VMware Tanzu Greenplum Platform Extension Framework v6.2 Documentation

VMware Tanzu Greenplum Platform Extension Framework 6.2

vmware®

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2022 VMware, Inc. All rights reserved. Copyright and trademark information.

Contents

VMware Tanzu™ Greenplum® Platform Extension Framework (PXF) 15

Tanzu Greenplum Platform Extension Framework 6.x Release Notes 17

- Supported Platforms 17
- Upgrading to PXF 6.2.x 17
- Release 6.2.3 17
  - Changed Features 17
  - Resolved Issues 18
- Release 6.2.2 18
  - Changed Features 18
  - Resolved Issues 18
- Release 6.2.1 18
  - Changed Features 18
  - Resolved Issues 19
- Release 6.2.0 19
  - New and Changed Features 19
  - Resolved Issues 20
- Release 6.1.0 20
  - New and Changed Features 20
  - Resolved Issues 21
- Release 6.0.1 21
  - Resolved Issues 21
- Release 6.0.0 21
  - New and Changed Features 21
  - Removed Features 24
  - Resolved Issues 24
- Deprecated Features 24
- Known Issues and Limitations 25

Installing PXF 26

- Prerequisites 26
- Downloading the PXF Package 27
- Installing the PXF Package 27
- Next Steps 28

Installing Java for PXF 29
- Prerequisites 29
- Procedure 29

Uninstalling PXF 30
- Prerequisites 30
- Uninstalling PXF 30

Upgrading to PXF 6 32

Upgrading from PXF 5.x 32
- Step 1: Performing the PXF Pre-Upgrade Actions 32
- Step 2: Installing PXF 6.x 32
- Step 3: Completing the Upgrade to PXF 6.x 33

Upgrading from an Earlier PXF 6 Release 36
- Step 1: Perform the PXF Pre-Upgrade Actions 36
- Step 2: Installing the New PXF 6.x 37
- Step 3: Completing the Upgrade to a Newer PXF 6.x 37

Greenplum Platform Extension Framework (PXF) 39
- Basic Usage 40
- Get Started Configuring PXF 40
- Get Started Using PXF 40

Introduction to PXF 42
- Supported Platforms 42
  - Operating Systems 42
  - Java 42
  - Hadoop 42
- Architectural Overview 42
- About Connectors, Servers, and Profiles 42
- Creating an External Table 43
- Other PXF Features 44

About PXF Filter Pushdown 44

About Column Projection in PXF 46

About the PXF Installation and Configuration Directories 48
- PXF Installation Directories 48
- Relocating $PXF_BASE 49

About the PXF Configuration Files 49
pxf-application.properties 50
pxf-env.sh 50
pxf-log4j2.xml 51
pxf-profiles.xml 51
Modifying the PXF Configuration 51
Configuring PXF 51
Configuring PXF Servers 52
About Server Template Files 53
About the Default Server 54
Configuring a Server 54
About the pxf-site.xml Configuration File 54
About the pxf.fs.basePath Property 56
Configuring a PXF User 56
Procedure 57
About Configuration Property Precedence 57
Using a Server Configuration 58
Configuring PXF Hadoop Connectors (Optional) 58
Prerequisites 59
Procedure 59
About Updating the Hadoop Configuration 60
Configuring the Hadoop User, User Impersonation, and Proxying 61
Use Cases and Configuration Scenarios 61
Accessing Hadoop as the Greenplum User Proxied by gpadmin 61
Accessing Hadoop as the Greenplum User Proxied by a User 62
Accessing Hadoop as the gpadmin User 62
Accessing Hadoop as a User 63
Configure the Hadoop User 63
Configure PXF User Impersonation 64
Configure Hadoop Proxying 65
Hive User Impersonation 66
HBase User Impersonation 66
Configuring PXF for Secure HDFS 66
Prerequisites 67

Use Cases and Configuration Scenarios 67

Accessing Hadoop as the Greenplum User Proxied by the Kerberos Principal 67

Accessing Hadoop as the Kerberos Principal 68

Accessing Hadoop as a User 68

Procedure 69

Configuring PXF with a Microsoft Active Directory Kerberos KDC Server 69

Configuring PXF with an MIT Kerberos KDC Server 71

Configuring Connectors to Minio and S3 Object Stores (Optional) 73

About Object Store Configuration 73

Minio Server Configuration 74

S3 Server Configuration 74

Configuring S3 Server-Side Encryption 74

Configuring SSE via an S3 Bucket Policy (Recommended) 75

Specifying SSE Options in a PXF S3 Server Configuration 75

Example Server Configuration Procedure 76

Configuring Connectors to Azure and Google Cloud Storage Object Stores (Optional) 77

About Object Store Configuration 77

Azure Blob Storage Server Configuration 77

Azure Data Lake Server Configuration 78

Google Cloud Storage Server Configuration 78

Example Server Configuration Procedure 78

Configuring the JDBC Connector (Optional) 79

About JDBC Configuration 80

JDBC Driver JAR Registration 80

JDBC Server Configuration 80

Connection-Level Properties 81

Connection Transaction Isolation Property 81

Statement-Level Properties 82

Session-Level Properties 82

About JDBC Connection Pooling 83

Tuning the Maximum Connection Pool Size 84

JDBC User Impersonation 84

Example Configuration Procedure 85

About Session Authorization 85

Session Authorization Considerations for Connection Pooling 86

JDBC Named Query Configuration 86

Defining a Named Query 87
Query Naming 87
Referencing a Named Query 87
Overriding the JDBC Server Configuration 87
Configuring Access to Hive 88
Example Configuration Procedure 88

Configuring the JDBC Connector for Hive Access (Optional) 89
JDBC Server Configuration 89
Example Configuration Procedure 90

Starting, Stopping, and Restarting PXF 93
Starting PXF 93
Prerequisites 93
Procedure 94
Stopping PXF 94
Prerequisites 94
Procedure 94
Restarting PXF 94
Prerequisites 94
Procedure 94

Granting Users Access to PXF 95
Enabling PXF in a Database 95
Disabling PXF in a Database 95
Granting a Role Access to PXF 96

Registering PXF Library Dependencies 96
Registering a JAR Dependency 96
Registering a Native Library Dependency 97
Procedure 97

Monitoring PXF 98
Viewing PXF Status on the Command Line 98
About PXF Service Runtime Monitoring 99
Examining PXF Metrics 99
Filtering Metric Data 100

PXF Service Host and Port 100
Procedure 101

Logging 101

Configuring the Log Directory 102

Configuring Service-Level Logging 102

Configuring for a Specific Host 103

Configuring for the Cluster 104

Configuring Client-Level Logging 104

Memory and Threading 105

Increasing the JVM Memory for PXF 105

Configuring Out of Memory Condition Actions 106

Auto-Killing the PXF Server 106

Dumping the Java Heap 106

Procedure 107

Another Option for Resource-Constrained PXF Segment Hosts 108

Accessing Hadoop with PXF 110

Architecture 110

Prerequisites 111

HDFS Shell Command Primer 112

Connectors, Data Formats, and Profiles 112

Choosing the Profile 113

Specifying the Profile 114

Reading and Writing HDFS Text Data 114

Prerequisites 114

Reading Text Data 114

Example: Reading Text Data on HDFS 115

Reading Text Data with Quoted Linefeeds 116

Example: Reading Multi-Line Text Data on HDFS 117

Writing Text Data to HDFS 118

Example: Writing Text Data to HDFS 119

Procedure 119

Reading and Writing HDFS Avro Data 121

Prerequisites 122

Working with Avro Data 122

Data Type Mapping 122

Read Mapping 122

Write Mapping 123

Avro Schemas and Data 124

Creating the External Table 124
Example: Reading Avro Data 126
Create Schema 126
Create Avro Data File (JSON) 127
Reading Avro Data 128
Writing Avro Data 128
Example: Writing Avro Data 129

Reading JSON Data from HDFS 130
Prerequisites 130
Working with JSON Data 130
JSON to Greenplum Database Data Type Mapping 131
JSON Data Read Modes 131
Loading the Sample JSON Data to HDFS 132
Creating the External Table 133
Example: Reading a JSON File with Single Line Records 134
Example: Reading a JSON file with Multi-Line Records 134
Other Methods to Read a JSON Array 135
Using Array Element Projection 135
Specifying a Single Text-type Column 136

Reading ORC Data 136
Prerequisites 137
About the ORC Data Format 137
Data Type Mapping 137
Creating the External Table 138
Example: Reading an ORC File on HDFS 139

Reading and Writing HDFS Parquet Data 141
Prerequisites 141
Data Type Mapping 141
Read Mapping 141
Write Mapping 142
Creating the External Table 143
Example 144

Reading and Writing HDFS SequenceFile Data 145
Prerequisites 145
Creating the External Table 145
Reading and Writing Binary Data 147

Example: Writing Binary Data to HDFS 147

Reading the Record Key 150

Example: Using Record Keys 151

Reading a Multi-Line Text File into a Single Table Row 151

Prerequisites 151

Reading Multi-Line Text and JSON Files 151

Example: Reading an HDFS Text File into a Single Table Row 153

Reading Hive Table Data 154

Prerequisites 155

Hive Data Formats 155

Data Type Mapping 155

Primitive Data Types 155

Complex Data Types 156

Sample Data Set 156

Hive Command Line 157

Example: Creating a Hive Table 157

Determining the HDFS Location of a Hive Table 157

Querying External Hive Data 158

Accessing TextFile-Format Hive Tables 159

Example: Using the hive Profile 159

Example: Using the hive:text Profile 159

Accessing RCFile-Format Hive Tables 160

Example: Using the hive:rc Profile 160

Accessing ORC-Format Hive Tables 161

Profiles Supporting the ORC File Format 161

Example: Using the hive:orc Profile 161

Example: Using the Vectorized hive:orc Profile 162

Accessing Parquet-Format Hive Tables 163

Accessing Avro-Format Hive Tables 163

Working with Complex Data Types 164

Example: Using the hive Profile with Complex Data Types 164

Example: Using the hive:orc Profile with Complex Data Types 166

Partition Pruning 167

Example: Using the hive Profile to Access Partitioned Homogenous Data 167

Example: Using the hive Profile to Access Partitioned Heterogeneous Data 169

Using PXF with Hive Default Partitions 171

Reading HBase Table Data 172

Prerequisites 172
HBase Primer 172
HBase Shell 172
Example: Creating an HBase Table 172
Querying External HBase Data 173
Data Type Mapping 174
Column Mapping 174
Direct Mapping 174
Indirect Mapping via Lookup Table 174
Row Key 175

