/kubernetes/ranger.html#ldap-user-synchronization>

Example Company



Example Company

Company Name: Example Corp

Departments: Marketing, Finance & Operations

Sub Departments: Marketing Interns

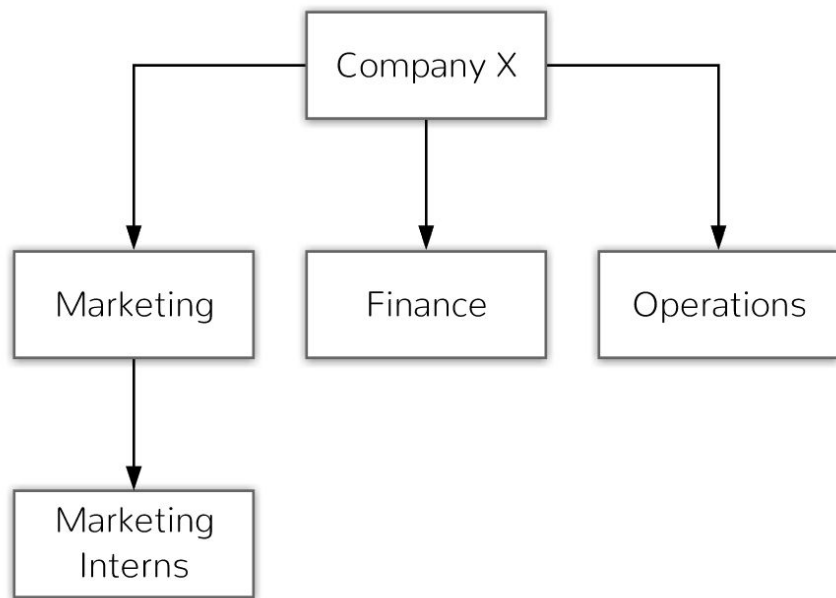
Users:

ted-marketing

vishal-operations

nancy-finance

kevin-marketing-intern



Data Sources

ADLS - Corporate Data Lake

SQL Server - Marketing Database

PostgreSQL - POS System Database



ADLS
Data Lake



Permissions Overview

Marketing



Read only
Write to Sandbox



Full Access

Finance



Read only
Write to Sandbox



Full Access



Full Access

Operations



Read only
Write to Sandbox



Full Access

Policy Layout

Group	Data Source	Policy Name	Description
Marketing	SQL Server	Marketing-MarketingDB	Full Access to SQL Server Marketing database
All	Data Lake	All-DataLake	Read-only access to the Data Lake for each department
Marketing	Data Lake	Marketing-Sandbox	Full access to the Marketing Sandbox - Read-only for interns
Finance	All sources	Finance-All	Read-only access to all sources
Finance	Data Lake	Finance-Sandbox	Full access to the Finance Sandbox
Operations	POS	Operations-POS	Full Access to POS PostgreSQL database
Operations	Data Lake	Operations-Sandbox	Full access to the Operations Sandbox
Marketing	Data Lake	MarketingInterns-CustMask	Mask customer name for Marketing Interns
Public	Data Lake	NationRowFilter	Restrict NationKey = 19 from anyone

Marketing Database - Full Access

Policy Name *

Marketing-MarketingDB

Policy Label

Policy Label

catalog ▼ *

✕ marketing

schema ▼ *

✕ *

table ▼ *

✕ *

Column *

✕ *

Description

Full Access to SQL Server Marketing database

▼

 Demo-Marketing

▶

 hive

▶

 marketing

▶

 system

Select Group	Select User	Permissions
<div>✕ Marketing</div>	<div>Select Users</div>	<div>SelectInsertDeleteUpdateOwnership</div> <div></div>

