



Release 3.8.0



Release 3.8.0

Copyright © 2020 Privacera, Inc. All rights reserved. Privacera, the Privacera logo, and Privacera Platform are registered trademarks of Privacera, Inc. in the United States and other countries. All other trademarks and service marks are the property of their respective owners.

Privacera

<http://www.privacera.com>

2020



Table of Contents

1. Privacera Release - 3.8.0.....	4
2. Supplementary Release Information.....	7
3. Supported Versions of Components.....	10



1. Privacera Release - 3.8.0

Release 3.8.0

08-Sep-2020

Versions

- Discovery Version - r.3.8.0.1
- Access Manager Version - rel.r.3.8.0.1.68
- Previous Discovery Version - r.3.7.0.1
- Previous Access Manager Version - rel.r.3.7.0.1.66

Features/Changelist

- EMR 5.30.x is now supported with Privacera platform.
- Access manager audit events now include the OrgId and WorkspaceId attributes for each access event in Databricks. These additional properties can help with supporting chargeback/billing models based on OrgId and/or WorkspaceId. For more supporting information see, [Databricks Chargeback Model](#).
- Databricks 7.x with Spark 3.0 is now supported for Databricks authorization and audits via Databricks plugin in Access Manager.
- Databricks 7.x and Databricks 6.x are supported using the same global init script.
- Process isolation feature for job clusters has been disabled in Databricks as it is not supported.

- A new service definition for authorizing files and folders via mount points has been added. This new service appears as FileSystem (privacera_files) in Access Manager under Resource Policies and allows READ / WRITE permissions for the files for DBFS, S3, ADLS, GCS when using mount points. The new service can be enabled by setting the following property to 'true' in Spark configuration in Databricks:
privacera.plugin.files.repository.enabled true
- [EMRFS] Security configurations in EMRFS have been improved to block STS endpoints for all other users except Hadoop, Hive, Presto, and Yarn. This prevents all other users (for example end-users and non-service users) from performing privilege escalation with an Assume Role Request via STS. For more supporting information see, [Block STS Endpoints](#) section.
- [EMRFS] Security Configurations for EMRFS within Privacera now support custom credential provider to Yarn, which is required by Yarn for handling large data read/write data jobs posted by Hive EMR. For more supporting information see, [Block STS Endpoints](#) section.
- Privacera Discovery now supports real-time scans on ADLS for multiple containers within a Storage account.
- [\[Preview\]](#) Athena Federated Query (Preview) is now supported within Privacera for authorization and audits in Access Manager.



© Privacera, Inc.. All rights reserved.

[Send feedback](#) on this topic to Privacera, Inc..



2. Supplementary Release Information

Introduction

This topic provides additional supporting information for changes and upgrades introduced in Privacera Release 3.8.0.

Databricks Chargeback Model

Docker

Bash (Unix Shell)

```
cd ~/privacera/downloads  
wget https://privacera.s3.amazonaws.com/public/managed-schema -O managed-schema
```

Bash (Unix Shell)

```
cd ~/privacera/docker  
mkdir -p solr_backup  
cp ~/privacera/docker/ranger/solr/conf/managed-schema solr_backup/  
cp ~/privacera/downloads/managed-schema  
~/privacera/docker/ranger/solr/conf/managed-schema
```

Bash (Unix Shell)

```
./privacera_services restart solr
```

To validate go to Solr UI -> Cloud -> Tree -> /configs -> ranger_audits -> managed-schema -> search for workspaceId.

Kubernetes

Bash (Unix Shell)



```
cd ~/path/to/kubernetes
mkdir downloads
cd downloads
wget https://privacera.s3.amazonaws.com/public/managed-schema -O managed-
schema
```

Bash (Unix Shell)

```
cd ~/path/to/kubernetes
mkdir backup
cp solr/conf/managed-schema backup/
cp downloads/managed-schema solr/conf/managed-schema
```

Bash (Unix Shell)

```
kubectl delete cm solr-config-0
kubectl create configmap solr-config-0 --from-file=solr/conf/
kubectl delete po solr-0 solr-1 solr-2
```

To validate go to solr UI -> Cloud -> Tree -> /configs -> ranger_audits -> managed-schema -> search for workspaceId.