Accessing Azure, Google Cloud Storage, Minio, and S3 Object Stores with PXF 176
Prerequisites 176
Connectors, Data Formats, and Profiles 176
Sample CREATE EXTERNAL TABLE Commands 177

About Accessing the S3 Object Store 178
Overriding the S3 Server Configuration with DDL 178
Using the Amazon S3 Select Service 179

Reading and Writing Text Data in an Object Store 179
Prerequisites 179
Reading Text Data 179
Example: Reading Text Data from S3 180
Reading Text Data with Quoted Linefeeds 182
Example: Reading Multi-Line Text Data from S3 182
Writing Text Data 184
Example: Writing Text Data to S3 185
Procedure 185

Reading and Writing Avro Data in an Object Store 187
Prerequisites 187
Working with Avro Data 187
Creating the External Table 187
Example 188

Reading JSON Data from an Object Store 189
Prerequisites 189
Working with JSON Data 189
Creating the External Table 189

Example 190

Reading ORC Data from an Object Store 190

Prerequisites 191

Data Type Mapping 191

Creating the External Table 191

Example 192

Reading and Writing Parquet Data in an Object Store 192

Prerequisites 192

Data Type Mapping 192

Creating the External Table 193

Example 194

Reading and Writing SequenceFile Data in an Object Store 194

Prerequisites 194

Creating the External Table 194

Example 195

Reading a Multi-Line Text File into a Single Table Row 196

Prerequisites 196

Creating the External Table 196

Example 197

Reading CSV and Parquet Data from S3 Using S3 Select 197

Enabling PXF to Use S3 Select 198

Reading Parquet Data with S3 Select 198

Specifying the Parquet Column Compression Type 198

Creating the External Table 199

Reading CSV files with S3 Select 199

Handling the CSV File Header 199

Specifying the CSV File Compression Type 200

Creating the External Table 200

Accessing an SQL Database with PXF (JDBC) 202

Prerequisites 202

Data Types Supported 202

About Accessing Hive via JDBC 203

Accessing an External SQL Database 203

JDBC Custom Options 204

Batching Insert Operations (Write) 205

Batching on Read Operations 205
Thread Pooling (Write) 205
Partitioning (Read) 206
Examples 207
About Using Named Queries 207
Overriding the JDBC Server Configuration with DDL 208

Example: Reading From and Writing to a PostgreSQL Table 209
Create a PostgreSQL Table 209
Configure the JDBC Connector 210
Read from the PostgreSQL Table 210
Write to the PostgreSQL Table 211

Example: Reading From and Writing to a MySQL Table 211
Create a MySQL Table 212
Configure the MySQL Connector 212
Read from the MySQL Table 213
Write to the MySQL Table 214

Example: Reading From and Writing to an Oracle Table 214
Create an Oracle Table 215
Configure the Oracle Connector 215
Read from the Oracle Table 216
Write to the Oracle Table 217

Example: Reading From and Writing to a Trino (formerly Presto SQL) Table 217
Create a Trino Table 218
Configure the Trino Connector 218
Read from a Trino Table 220
Write to the Trino Table 220

Example: Using a Named Query with PostgreSQL 221
Create the PostgreSQL Tables and Assign Permissions 221
Configure the Named Query 222
Read the Query Results 222

Accessing Files on a Network File System with PXF 225
Prerequisites 225
Configuring a PXF Network File System Server 225
Creating the External Table 227

Example: Reading From and Writing to a CSV File on a Network File System 227
Create a CSV File 228
Create the Network File System Server 228
Read Data 228
Write Data and Read Again 228

Troubleshooting PXF 230
PXF Errors 230
PXF Logging 230
Addressing PXF JDBC Connector Time Zone Errors 231
About PXF External Table Child Partitions 231
Addressing Hive MetaStore Connection Errors 232

PXF Utility Reference 233

pxf cluster 233
Synopsis 233
Description 233
Commands 234
Options 235
Examples 235
See Also 235

pxf 235
Synopsis 235
Description 236
Commands 236
Options 237
Examples 237
See Also 237

VMware Tanzu™ Greenplum® Platform Extension Framework (PXF)

Revised: 2022-02-02

The VMware Tanzu Greenplum Platform Extension Framework (PXF) provides parallel, high throughput data access and federated queries across heterogeneous data sources via built-in connectors that map a Greenplum Database external table definition to an external data source. PXF has its roots in the Apache HAWQ project.

- Release Notes
- Installing PXF
- Uninstalling PXF
- Upgrading to PXF 6
- Overview of PXF
- Introduction to PXF

  This topic introduces PXF concepts and usage.
- Administering PXF

  This set of topics details the administration of PXF including configuration and management procedures.
- Accessing Hadoop with PXF

  This set of topics describe the PXF Hadoop connectors, the data types they support, and the profiles that you can use to read from and write to HDFS.
- Accessing Azure, Google Cloud Storage, Minio, and S3 Object Stores with PXF

  This set of topics describe the PXF object storage connectors, the data types they support, and the profiles that you can use to read data from and write data to the object stores.
- Accessing an SQL Database with PXF (JDBC)

  This topic describes how to use the PXF JDBC connector to read from and write to an external SQL database such as Postgres or MySQL.
- Accessing Files on a Network File System with PXF

  This topic describes how to use PXF to access files on a network file system that is mounted on your Greenplum Database hosts.
- Troubleshooting PXF

  This topic details the service- and database- level logging configuration procedures for PXF. It also identifies some common PXF errors and describes how to address PXF memory

issues.

- PXF Utility Reference

  The PXF utility reference.

Tanzu Greenplum Platform Extension Framework 6.x Release Notes

The Tanzu Greenplum Platform Extension Framework (PXF) is included in the Tanzu Greenplum Database Server distribution in Greenplum version 6.18.x and older, and in version 5.28.0 and older. PXF for Redhat/CentOS and Oracle Enterprise Linux is updated and distributed independently of Greenplum Database starting with PXF version 5.13.0. PXF version 5.16.0 is the first independent release that includes an Ubuntu distribution.

You may need to download and install the PXF package to obtain the most recent version of this component.

Supported Platforms

The independent PXF 6.x distribution is compatible with these operating system platform and Greenplum Database versions:

OS Version	Greenplum Version
RHEL 7.x, CentOS 7.x	5.21.2+, 6.x
OEL 7.x, Ubuntu 18.04 LTS	6.x

PXF is compatible with these Java and Hadoop component versions:

PXF Version	Java Versions	Hadoop Versions	Hive Server Versions	HBase Server Version
6.2.x, 6.1.0, 6.0.x	8, 11	2.x, 3.1+	1.x, 2.x, 3.1+	1.3.2
5.16.x, 5.15.x, 5.14, 5.13	8, 11	2.x, 3.1+	1.x, 2.x, 3.1+	1.3.2

Upgrading to PXF 6.2.x

If you are currently using PXF with Greenplum Database, you may be required to perform upgrade actions for this release. Review Upgrading from PXF 5 or Upgrading from an Earlier PXF 6 Release to plan your upgrade to PXF version 6.2.x.

Release 6.2.3

Release Date: February 2, 2022

Changed Features

PXF 6.2.3 includes these changes:

- PXF bundles version 2.17.1 of the `log4j2` library to mitigate CVE-2021-44832.

- PXF updates the version of `go` that it uses to build the `pxf` CLI tool to version 1.17.6 to mitigate CVE-2021-44716.
- PXF now writes early startup messages that were previously directed to `stdout/stderr` and ignored to the file `$PXF_LOG_DIR/pxf_app.out`.
- PXF introduces a performance improvement when it iterates over a list of fragments.

Resolved Issues

PXF 6.2.3 resolves these issues:

Issue #	Summary
CVE-2021-44832	Updates the bundled `log4j2` library to version 2.17.1. (Resolved by PR-735.)
CVE-2021-44716	Updates the `go` library to version 1.17.6. (Resolved by PR-740.)

Release 6.2.2

Release Date: December 22, 2021

Changed Features

PXF 6.2.2 includes these changes:

- PXF bundles version 2.17.0 of the `log4j2` library to mitigate CVE-2021-45105.
- PXF downgrades the bundled version of Spring Boot to resolve issue 31927.

Resolved Issues

PXF 6.2.2 resolves these issues:

Issue #	Summary
CVE-2021-45105	Updates the bundled `log4j2` library to version 2.17.0. (Resolved by PR-733.)
31927	Resolves an issue where the PXF C extension reported a `partial file transfer` error when a data-less response that the PXF server sent to Greenplum Database failed to include a zero-length chunk. PXF 6.2.2 downgrades the bundled version of Spring Boot to 2.4.3 which does not exhibit the error behavior. (Resolved by PR-732.)

Release 6.2.1

Release Date: December 17, 2021

Changed Features

PXF 6.2.1 includes these changes:

- PXF bundles version 2.16.0 of the `log4j2` library to mitigate CVE-2021-44228 and CVE-2021-45046.
- PXF now returns an `UnsupportedOperationException` when it accesses a Hive transactional table.