Data Lake - Read Only Access

Policy Name *

All-DataLake

Policy Label

Policy Label

catalog ▾ *

× hive

schema ▾ *

× adlsgen2

table ▾ *

× *

Column *

× *

Description

Read only access to the Data Lake

▼ Demo-Finance

▼ hive

> adlsgen2

> financesb

> marketing

> pos

> system

▼ Demo-Marketing -

▼ hive

> adlsgen2

> marketingsb

> marketing

> system

▼ Demo-Operations -

▼ hive

> adlsgen2

> operationssb

> pos

> system

Select Group	Select User	Permissions
<div>× Marketing</div> <div>× Operations</div> <div>× Finance</div>	<div>Select Users</div>	<div>Select</div> <div></div>

Marketing Sandbox

Policy Type **Access**

Policy ID **10**

Policy Name * Marketing-Sandbox ⓘ

Policy Label Policy Label

catalog * × hive

schema * × marketingsb

table * × *

Column * × *

Description Full access to the Marketing Sandbox but deny other groups. ↗






Allow full access to only the Marketing group
and read-only for interns

Select Group	Select User	Permissions
× Marketing	Select Users	Select Insert Delete Update Ownership ⓘ
× Marketing-Interns	Select Users	Select ⓘ

Deny other groups access to the sandbox

Deny All Other Accesses :

True

- ▼  hive
 - >  adlsgen2
 - >  marketingsb
- >  marketing
- >  system

Masking Customer Name for Marketing Interns

Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups
MarketingInterns-CustMask	--	Enabled	Enabled	--	Marketing-Interns

Operations sees customer name

```
select * from hive.adlsgen2.customer limit 10;
```

123 custkey	ABC name	ABC address
75,001	Customer#00007500	iQyegZCktrX8jMFs9ip
75,002	Customer#00007500	TRzWtXys54mXmbNLIZQ4UR,5VkzA4Ycjsx
75,003	Customer#00007500	OVaJQHekQKFzsjqYpkLD
75,004	Customer#00007500	DM8SVyMtxAqxUhSFtMoYXSwAPri,rCxLHaC
75,005	Customer#00007500	UFaglit4yLXrIK4,KpBFwt7Pa5egjXuWw8 P6u
75,006	Customer#00007500	TCS8unq57G
75,007	Customer#00007500	iPQBVCkNZlwa,g3VUeFzV
75,008	Customer#00007500	imzHhGJHRabMisy43J4P,RAo3K
75,009	Customer#00007500	bWv8PnlGSz
75,010	Customer#00007501	VYzjCpkja24Kbnvv0Oa1,b2pVU1pd45AYCoZK

Marketing interns do not

```
select * from hive.adlsgen2.customer limit 10;
```

123 custkey	ABC name	ABC address
75,001	xxxxxxxxxxxxxxxxxx	iQyegZCktrX8jMFs9ip
75,008	xxxxxxxxxxxxxxxxxx	imzHhGJHRabMisy43J4P,RAo3K
75,009	xxxxxxxxxxxxxxxxxx	bWv8PnlGSz
75,010	xxxxxxxxxxxxxxxxxx	VYzjCpkja24Kbnvv0Oa1,b2pVU1pd45AYCoZK
75,002	xxxxxxxxxxxxxxxxxx	TRzWtXys54mXmbNLIZQ4UR,5VkzA4Ycjsx
75,003	xxxxxxxxxxxxxxxxxx	OVaJQHekQKFzsjqYpkLD
75,004	xxxxxxxxxxxxxxxxxx	DM8SVyMtxAqxUhSFtMoYXSwAPri,rCxLHaC
75,005	xxxxxxxxxxxxxxxxxx	UFaglit4yLXrIK4,KpBFwt7Pa5egjXuWw8 P6u
75,006	xxxxxxxxxxxxxxxxxx	TCS8unq57G
75,007	xxxxxxxxxxxxxxxxxx	iPQBVCkNZlwa,g3VUeFzV

Row Filter Example

Policy Name * NationRowFilter18 ⓘ

Policy Label Policy Label

Catalog * ✕ hive

Schema * ✕ adlsgen2

Table * ✕ nation

Description Filter out nationkey = 18

Filter out nationkey = 18
to all users/groups

```
select * from hive.adlsgen2.nation
```