Sample Ranger Audit showing WorkspaceId in the audits.

http://{privacera_portal}:8983/solr/ranger_audits/select?q=%3A*

Bash (Unix Shell)

```
{
  "id": "3618abc2-8f6b-48df-bda6-2d42c20bda5e-17",
  "access": "SELECT",
  "enforcer": "ranger-acl",
  "agent": "spark",
  "repo": "privacera_hive",
  "reqUser": "sushil.sani",
  "resource": "s_test/sales_data",
  "cliIP": "49.32.62.19",
  "logType": "RangerAudit",
  "result": 0,
  "policy": -1,
  "repoType": 3,
  "resType": "@table",
  "action": "select",
  "evtTime": "2020-09-10T14:05:36.616Z",
```



```
"seq_num": 29,  
"event_count": 1,  
"event_dur_ms": 0,  
"cluster": "Sushil-Ranger-Plugin-7.1",  
"agentHost": "0910-123329-purls895-10-55-250-184",  
"workspace_id": "7229028481585261",  
"account_name": "nvirginia.cloud.databricks.com",  
"_ttl_": "+90DAYS",  
"_expire_at_": "2020-12-09T14:05:39.495Z",  
"_version_": 1677456037120049200  
}
```

Block STS Endpoints

Currently in EMRFS, any user who can login to EMR Cluster, can use the STS service to assume a higher privileged IAM Role and read S3 data. While using Privacera, we have added a configuration that uses IPTables to block STS Endpoints for all users on the cluster except the application superusers such as Hadoop, Hive, Presto, Yarn. This feature can help avoid the privilege escalation.

Service Definition for FileSystem

Service Definition for FileSystem (privacera_files) which allows READ / WRITE for the files for DBFS, S3, ADLS, GCS when using mount points.

In Databricks, you need to set the below property in Spark configs:

```
privacera.plugin.files.repository.enabled true
```

© Privacera, Inc.. All rights reserved.

[Send feedback](#) on this topic to Privacera, Inc..

3. Supported Versions of Components

Supported Versions of Components

The following table provides the list of supported versions of components.

Components	Versions Supported in Each Release		
	3.6.x	3.7.x	3.8.x
MariaDB	10.3	10.3	10.3
Solr	8.5.1	8.5.1	8.5.1
Zookeeper	3.4	3.4	3.4
PostgreSQL	10	10	10
Databricks Runtime Version	6.3	6.3 - 6.6	6.3-7.2
Databricks Spark Version	2.4.4	2.4.4 - 2.4.5	2.4.4 - 3.0.0
AKS	1.14.8	1.15.11	1.15.11
EKS	1.15	1.15	1.15
EMR	5.29	5.29	5.29 - 5.30.x
RDS (Aurora DB MySQL flavor)	5.6	5.7/5.6	5.7/5.6
Databricks Spark Version for Discovery	-	2.4.4 - 2.4.5	2.4.4 - 2.4.5
Discovery Docker Spark Version	-	2.4.3	2.4.3
Discovery Kubernetes Spark Version	-	2.4.5	2.4.5
Hive	-	2.3.6	2.3.6
Spark	2.4.4	2.4.4	2.4.4
PrestoDB	0.227	0.227	0.227 - 0.232
Zeppelin	0.8.2	0.8.2	0.8.2
Hue	4.4.0	4.4.0	4.4.0



Livy	0.6.0	0.6.0	0.6.0
Ranger	2	2	2.1/3.x
Starburst PrestoSQL	332-e	332-e	332-e

© Privacera, Inc.. All rights reserved.

[Send feedback](#) on this topic to Privacera, Inc..

- 1. Privacera Release - 3.8.0..... 4**
- 2. Supplementary Release Information..... 7**
- 3. Supported Versions of Components..... 10**