- PXF now supports the `SKIP_HEADER_COUNT` option for external tables that specified a `*:text:multi` profile.
- When reading from a MySQL database, PXF now uses a `jdbc.statement.fetchSize` default value of `-2147483648` (`Integer.MIN_VALUE`). This setting enables the MySQL JDBC driver to stream the results from a MySQL server, lessening the memory requirements when reading large data sets.
- The PXF Hive connector now uses the `hive-site.xml` `hive.metastore.failure.retries` property setting to identify the maximum number of times to retry a failed connection to the Hive MetaStore. The default value is one retry. Addressing Hive MetaStore Connection Errors describes when and how to configure this property.

Resolved Issues

PXF 6.2.1 resolves these issues:

Issue #	Summary
CVE-2021-45046	Updates the bundled `log4j2` library to version 2.16.0. (Resolved by PR-727.)
CVE-2021-44228	Updates the bundled `log4j2` library to version 2.15.0. (Resolved by PR-723.)
31955	Resolves an issue where PXF failed to access a Hive table due to a MetaStore connection issue. PXF now includes retry logic for the MetaStore connection based on the `hive.metastore.failure.retries` property setting in the `hive-site.xml` file. (Resolved by PR-726.)
31948	Resolves an issue where PXF ran out of memory when it read a large data set from a MySQL database. PXF now uses a `jdbc.statement.fetchSize` default value of `-2147483648` (`Integer.MIN_VALUE`) when it accesses MySQL, which streams the results from a MySQL server to PXF. (Resolved by PR-721.)
31906	Resolves an issue where PXF returned 0 rows when a query was performed on a Hive transactional table instead of reporting that transactional tables are unsupported. PXF now more clearly identifies the problem by returning an `UnsupportedOperationException` and the error: `PXF does not support Hive transactional tables`. (Resolved by PR-719.)
31791	Resolves an issue where PXF ignored the `SKIP_HEADER_COUNT` custom option when it read from an external data source via an external table that specified a `*:text:multi` profile. PXF now recognizes and implements this option for `*:text:multi` profiles. (Resolved by PR-710.)

Release 6.2.0

Release Date: September 13, 2021

New and Changed Features

PXF 6.2.0 includes these new and changed features:

- PXF adds support for reading a JSON array into a Greenplum Database text array (`TEXT[]`). Refer to Working with JSON Data for additional information.
- PXF adds support for reading lists of certain ORC scalar types into a Greenplum Database array of native type. Refer to the PXF ORC data type mapping documentation for more information about the data type mapping.
- PXF bundles newer versions of ORC, Spring Boot, and other dependent libraries.

- PXF improves its message logging by:
    - Better aligning the log message text.
    - Also logging the affected fragment when it encounters a read error.
- PXF introduces a new property to the pxf-site.xml per-server configuration file. PXF uses this property, `pxf.sasl.connection.retries`, to specify the maximum number of times that it retries a SASL connection request to an external data source after a refused connection returns a `GSS initiate failed` error.
- PXF introduces a new PXF Service application property, `pxf.fragmenter-cache.expiration`, to specify the amount of time after which an entry expires and is removed from the fragment cache.

Resolved Issues

PXF 6.2.0 resolves these issues:

Issue #	Summary
	Resolves an issue when using the `jdbc` profile to write data to a Hive table. The Hive JDBC driver always returned `0` when executing an update, and PXF would return an error even if the `INSERT` executed correctly. (Resolved by PR-662.)
31675	Resolves a fragment cache issue that appeared when an external table was re-created within the same transaction in a stored procedure, and the new external table referenced a different `LOCATION`. (Resolved by PR-691.)
31657	Queries on an external table intermittently failed in some Kerberos-secured environments because the Hadoop NameNode erroneously detected a replay attack during Kerberos authentication. This issue is resolved by PR-688.
31571	PXF did not support ORC lists. PXF 6.2.0 includes support for reading lists of certain ORC scalar types into a Greenplum Database array of native type. (Resolved by PR-675.)
31326	PXF did not support reading a JSON array into a Greenplum Database array-type column. PXF 6.2.0 includes support for reading a JSON array into a text array (`TEXT[]`). (Resolved by PR-646.)
683	Resolves an issue where PXF incorrectly casted an `enum` value from the external data source to a `string`. (Resolved by PR-696.)

Release 6.1.0

Release Date: June 24, 2021

New and Changed Features

PXF 6.1.0 includes these new and changed features:

- PXF now natively supports reading and writing Avro arrays.
- PXF adds support for reading JSON objects, such as embedded arrays, as `text`. The data returned by PXF is a valid JSON string that you can manipulate with the existing Greenplum Database JSON functions and operators.
- PXF improves its error reporting by displaying the exception class when there is no error message available.

- PXF introduces a new property that you can use to configure the connection timeout for data upload/write operations to an external datastore. This property is named `pxf.connection.upload-timeout`, and is located in the pxf-application.properties file.
- PXF now uses the `pxf.connection.timeout` configuration property to set the connection timeout only for read operations. If you previously set this property to specify the write timeout, you should now use `pxf.connection.upload-timeout` instead.
- PXF bundles a newer `gp-common-go-libs` supporting library along with its dependencies.

Resolved Issues

PXF 6.1.0 resolves these issues:

Issue #	Summary
31389	Resolves an issue where certain `pxf cluster` commands returned the error `connect: no such file or directory` when the current working directory contained a directory with the same name as the hostname. This issue was resolved by upgrading a dependent library. (Resolved by PR-633.)
31317	PXF did not support writing Avro arrays. PXF 6.1.0 includes native support for reading and writing Avro arrays. (Resolved by PR-636.)

Release 6.0.1

Release Date: May 11, 2021

Resolved Issues

PXF 6.0.1 resolves these issues:

Issue #	Summary
	Resolves an issue where PXF returned wrong results for batches of ORC data that were shorter than the default batch size. (Resolved by PR-630.)
	Resolves an issue where PXF threw a `NullPointerException` when it encountered a repeating ORC column value of type `string`. (Resolved by PR-627.)
178013439	Resolves an issue where using the profile `HiveVectorizedORC` did not result in vectorized execution. (Resolved by PR-624.)
31409	Resolves an issue where PXF intermittently failed with the error `ERROR: PXF server error(500) : Failed to initialize HiveResolver` when it accessed Hive tables `STORED AS ORC`. (Resolved by PR-626.)

Release 6.0.0

Release Date: March 29, 2021

New and Changed Features

PXF 6.0.0 includes these new and changed features:

Architecture and Bundled Libraries

- PXF 6.0.0 is built on the Spring Boot framework:

    - PXF distributes a single JAR file that includes all of its dependencies.
    - PXF no longer installs and uses a standalone Tomcat server; it uses the Tomcat version 9.0.43 embedded in the PXF Spring Boot application.
- PXF bundles the `postgresql-42.2.14.jar` PostgreSQL driver JAR file.
- PXF library dependencies have changed with new, updated, and removed libraries.
- The PXF API has changed. If you are upgrading from PXF 5.x, you must update the PXF extension in each database in which it is registered as described in Upgrading from PXF 5.
- PXF 6 moves fragment allocation from its C extension to the PXF Service running on each segment host.
- The PXF Service now also runs on the Greenplum Database master and standby master hosts. If you used PXF 5.x to access Kerberos-secured HDFS, you must now generate principals and keytabs for the master and standby master as described in Upgrading from PXF 5.

Files, Configuration, and Commands

- PXF 6 uses the `$PXF_BASE` environment variable to identify its runtime configuration directory; it no longer uses `$PXF_CONF` for this purpose.
- By default, PXF installs its executables and runtime configuration into the same directory, `$PXF_HOME`, and `PXF_BASE=$PXF_HOME`. See About the PXF Installation and Configuration Directories for the new installation file layout.
- You can relocate the `$PXF_BASE` runtime configuration directory to a different directory after you install PXF by running the new `pxf [cluster] prepare` command as described in Relocating $PXF_BASE.
- PXF template server configuration files now reside in `$PXF_HOME/templates`; they were previously located in the `$PXF_CONF/templates` directory.
- The `pxf [cluster] register` command now copies only the PXF `pxf.control` extension file to the Greenplum Database installation. Run this command after your first installation of PXF, and/or after you upgrade your Greenplum Database installation.
- PXF 6 no longer requires initialization, and deprecates the `init` and `reset` commands. `pxf [cluster] init` is now equivalent to `pxf [cluster] register`, and `pxf [cluster] reset` is a no-op.
- PXF 6 includes new and changed configuration; see About the PXF Configuration Files for more information:
    - PXF 6 integrates with Apache Log4j 2; the PXF logging configuration file is now named `pxf-log4j2.xml`, and is in `xml` format.
    - PXF 6 adds a new configuration file for the PXF server application, `pxf-application.properties`; this file includes:
        - New properties to configure the PXF streaming thread pool.
        - New `pxf.log.level` property to set the PXF logging level.
        - Configuration properties moved from the PXF 5 `pxf-env.sh` file and

renamed:

pxf-env.sh Property Name	pxf-application.properties Property Name
PXF_MAX_THREADS	pxf.max.threads

- PXF 6 adds new configuration environment variables to `pxf-env.sh` to simplify the registration of external library dependencies:

New Property Name	Description
PXF_LOADER_PATH	Additional directories and JARs for PXF to class-load.
LD_LIBRARY_PATH	Additional directories and native libraries for PXF to load.

  See Registering PXF Library Dependencies for more information.

- PXF 6 deprecates the `PXF_FRAGMENTER_CACHE` configuration property; fragment metadata caching is no longer configurable and is now always enabled.

Profiles

- PXF 6 introduces new profile names and deprecates some older profile names. The old profile names still work, but it is highly recommended to switch to using the new profile names:

New Profile Name	Old/Deprecated Profile Name
hive	Hive
hive:rc	HiveRC
hive:orc	HiveORC
hive:orc	HiveVectorizedORC[1]
hive:text	HiveText
jdbc	Jdbc
hbase	HBase

  [1] To use the `HiveVectorizedORC` profile in PXF 6, specify the `hive:orc` profile name with the new `VECTORIZE=true` custom option.