123 nationkey	ABC name	123 reg
0	ALGERIA	
1	ARGENTINA	
2	BRAZIL	
3	CANADA	
4	EGYPT	
5	ETHIOPIA	
6	FRANCE	
7	GERMANY	
8	INDIA	
9	INDONESIA	
10	IRAN	
11	IRAQ	
12	JAPAN	
13	JORDAN	
14	KENYA	
15	MOROCCO	
16	MOZAMBIQUE	
17	PERU	
19	ROMANIA	

Select Group	Select User	Access Types	Row Level Filter
✕ public	Select Users	Select	nationkey<>18

Policies are applied to all SQL clients

This example shows
Power BI and
masking policy for
the customer name

Navigator

Display Options ▾

- presto-demo.az.starburstdata.net: 8080 [5]
- hive [1]
- adlsgen2 [8]
 - ☒ customer
 - ☐ lineitem
 - ☐ nation
 - ☐ orders
 - ☐ part
 - ☐ partsupp
 - ☐ region
 - ☐ supplier
- marketing
- pos
- system
- tpch

customer

custkey	name	address	nationke
75064	xxxxxxxxxxxxxxxx	KqHtS9rwWm	
75065	xxxxxxxxxxxxxxxx	xqbTJ, HvYyQotg3o6vldnQ1qb2YT	
75066	xxxxxxxxxxxxxxxx	wTcyGeJ6AyTug8Z, XpwFj	
75067	xxxxxxxxxxxxxxxx	AdkEMOeSVoVydZ9	
75068	xxxxxxxxxxxxxxxx	7Flh7Asxeq4GmMrPp3UY4vIHtPRsPHg606y	
75069	xxxxxxxxxxxxxxxx	ZgWUoiRxVjeNeK9vkWjMmV0,zols8	
75070	xxxxxxxxxxxxxxxx	zPlzOwRnS5XiAogAAcFCkytV5zrNJm	
75071	xxxxxxxxxxxxxxxx	xCToDGqgXu	
75072	xxxxxxxxxxxxxxxx	6CBUZqm a9mlRf	
75073	xxxxxxxxxxxxxxxx	D3XVljorMgp16d VfkDFxsst1wq0prqZ	
75074	xxxxxxxxxxxxxxxx	T4y3dkyFEqhyxOt1nuScl	
75075	xxxxxxxxxxxxxxxx	spMScTrbdaut1RaW0uUL, B3U3MGJyX	
75076	xxxxxxxxxxxxxxxx	soeCBHELM67EfVQlnP4wXDWkx1	
75077	xxxxxxxxxxxxxxxx	bjKa1wEfbGv oVlgKhan	
75078	xxxxxxxxxxxxxxxx	hI3JEftzFZHUZEBQtu5rhVC3WWTOAPgOBtQlq	
75079	xxxxxxxxxxxxxxxx	kyYRNfnYp, Dvu1VPM0Kpc0oQ89unNTpWEMN	
75080	xxxxxxxxxxxxxxxx	3zURKq4w2VDW OBpezCNvDjlLp1TlBjknf	
75081	xxxxxxxxxxxxxxxx	U, gb nLSaSp67yUUIyRmkY4na6B2nM	
75082	xxxxxxxxxxxxxxxx	bubleAQuvVYT5rx9pHChmGXVerRmkq	
75083	xxxxxxxxxxxxxxxx	kaQsgmAz3nPvkmxJlKwHofhi0lq	
75084	xxxxxxxxxxxxxxxx	Rnl39cQBxajMls, sO4pGOjvX5GW32g	
75085	xxxxxxxxxxxxxxxx	GD3tDB0GAB27D1K6Tbc	
75086	xxxxxxxxxxxxxxxx	E6usqXGwAFE	

Load Transform Data Cancel

Global Security in Starburst Presto

- Flexible fine-grained access control policies across all connectors
- Column masking for any connectors
- Row-level filtering covering all connectors
- User/Group LDAP/AD synchronization
- Security Zones allow policy delegation to users and groups



Thank You

www.starburstdata.com