- PXF adds support for natively reading an ORC file located in Hadoop, an object store, or a network file system. See the Hadoop ORC and Object Store ORC documentation for prerequisites and usage information.
- PXF adds support for reading and writing comma-separated value form text data located in Hadoop, an object store, or a network file system though a separate `csv` profile. See the Hadoop Text and Object Store Text documentation for usage information.
- PXF supports predicate pushdown on `VARCHAR` data types.
- PXF supports predicate pushdown for the `IN` operator when you specify one of the `*:parquet` profiles to read a parquet file.

- PXF supports specifying a codec short name (alias) rather than the Java class name when you create a writable external table for a `*:text`, `*:csv`, or `*:SequenceFile` profile that includes a `COMPRESSION_CODEC`.

Monitoring

- PXF now supports monitoring of the PXF Service process at runtime. Refer to About PXF Service Runtime Monitoring for more information.

Logging

- PXF improves the display of error messages in the `psql` client, in some cases including a `HINT` that provides possible error resolution actions.
- When PXF is configured to auto-terminate on detection of an out of memory condition, it now logs messages to `$PXF_LOGDIR/pxf-oom.log` rather than `catalina.out`.

Removed Features

PXF version 6.0.0 removes:

- The `THREAD-SAFE` external table custom option (deprecated since 5.10.0).
- The `PXF_USER_IMPERSONATION`, `PXF_PRINCIPAL`, and `PXF_KEYTAB` configuration properties in `pxf-env.sh` (deprecated since 5.10.0).
- The `jdbc.user.impersonation` configuration property in `jdbc-site.xml` (deprecated since 5.10.0).
- The Hadoop profile names `HdfsTextSimple`, `HdfsTextMulti`, `Avro`, `Json`, `Parquet`, and `SequenceWritable` (deprecated since 5.0.1).

Resolved Issues

PXF 6.0.0 resolves these issues:

Issue #	Summary
30987	Resolves an issue where PXF returned an `out of memory` error while executing a query on a Hive table backed by a large number of files when it could not enlarge a string buffer during the fragmentation process. PXF 6.0.0 moves fragment distribution logic and fragment allocation to the PXF Service running on each segment host.

Deprecated Features

Deprecated features may be removed in a future major release of PXF. PXF version 6.x deprecates:

- The `PXF_FRAGMENTER_CACHE` configuration property (deprecated since PXF version 6.0.0).
- The `pxf [cluster] init` commands (deprecated since PXF version 6.0.0).
- The `pxf [cluster] reset` commands (deprecated since PXF version 6.0.0).
- The Hive profile names `Hive`, `HiveText`, `HiveRC`, `HiveORC`, and `HiveVectorizedORC` (deprecated since PXF version 6.0.0). Refer to Connectors, Data Formats, and Profiles in the PXF Hadoop documentation for the new profile names.
- The `HBase` profile name (now `hbase`) (deprecated since PXF version 6.0.0).

- The `Jdbc` profile name (now `jdbc`) (deprecated since PXF version 6.0.0).
- Specifying a `COMPRESSION_CODEC` using the Java class name; use the codec short name instead.

Known Issues and Limitations

PXF 6.x has these known issues and limitations:

Issue #	Description
178013439	(**Resolved in 6.0.1**) Using the deprecated `HiveVectorizedORC` profile does not result in vectorized execution. **Workaround:** Use the `hive:orc` profile with the option `VECTORIZE=true`.
31409	(**Resolved in 6.0.1**) PXF can intermittently fail with the following error when it accesses Hive tables `STORED AS ORC`: `ERROR: PXF server error(500) : Failed to initialize HiveResolver` **Workaround:** Use vectorized query execution by adding the `VECTORIZE=true` custom option to the `LOCATION` URL. (Note that PXF does not support predicate pushdown, complex types, and the `timestamp` data type with ORC vectorized execution.)
168957894	The PXF Hive Connector does not support using the `hive[:*]` profiles to access Hive 3 managed (CRUD and insert-only transactional, and temporary) tables. **Workaround**: Use the PXF JDBC Connector to access Hive 3 managed tables.

Installing PXF

The VMware Tanzu Greenplum Platform Extension Framework (PXF) is available as a separate VMware Tanzu Network download for:

- Tanzu Greenplum Database 5.x for CentOS 7.x and RHEL 7.x platforms
- Tanzu Greenplum Database 6.x for CentOS 7.x, RHEL 7.x, and Ubuntu 18.04 LTS platforms

The PXF download package is an `.rpm` or `.deb` file that installs PXF libraries, executables, and script files on a Greenplum Database host.

When you install PXF, you will:

1. Satisfy the prerequisites.
2. Download the PXF package.
3. Install the PXF package on every host in your Greenplum Database cluster.
4. Check out Next Steps for post-install topics.

Prerequisites

The recommended deployment model is to install PXF on all Greenplum Database hosts. Before you install PXF 6, ensure that you meet the following prerequisites:

- Tanzu Greenplum version 5.21.2 or later or 6.x is installed in the cluster.
- You have access to all hosts (master, standby master, and segment hosts) in your Greenplum Database cluster.
- You must be an operating system superuser, or have `sudo` privileges, to install the PXF package. If you are installing on CentOS/RHEL, you can choose to install the package into a custom file system location.
- You have installed Java 8 or 11 on all Greenplum Database hosts as described in Installing Java for PXF.
- You can identify the operating system user that will own the PXF installation. This user must be the same user that owns the Greenplum Database installation, or a user that has write privileges to the Greenplum Database installation directory.
- If you have previously configured and are using PXF in your Greenplum installation:
    1. Identify and note the current PXF version number.
    2. Stop PXF as described in Stopping PXF.

If this is your first installation of a PXF package, and the `$GPHOME/pxf` directory exists in your Greenplum installation, you may choose to remove the directory on all Greenplum hosts **after** you confirm that you have installed and configured PXF correctly and that it is working as expected.

If you choose to remove this directory, you may encounter `warning: <pxf-filename>: remove failed: No such file or directory` messages when you upgrade Greenplum. You can ignore these warnings for PXF files.

Downloading the PXF Package

Follow this procedure to download PXF:

1. Navigate to VMware Tanzu Network and locate and select the *Release Download* directory named *Greenplum Platform Extension Framework*.

   The format of the PXF download file name is `pxf-gp<greenplum-major-version>-<pxf-version>-<pkg-version>.<platform>.<file_type>`. For example:

   ```
   pxf-gp6-6.2.3-2.el7.x86_64.rpm
   ```

   or

   ```
   pxf-gp6-6.2.3-2-ubuntu18.04-amd64.deb
   ```

2. Select the appropriate PXF package for your Greenplum Database major version and operating system platform.
3. Make note of the directory to which the file was downloaded.

Installing the PXF Package

You must install the PXF package on the Greenplum Database master and standby master hosts, and on each segment host.

If you installed an older version of the PXF package on your hosts, installing a newer package removes the existing PXF installation, and installs the new version.

The install procedure follows:

1. Locate the installer file that you downloaded from VMware Tanzu Network.
2. Create a text file that lists your Greenplum Database standby master host and segment hosts, one host name per line. For example, a file named `gphostfile` may include:

   ```
   gpmaster
   mstandby
   seghost1
   seghost2
   seghost3
   ```

3. Copy the downloaded PXF package file to all hosts in your Greenplum cluster. For example, to copy the `rpm` to the `/tmp` directory on each host:

   ```
   gphost$ gpscp -f gphostfile pxf-gp6-6.2.3-2.el7.x86_64.rpm =:/tmp/
   ```

4. Install the package on each Greenplum Database host using your package management utility. If a previous installation of PXF exists for the same Greenplum version, the files and runtime directories from the older version are removed before the current package is installed.

1. **To install PXF into the default location on all Greenplum hosts**:

   On a CentOS/RHEL system:

   ```
   gphost$ gpssh -e -v -f gphostfile "sudo rpm -Uvh /tmp/pxf-gp6-6.2.3-2.el7.x86_64.rpm"
   ```

   On an Ubuntu system:

   ```
   gphost$ gpssh -e -v -f gphostfile "sudo dpkg --install /tmp/pxf-gp6-6.2.3-2-ubuntu18.04-amd64.deb"
   ```

   The default PXF package installation directory is `/usr/local/pxf-gp<greenplum-major-version>`.

2. **To install PXF into a custom location on all Greenplum hosts** (CentOS/RHEL only):

   ```
   gpadmin@gphost$ gpssh -e -v -f gphostfile "sudo rpm -Uvh --prefix <install-location> pxf-gp6-6.2.3-2.el7.x86_64.rpm"
   ```

5. Set the ownership and permissions of the PXF installation files to enable access by the `gpadmin` user. For example, if you installed PXF to the default location:

   ```
   gphost$ gpssh -e -v -f gphostfile "sudo chown -R gpadmin:gpadmin /usr/local/pxf-gp*"
   ```

   If you installed PXF to a custom `<install-location>` on CentOS/RHEL, specify that location in the command.

6. (Optional) Add the PXF `bin` directory to the PXF owner's `$PATH`. For example, if you installed PXF for Greenplum 6 in the default location, you could add the following text to the `.bashrc` shell initialization script for the `gpadmin` user:

   ```
   export PATH=$PATH:/usr/local/pxf-gp6/bin
   ```

   Be sure to remove any previously-added `$PATH` entries for PXF in `$GPHOME/pxf/bin`.

7. Remove the PXF package download file that you copied to each system. For example, to remove the `rpm` from `/tmp`:

   ```
   gpadmin@gphost$ gpssh -e -v -f gphostfile "rm -f /tmp/pxf-gp6-6.2.3-2.el7.x86_64.rpm"
   ```

Next Steps

PXF is not active after installation. You must explicitly initialize and start the PXF server before you can use PXF.

- See About the PXF Installation and Configuration Directories for a list and description of important PXF files and directories. This topic also provides instructions about relocating the PXF runtime configuration directories.
- If this is your first time using PXF, review Configuring PXF for a description of the initialization

and configuration procedures that you must perform before you can use PXF.

- If you installed the PXF `rpm` or `deb` as part of a Greenplum Database upgrade procedure, return to those upgrade instructions.
- If you installed the PXF `rpm` or `deb` into a Greenplum cluster in which you had already configured and were using PXF 5, you are required to perform some upgrade actions. Recall the original version of PXF (before you installed the `rpm` or `deb`), and perform Step 3 of the PXF upgrade procedure.

Installing Java for PXF

PXF is a Java service. It requires a Java 8 or Java 11 installation on each Greenplum Database host.

Prerequisites

Ensure that you have access to, or superuser permissions to install, Java 8 or Java 11 on each Greenplum Database host.

Procedure

Perform the following procedure to install Java on the master, standby master, and on each segment host in your Greenplum Database cluster. You will use the `gpssh` utility where possible to run a command on multiple hosts.

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Determine the version(s) of Java installed on the system:

   ```
   gpadmin@gpmaster$ rpm -qa | grep java
   ```

3. If the system does not include a Java version 8 or 11 installation, install one of these Java versions on the master, standby master, and on each Greenplum Database segment host.

   1. Create a text file that lists your Greenplum Database standby master host and segment hosts, one host name per line. For example, a file named `gphostfile` may include:

      ```
      gpmaster
      mstandby
      seghost1
      seghost2
      seghost3
      ```

   2. Install the Java package on each host. For example, to install Java version 8:

      ```
      gpadmin@gpmaster$ gpssh -e -v -f gphostfile sudo yum -y install java-1.8.0-openjdk-1.8.0*
      ```

4. Identify the Java 8 or 11 `$JAVA_HOME` setting for PXF. For example:

   If you installed Java 8:

```
JAVA_HOME=/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.x86_64/jre
```

If you installed Java 11:

```
JAVA_HOME=/usr/lib/jvm/java-11-openjdk-11.0.4.11-0.el7_6.x86_64
```

If the superuser configures the newly-installed Java alternative as the system default:

```
JAVA_HOME=/usr/lib/jvm/jre
```

5. Note the `$JAVA_HOME` setting; you will need this value when you configure PXF.

Uninstalling PXF

The PXF download package is an `.rpm` or `.deb` file that installs PXF libraries, executables, and script files on a Greenplum Database host.

If you want to remove PXF from the Greenplum cluster and from your hosts, you will:

1. Satisfy the prerequisites.
2. Uninstall PXF from every host in your Greenplum Database cluster.

Prerequisites

Before you uninstall PXF, ensure that you meet the following prerequisites:

- You have access to all hosts (master, standby master, and segment hosts) in your Greenplum Database cluster.
- You must be an operating system superuser, or have `sudo` privileges, to remove the PXF package.

Uninstalling PXF

Follow these steps to remove PXF from your Greenplum Database cluster:

1. Log in to the Greenplum Database master node. For example:

```
$ ssh gpadmin@<gpmaster>
```

2. Stop PXF as described in Stopping PXF.
3. Remove the PXF library and extension files from your Greenplum installation:

```
gpadmin@gpmaster$ rm $GPHOME/lib/postgresql/pxf.so
gpadmin@gpmaster$ rm $GPHOME/share/postgresql/extension/pxf*
```

4. Remove PXF from each Greenplum Database master, standby, and segment host. You must be an operating system superuser, or have `sudo` privileges, to remove the package. For example, if you installed PXF for Greenplum 6 in the default location on a CentOS 7 system, the following command removes the PXF package on all hosts listed in `gphostfile`:

   On a CentOS/RHEL system:

```
gpadmin@gpmaster$ gpssh -e -v -f gphostfile "sudo rpm -e pxf-gp6"
```

On an Ubuntu system:

```
gpadmin@gpmaster$ gpssh -e -v -f gphostfile "sudo dpkg --remove pxf-gp6"
```

The command removes the PXF install files on all Greenplum hosts. The command also removes the PXF runtime directories on all hosts.

The PXF configuration directory `$PXF_CONF` is not affected by this command and remains on the Greenplum hosts.

Upgrading to PXF 6

PXF 6.x supports these upgrade paths:

- Upgrading from PXF 5
- Upgrading from an earlier PXF 6 Release

Upgrading from PXF 5.x

If you have installed, configured, and are using PXF 5.x in your Greenplum Database 5 or 6 cluster, you must perform some upgrade actions when you install PXF 6.x.

If you are using PXF with Greenplum Database 5, you must upgrade Greenplum to version 5.21.2 or newer before you upgrade to PXF 6.x.

The PXF upgrade procedure has three steps. You perform one pre-install procedure, the install itself, and then a post-install procedure to upgrade to PXF 6.x:

- Step 1: Perform the PXF Pre-Upgrade Actions
- Step 2: Install PXF 6.x
- Step 3: Complete the Upgrade to PXF 6.x

Step 1: Performing the PXF Pre-Upgrade Actions

Perform this procedure before you upgrade to a new version of PXF:

1. Log in to the Greenplum Database master node. For example:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Identify and note the version of PXF currently running in your Greenplum cluster:

   ```
   gpadmin@gpmaster$ pxf version
   ```

3. Identify the file system location of the `$PXF_CONF` setting in your PXF 5.x PXF installation; you will need this later. If you are unsure of the location, you can find the value in `pxf-env-default.sh`.
4. Stop PXF on each Greenplum host as described in Stopping PXF.

Step 2: Installing PXF 6.x

1. Install PXF 6.x and identify and note the new PXF version number.
2. Check out the new installation layout in About the PXF Installation and Configuration Directories.

Step 3: Completing the Upgrade to PXF 6.x

After you install the new version of PXF, perform the following procedure:

1. Log in to the Greenplum Database master node. For example:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. You must run the `pxf` commands specified in subsequent steps using the binaries from your PXF 6.x installation. Ensure that the PXF 6.x installation `bin/` directory is in your `$PATH`, or provide the full path to the `pxf` command. You can run the following command to check the `pxf` version:

   ```
   gpadmin@gpmaster$ pxf version
   ```

3. (Optional, Advanced) If you want to relocate `$PXF_BASE` outside of `$PXF_HOME`, perform the procedure described in Relocating $PXF_BASE.

4. Auto-migrate your PXF 5.x configuration to PXF 6.x `$PXF_BASE`:

   1. Recall your PXF 5.x `$PXF_CONF` setting.
   2. Run the `migrate` command (see pxf cluster migrate). You must provide `PXF_CONF`. If you relocated `$PXF_BASE`, provide that setting as well.

      ```
      gpadmin@gpmaster$ PXF_CONF=/path/to/dir pxf cluster migrate
      ```

      Or:

      ```
      gpadmin@gpmaster$ PXF_CONF=/path/to/dir PXF_BASE=/new/dir pxf cluster mig
      rate
      ```

      The command copies PXF 5.x `conf/pxf-profiles.xml`, `servers/*`, `lib/*`, and `keytabs/*` to the PXF 6.x `$PXF_BASE` directory. The command also merges configuration changes in the PXF 5.x `conf/pxf-env.sh` into the PXF 6.x file of the same name and into `pxf-application.properties`.

   3. The `migrate` command does not migrate PXF 5.x `$PXF_CONF/conf/pxf-log4j.properties` customizations; you must manually migrate any changes that you made to this file to `$PXF_BASE/conf/pxf-log4j2.xml`. Note that PXF 5.x `pxf-log4j.properties` is in properties format, and PXF 6 `pxf-log4j2.xml` is `xml` format. See the Configuration with XML topic in the Apache Log4j 2 documentation for more information.

5. *If you migrated your PXF 6.x `$PXF_BASE` configuration (see previous step), be sure to apply any changes identified in subsequent steps to the new, migrated directory.*

6. **If you are upgrading from PXF version 5.9.x or earlier** and you have configured any JDBC servers that access Kerberos-secured Hive, you must now set the `hadoop.security.authentication` property to the `jdbc-site.xml` file to explicitly identify use of the Kerberos authentication method. Perform the following for each of these server configs:

1. Navigate to the server configuration directory.
2. Open the `jdbc-site.xml` file in the editor of your choice and uncomment or add the following property block to the file:

```
<property>
    <name>hadoop.security.authentication</name>
    <value>kerberos</value>
</property>
```

3. Save the file and exit the editor.

7. **If you are upgrading from PXF version 5.11.x or earlier**: The PXF `Hive` and `HiveRC` profiles (named `hive` and `hive:rc` in PXF version 6.x) now support column projection using column name-based mapping. If you have any existing PXF external tables that specify one of these profiles, and the external table relied on column index-based mapping, you may be required to drop and recreate the tables:
   1. Identify all PXF external tables that you created that specify a `Hive` or `HiveRC` profile.
   2. For *each* external table that you identify in step 1, examine the definitions of both the PXF external table and the referenced Hive table. If the column names of the PXF external table *do not* match the column names of the Hive table:
      1. Drop the existing PXF external table. For example:

```
DROP EXTERNAL TABLE pxf_hive_table1;
```

      2. Recreate the PXF external table using the Hive column names. For example:

```
CREATE EXTERNAL TABLE pxf_hive_table1( hivecolname int, hivecolnam
e2 text )
  LOCATION( 'pxf://default.hive_table_name?PROFILE=hive')
FORMAT 'custom' (FORMATTER='pxfwritable_import');
```

      3. Review any SQL scripts that you may have created that reference the PXF external table, and update column names if required.

8. **If you are upgrading from PXF version 5.15.x or earlier**:
   1. The `pxf.service.user.name` property in the `pxf-site.xml` template file is now commented out by default. Keep this in mind when you configure new PXF servers.
   2. The default value for the `jdbc.pool.property.maximumPoolSize` property is now `15`. If you have previously configured a JDBC server and want that server to use the new default value, you must manually change the property value in the server's `jdbc-site.xml` file.
   3. PXF 5.16 disallows specifying relative paths and environment variables in the `CREATE EXTERNAL TABLE LOCATION` clause file path. If you previously created any external tables that specified a relative path or environment variable, you must drop each external table, and then re-create it without these constructs.
   4. Filter pushdown is enabled by default for queries on external tables that specify the `Hive`, `HiveRC`, or `HiveORC` profiles (named `hive`, `hive:rc`, and `hive:orc` in PXF version 6.x). If you have previously created an external table that specifies one of these

profiles and queries are failing with PXF v5.16+, you can disable filter pushdown at the external table-level or at the server level:

1. (External table) Drop the external table and re-create it, specifying the `&PPD=false` option in the `LOCATION` clause.
2. (Server) If you do not want to recreate the external table, you can disable filter pushdown *for all* `Hive*` (named as described here in PXF version 6.x) *profile queries using the server* by setting the `pxf.ppd.hive` property in the `pxf-site.xml` file to `false`:

```
<property>
    <name>pxf.ppd.hive</name>
    <value>false</value>
</property>
```

   You may need to add this property block to the `pxf-site.xml` file.

9. Register the PXF 6.x extension files with Greenplum Database (see pxf cluster register). `$GPHOME` must be set when you run this command.

```
gpadmin@gpmaster$ pxf cluster register
```

   The `register` command copies only the `pxf.control` extension file to the Greenplum cluster. In PXF 6.x, the PXF extension `.sql` file and library `pxf.so` reside in `$PXF_HOME/gpextable`. You may choose to remove these now-unused files from the Greenplum Database installation *on the Greenplum Database master, standby master, and all segment hosts*. For example, to remove the files on the master host:

```
gpadmin@gpmaster$ rm $GPHOME/share/postgresql/extension/pxf--1.0.sql
gpadmin@gpmaster$ rm $GPHOME/lib/postgresql/pxf.so
```

10. PXF 6.x includes a new version of the `pxf` extension. You must update the extension in every Greenplum database in which you are using PXF. A database superuser or the database owner must run this SQL command in the `psql` subsystem or in an SQL script:

```
ALTER EXTENSION pxf UPDATE;
```

11. Ensure that you no longer reference previously-deprecated features that were removed in PXF 6.0:

Deprecated Feature	Use Instead
Hadoop profile names	`hdfs:<profile>` as noted here
`jdbc.user.impersonation` property	`pxf.service.user.impersonation` property in the jdbc-site.xml server configuration file
`PXF_KEYTAB` configuration property	`pxf.service.kerberos.keytab` property in the pxf-site.xml server configuration file
`PXF_PRINCIPAL` configuration property	`pxf.service.kerberos.principal` property in the pxf-site.xml server configuration file

Deprecated Feature	Use Instead
`PXF_USER_IMPERSONATION` configuration property	`pxf.service.user.impersonation` property in the pxf-site.xml server configuration file

12. PXF 6.x distributes a single JAR file that includes all of its dependencies, and separately makes its HBase JAR file available in `$PXF_HOME/share`. If you have configured a PXF Hadoop server for HBase access, you must register the new `pxf-hbase-<version>.jar` with Hadoop and HBase as follows:
    1. Copy `$PXF_HOME/share/pxf-hbase-<version>.jar` to each node in your HBase cluster.
    2. Add the location of this JAR to `$HBASE_CLASSPATH` on each HBase node.
    3. Restart HBase on each node.
13. In PXF 6.x, the PXF Service runs on all Greenplum Database hosts. If you used PXF 5.x to access Kerberos-secured HDFS, you must now generate principals and keytabs for the Greenplum master and standby master hosts, and distribute these to the hosts as described in Configuring PXF for Secure HDFS.
14. Synchronize the PXF 6.x configuration from the master host to the standby master and each Greenplum Database segment host. For example:

    ```
    gpadmin@gpmaster$ pxf cluster sync
    ```

15. Start PXF on each Greenplum host. For example:

    ```
    gpadmin@gpmaster$ pxf cluster start
    ```

16. Verify that PXF can access each external data source by querying external tables that specify each PXF server.

Upgrading from an Earlier PXF 6 Release

If you have installed a PXF 6.x `rpm` or `deb` package and have configured and are using PXF in your current Greenplum Database 5.21.2+ or 6.x installation, you must perform some upgrade actions when you install a new version of PXF 6.x.

The PXF upgrade procedure has three steps. You perform one pre-install procedure, the install itself, and then a post-install procedure to upgrade to PXF 6.x

- Step 1: Perform the PXF Pre-Upgrade Actions
- Step 2: Install the New PXF 6.x
- Step 3: Complete the Upgrade to a Newer PXF 6.x

Step 1: Perform the PXF Pre-Upgrade Actions

Perform this procedure before you upgrade to a new version of PXF 6.x:

1. Log in to the Greenplum Database master node. For example:

    ```
    $ ssh gpadmin@<gpmaster>
    ```

2. Identify and note the version of PXF currently running in your Greenplum cluster:

```
gpadmin@gpmaster$ pxf version
```

3. Stop PXF on each Greenplum host as described in Stopping PXF:

```
gpadmin@gpmaster$ pxf cluster stop
```

4. (Optional, Recommended) Back up the PXF user configuration files; for example, if `PXF_BASE=/usr/local/pxf-gp6`:

```
gpadmin@gpmaster$ cp -avi /usr/local/pxf-gp6 pxf_base.bak
```

Step 2: Installing the New PXF 6.x

Install PXF 6.x and identify and note the new PXF version number.

Step 3: Completing the Upgrade to a Newer PXF 6.x

After you install the new version of PXF, perform the following procedure:

1. Log in to the Greenplum Database master node. For example:

```
$ ssh gpadmin@<gpmaster>
```

2. PXF 6.x includes a new version of the `pxf` extension. Register the extension files with Greenplum Database (see pxf cluster register). `$GPHOME` must be set when you run this command:

```
gpadmin@gpmaster$ pxf cluster register
```

3. You must update the `pxf` extension in every Greenplum database in which you are using PXF. A database superuser or the database owner must run this SQL command in the `psql` subsystem or in an SQL script:

```
ALTER EXTENSION pxf UPDATE;
```

4. **If you are upgrading *from* PXF version 6.0.x**:

    - If you previously set the `pxf.connection.timeout` property to change the write/upload timeout, you must now set the `pxf.connection.upload-timeout` property for this purpose.
    - Existing external tables that access Avro arrays and JSON objects will continue to work as-is. If you want to take advantage of the new Avro array read/write functionality or the new JSON object support, create a new external table with the adjusted DDL. If you can access the data with the new external table as you expect, you may choose to drop and recreate the existing external table.

5. **If you are upgrading *to* PXF version 6.2.0 to resolve an erroneous replay attack issue in a Kerberos-secured environment**:

1. If you want to change the default value of the new `pxf.sasl.connection.retries` property, add the following to the `pxf-site.xml` file for your PXF server:

```
<property>
  <name>pxf.sasl.connection.retries</name>
  <value><new-value></value>
  <description>
    Specifies the number of retries to perform when a SASL connection is refused by a Namenode due to 'GSS initiate failed' error.
  </description>
</property>
```

2. (Recommended) Configure PXF to use a host-specific Kerberos principal for each segment host. If you specify the following `pxf.service.kerberos.principal` property setting in the PXF server's `pxf-site.xml` file, PXF automatically replaces `_HOST` with the FQDN of the segment host:

```
<property>
  <name>pxf.service.kerberos.principal</name>
  <value>gpadmin/_HOST@REALM.COM</value>
</property>
```

6. (Recommended) **If you are upgrading *from* PXF version 6.2.2 or earlier *to* PXF version 6.2.3 or later**, update your `$PXF_BASE/conf/pxf-log4j2.xml` file to fully configure the logging changes introduced in version 6.2.3:

   1. Remove the following line from the initial `<Properties>` block:

```
<Property name="PXF_LOG_LEVEL">${bundle:pxf-application:pxf.log.level}</Property>
```

   2. Change the following line:

```
<Logger name="org.greenplum.pxf" level="${env:PXF_LOG_LEVEL:-${sys:PXF_LOG_LEVEL:-info}}"/>
```

   to:

```
<Logger name="org.greenplum.pxf" level="${env:PXF_LOG_LEVEL:-${spring:pxf.log.level}}"/>
```

7. Synchronize the PXF configuration from the master host to the standby master and each Greenplum Database segment host. For example:

```
gpadmin@gpmaster$ pxf cluster sync
```

8. Start PXF on each Greenplum host as described in Starting PXF:

```
gpadmin@gpmaster$ pxf cluster start
```

Greenplum Platform Extension Framework (PXF)

With the explosion of data stores and cloud services, data now resides across many disparate systems and in a variety of formats. Often, data is classified both by its location and the operations performed on the data, as well as how often the data is accessed: real-time or transactional (hot), less frequent (warm), or archival (cold).

The diagram below describes a data source that tracks monthly sales across many years. Real-time operational data is stored in MySQL. Data subject to analytic and business intelligence operations is stored in Greenplum Database. The rarely accessed, archival data resides in AWS S3.

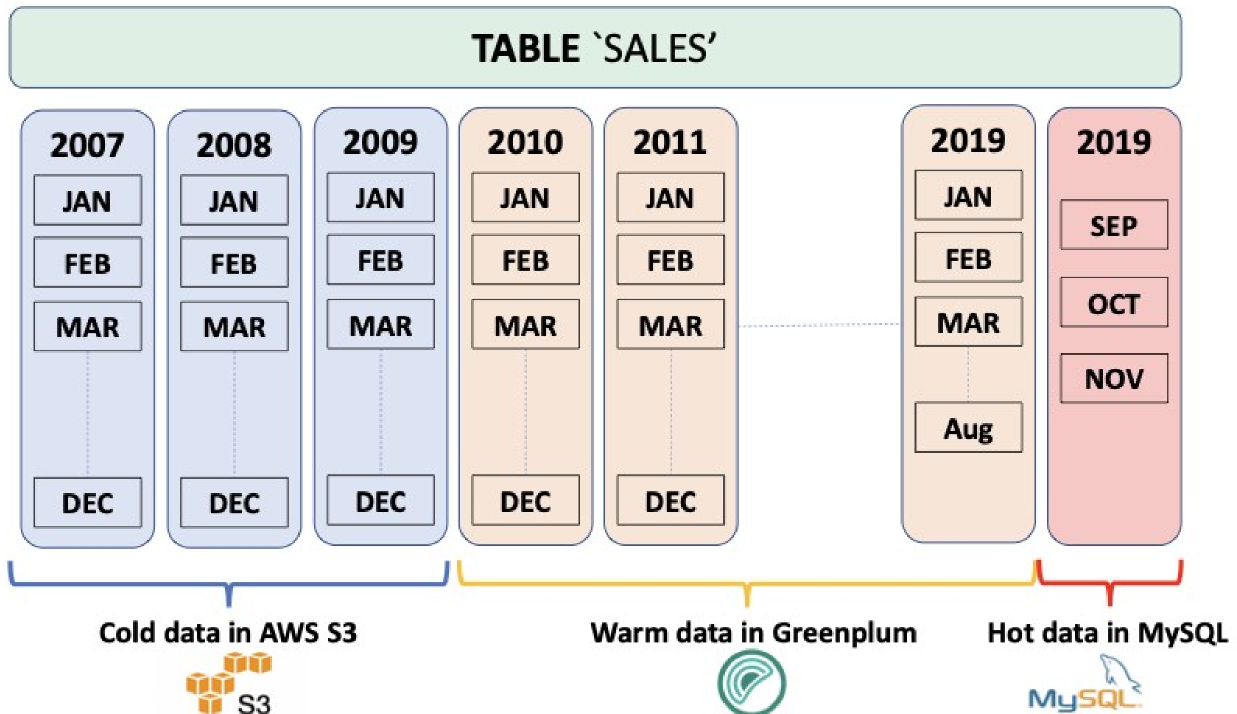

When multiple, related data sets exist in external systems, it is often more efficient to join data sets remotely and return only the results, rather than negotiate the time and storage requirements of performing a rather expensive full data load operation. The *Greenplum Platform Extension Framework (PXF)*, a Greenplum extension that provides parallel, high throughput data access and federated query processing, provides this capability.

With PXF, you can use Greenplum and SQL to query these heterogeneous data sources:

- Hadoop, Hive, and HBase
- Azure Blob Storage and Azure Data Lake
- AWS S3
- Minio

- Google Cloud Storage
- SQL databases including Apache Ignite, Hive, MySQL, ORACLE, Microsoft SQL Server, DB2, and PostgreSQL (via JDBC)
- Network file systems

And these data formats:

- Avro, AvroSequenceFile
- JSON
- ORC
- Parquet
- RCFile
- SequenceFile
- Text (plain, delimited, embedded line feeds)

Basic Usage

You use PXF to map data from an external source to a Greenplum Database *external table* definition. You can then use the PXF external table and SQL to:

- Perform queries on the external data, leaving the referenced data in place on the remote system.
- Load a subset of the external data into Greenplum Database.
- Run complex queries on local data residing in Greenplum tables and remote data referenced via PXF external tables.
- Write data to the external data source.

Check out the PXF introduction for a high level overview important PXF concepts.

Get Started Configuring PXF

The Greenplum Database administrator manages PXF, Greenplum Database user privileges, and external data source configuration. Tasks include:

- Installing, configuring, starting, monitoring, and troubleshooting the PXF Service.
- Managing PXF upgrade.
- Configuring and publishing one or more server definitions for each external data source. This definition specifies the location of, and access credentials to, the external data source.
- Granting Greenplum user access to PXF and PXF external tables.

Get Started Using PXF

A Greenplum Database user creates a PXF external table that references a file or other data in the external data source, and uses the external table to query or load the external data in Greenplum. Tasks are external data store-dependent:

- See Accessing Hadoop with PXF when the data resides in Hadoop.
- See Accessing Azure, Google Cloud Storage, Minio, and S3 Object Stores with PXF when the data resides in an object store.
- See Accessing an SQL Database with PXF when the data resides in an external SQL database.

Introduction to PXF

The Greenplum Platform Extension Framework (PXF) provides *connectors* that enable you to access data stored in sources external to your Greenplum Database deployment. These connectors map an external data source to a Greenplum Database *external table* definition. When you create the Greenplum Database external table, you identify the external data store and the format of the data via a *server* name and a *profile* name that you provide in the command.

You can query the external table via Greenplum Database, leaving the referenced data in place. Or, you can use the external table to load the data into Greenplum Database for higher performance.

Supported Platforms

Operating Systems

PXF supports the Red Hat Enterprise Linux 64-bit 7.x, CentOS 64-bit 7.x, and Ubuntu 18.04 LTS operating system platforms.

Java

PXF supports Java 8 and Java 11.

Hadoop

PXF bundles all of the Hadoop JAR files on which it depends, and supports the following Hadoop component versions:

PXF Version	Hadoop Version	Hive Server Version	HBase Server Version
6.x	2.x, 3.1+	1.x, 2.x, 3.1+	1.3.2
5.9+	2.x, 3.1+	1.x, 2.x, 3.1+	1.3.2
5.8	2.x	1.x	1.3.2

Architectural Overview

Your Greenplum Database deployment consists of a master host, a standby master host, and multiple segment hosts. A single PXF Service process runs on each Greenplum Database host. The PXF Service process running on a segment host allocates a worker thread for each segment instance on the host that participates in a query against an external table. The PXF Services on multiple segment hosts communicate with the external data store in parallel. The PXF Service process running on the master and standby master hosts are not currently involved in data transfer; these processes may be used for other purposes in the future.

About Connectors, Servers, and Profiles

Connector is a generic term that encapsulates the implementation details required to read from or write to an external data store. PXF provides built-in connectors to Hadoop (HDFS, Hive, HBase), object stores (Azure, Google Cloud Storage, Minio, S3), and SQL databases (via JDBC).

A PXF *Server* is a named configuration for a connector. A server definition provides the information required for PXF to access an external data source. This configuration information is data-store-specific, and may include server location, access credentials, and other relevant properties.

The Greenplum Database administrator will configure at least one server definition for each external data store that they will allow Greenplum Database users to access, and will publish the available server names as appropriate.

You specify a `SERVER=<server_name>` setting when you create the external table to identify the server configuration from which to obtain the configuration and credentials to access the external data store.

The default PXF server is named `default` (reserved), and when configured provides the location and access information for the external data source in the absence of a `SERVER=<server_name>` setting.

Finally, a PXF *profile* is a named mapping identifying a specific data format or protocol supported by a specific external data store. PXF supports text, Avro, JSON, RCFile, Parquet, SequenceFile, and ORC data formats, and the JDBC protocol, and provides several built-in profiles as discussed in the following section.

Creating an External Table

PXF implements a Greenplum Database protocol named `pxf` that you can use to create an external table that references data in an external data store. The syntax for a `CREATE EXTERNAL TABLE` command that specifies the `pxf` protocol follows:

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
        ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION('pxf://<path-to-data>?PROFILE=<profile_name>[&SERVER=<server_name>][&<custom-
option>=<value>[...]]')
FORMAT '[TEXT|CSV|CUSTOM]' (<formatting-properties>);
```

The `LOCATION` clause in a `CREATE EXTERNAL TABLE` statement specifying the `pxf` protocol is a URI. This URI identifies the path to, or other information describing, the location of the external data. For example, if the external data store is HDFS, the <path-to-data> identifies the absolute path to a specific HDFS file. If the external data store is Hive, <path-to-data> identifies a schema-qualified Hive table name.

You use the query portion of the URI, introduced by the question mark (?), to identify the PXF server and profile names.

PXF may require additional information to read or write certain data formats. You provide profile-specific information using the optional <custom-option>=<value> component of the `LOCATION` string and formatting information via the <formatting-properties> component of the string. The custom options and formatting properties supported by a specific profile vary; they are identified in usage documentation for the profile.

Table 1. CREATE EXTERNAL TABLE Parameter Values and Descriptions

Keyword	Value and Description
<path-to-data>	A directory, file name, wildcard pattern, table name, etc. The syntax of <path-to-data> is dependent upon the external data source.
PROFILE=<profile_name>	The profile that PXF uses to access the data. PXF supports profiles that access text, Avro, JSON, RCFile, Parquet, SequenceFile, and ORC data in Hadoop services, object stores, network file systems, and other SQL databases.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>=<value>	Additional options and their values supported by the profile or the server.
FORMAT <value>	PXF profiles support the `TEXT`, `CSV`, and `CUSTOM` formats.
<formatting-properties>	Formatting properties supported by the profile; for example, the `FORMATTER` or `delimiter`.

Note: When you create a PXF external table, you cannot use the `HEADER` option in your formatter specification.

Other PXF Features

Certain PXF connectors and profiles support filter pushdown and column projection. Refer to the following topics for detailed information about this support:

- About PXF Filter Pushdown
- About Column Projection in PXF

About PXF Filter Pushdown

PXF supports filter pushdown. When filter pushdown is enabled, the constraints from the `WHERE` clause of a `SELECT` query can be extracted and passed to the external data source for filtering. This process can improve query performance, and can also reduce the amount of data that is transferred to Greenplum Database.

You enable or disable filter pushdown for all external table protocols, including `pxf`, by setting the `gp_external_enable_filter_pushdown` server configuration parameter. The default value of this configuration parameter is `on`; set it to `off` to disable filter pushdown. For example:

```
SHOW gp_external_enable_filter_pushdown;
SET gp_external_enable_filter_pushdown TO 'on';
```

Note: Some external data sources do not support filter pushdown. Also, filter pushdown may not be supported with certain data types or operators. If a query accesses a data source that does not support filter push-down for the query constraints, the query is instead executed without filter pushdown (the data is filtered after it is transferred to Greenplum Database).

PXF filter pushdown can be used with these data types (connector- and profile-specific):

- `INT2`, `INT4`, `INT8`
- `CHAR`, `TEXT`, `VARCHAR`
- `FLOAT`

- `NUMERIC` (not available with the S3 connector when using S3 Select, nor with the `hive` profile when accessing `STORED AS Parquet`)
- `BOOL`
- `DATE`, `TIMESTAMP` (available only with the JDBC connector, the S3 connector when using S3 Select, the `hive:rc` and `hive:orc` profiles, and the `hive` profile when accessing `STORED AS RCFile` or `ORC`)

PXF accesses data sources using profiles exposed by different connectors, and filter pushdown support is determined by the specific connector implementation. The following PXF profiles support some aspects of filter pushdown as well as different arithmetic and logical operations:

Profile	<, >, <=, >=, =, <>	LIKE	IS [NOT] NULL	IN	AND	OR	NOT
jdbc	Y	Y[4]	Y	N	Y	Y	Y
*:parquet	Y[1]	N	Y[1]	Y[1]	Y[1]	Y[1]	Y[1]
*:orc (all except hive:orc)	Y[1,3]	N	Y[1,3]	Y[1,3]	Y[1,3]	Y[1,3]	Y[1,3]
s3:parquet and s3:text with S3-Select	Y	N	Y	Y	Y	Y	Y
hbase	Y	N	Y	N	Y	Y	N
hive:text	Y[2]	N	N	N	Y[2]	Y[2]	N
hive:rc, hive (accessing stored as RCFile)	Y[2]	N	Y	Y	Y, Y[2]	Y, Y[2]	Y
hive:orc, hive (accessing stored as ORC)	Y, Y[2]	N	Y	Y	Y, Y[2]	Y, Y[2]	Y
hive (accessing stored as Parquet)	Y, Y[2]	N	N	Y	Y, Y[2]	Y, Y[2]	Y
hive:orc and VECTORIZE=true	Y[2]	N	N	N	Y[2]	Y[2]	N

[1] PXF applies the predicate, rather than the remote system, reducing CPU usage and the memory footprint.

[2] PXF supports partition pruning based on partition keys.

[3] PXF filtering is based on file-level, stripe-level, and row-level ORC statistics.

[4] The PXF `jdbc` profile supports the `LIKE` operator only for `TEXT` fields.

PXF does not support filter pushdown for any profile not mentioned in the table above, including: *:avro, *:AvroSequenceFile, *:SequenceFile, *:json, *:text, *:csv, and *:text:multi.

To summarize, all of the following criteria must be met for filter pushdown to occur:

- You enable external table filter pushdown by setting the `gp_external_enable_filter_pushdown` server configuration parameter to `'on'`.
- The Greenplum Database protocol that you use to access external data source must support filter pushdown. The `pxf` external table protocol supports pushdown.
- The external data source that you are accessing must support pushdown. For example, HBase and Hive support pushdown.

- For queries on external tables that you create with the `pxf` protocol, the underlying PXF connector must also support filter pushdown. For example, the PXF Hive, HBase, and JDBC connectors support pushdown, as do the PXF connectors that support reading ORC and Parquet data.
  - Refer to Hive Partition Pruning for more information about Hive support for this feature.

About Column Projection in PXF

PXF supports column projection, and it is always enabled. With column projection, only the columns required by a `SELECT` query on an external table are returned from the external data source. This process can improve query performance, and can also reduce the amount of data that is transferred to Greenplum Database.

Note: Some external data sources do not support column projection. If a query accesses a data source that does not support column projection, the query is instead executed without it, and the data is filtered after it is transferred to Greenplum Database.

Column projection is automatically enabled for the `pxf` external table protocol. PXF accesses external data sources using different connectors, and column projection support is also determined by the specific connector implementation. The following PXF connector and profile combinations support column projection on read operations:

Data Source	Connector	Profile(s)
External SQL database	JDBC Connector	jdbc
Hive	Hive Connector	hive (accessing tables stored as Text, Parquet, RCFile, and ORC), hive:rc, hive:orc
Hadoop	HDFS Connector	hdfs:orc, hdfs:parquet
Network File System	File Connector	file:orc, file:parquet
Amazon S3	S3-Compatible Object Store Connectors	s3:orc, s3:parquet
Amazon S3 using S3 Select	S3-Compatible Object Store Connectors	s3:parquet, s3:text
Google Cloud Storage	GCS Object Store Connector	gs:orc, gs:parquet
Azure Blob Storage	Azure Object Store Connector	wasbs:orc, wasbs:parquet
Azure Data Lake	Azure Object Store Connector	adl:orc, adl:parquet

Note: PXF may disable column projection in cases where it cannot successfully serialize a query filter; for example, when the `WHERE` clause resolves to a `boolean` type.

To summarize, all of the following criteria must be met for column projection to occur:

- The external data source that you are accessing must support column projection. For example, Hive supports column projection for ORC-format data, and certain SQL databases support column projection.

- The underlying PXF connector and profile implementation must also support column projection. For example, the PXF Hive and JDBC connector profiles identified above support column projection, as do the PXF connectors that support reading Parquet data.
- PXF must be able to serialize the query filter.

About the PXF Installation and Configuration Directories

This documentation uses `$PXF_HOME` to refer to the PXF installation directory. Its value depends on how you have installed PXF:

- If you installed PXF as part of Greenplum Database, its value is `$GPHOME/pxf`.
- If you installed the PXF `rpm` or `deb` package, its value is `/usr/local/pxf-gp<greenplum-major-version>`, or the directory of your choosing (CentOS/RHEL only).

`$PXF_HOME` includes both the PXF executables and the PXF runtime configuration files and directories. In PXF 5.x, you needed to specify a `$PXF_CONF` directory for the runtime configuration when you initialized PXF. In PXF 6.x, however, no initialization is required: `$PXF_BASE` now identifies the runtime configuration directory, and the default `$PXF_BASE` is `$PXF_HOME`.

If you want to store your configuration and runtime files in a different location, see Relocating $PXF_BASE.

This documentation uses the `$PXF_HOME` environment variable to reference the PXF installation directory. PXF uses this variable internally at runtime; it is not set in your shell environment, and will display as empty if you attempt to `echo` its value. Similarly, this documentation uses the `$PXF_BASE` environment variable to reference the PXF runtime configuration directory. PXF uses the variable internally. It only needs to be set in your shell environment if you explicitly relocate the directory.

PXF Installation Directories

The following PXF files and directories are installed to `$PXF_HOME` when you install Greenplum Database or the PXF 6.x `rpm` or `deb` package:

Directory	Description
application/	The PXF Server application JAR file.
bin/	The PXF command line executable directory.
commit.sha	The commit identifier for this PXF release.
gpextable/	The PXF extension files. PXF copies the `pxf.control` file from this directory to the Greenplum installation (`$GPHOME`) on a single host when you run the `pxf register` command, or on all hosts in the cluster when you run the `pxf [cluster] register` command from the Greenplum master host.
share/	The directory for shared PXF files that you may require depending on the external data stores that you access. `share/` initially includes only the PXF HBase JAR file.
templates/	The PXF directory for server configuration file templates.
version	The PXF version.

The following PXF directories are installed to `$PXF_BASE` when you install Greenplum Database or the PXF 6.x `rpm` or `deb` package:

Directory	Description
conf/	The location of user-customizable PXF configuration files for PXF runtime and logging configuration settings. This directory contains the `pxf-application.properties`, `pxf-env.sh`, `pxf-log4j2.xml`, and `pxf-profiles.xml` files.
keytabs/	The default location of the PXF Service Kerberos principal keytab file. The `keytabs/` directory and contained files are readable only by the Greenplum Database installation user, typically `gpadmin`.
lib/	The location of user-added runtime dependencies. The `native/` subdirectory is the default PXF runtime directory for native libraries.
logs/	The PXF runtime log file directory. The `logs/` directory and log files are readable only by the Greenplum Database installation user, typically `gpadmin`.
run/	The default PXF run directory. After starting PXF, this directory contains a PXF process id file, `pxf-app.pid`. `run/` and contained files and directories are readable only by the Greenplum Database installation user, typically `gpadmin`.
servers/	The configuration directory for PXF servers; each subdirectory contains a server definition, and the name of the subdirectory identifies the name of the server. The default server is named `default`. The Greenplum Database administrator may configure other servers.

Refer to Configuring PXF and Starting PXF for detailed information about the PXF configuration and startup commands and procedures.

Relocating $PXF_BASE

If you require that `$PXF_BASE` reside in a directory distinct from `$PXF_HOME`, you can change it from the default location to a location of your choosing after you install PXF 6.x.

PXF provides the pxf [cluster] prepare command to prepare a new `$PXF_BASE` location. The command copies the runtime and configuration directories identified above to the file system location that you specify in a `PXF_BASE` environment variable.

For example, to relocate `$PXF_BASE` to the `/path/to/dir` directory on all Greenplum hosts, run the command as follows:

```
gpadmin@gpmaster$ PXF_BASE=/path/to/dir pxf cluster prepare
```

When your `$PXF_BASE` is different than `$PXF_HOME`, inform PXF by setting the `PXF_BASE` environment variable when you run a `pxf` command:

```
gpadmin@gpmaster$ PXF_BASE=/path/to/dir pxf cluster start
```

Set the environment variable in the `.bashrc` shell initialization script for the PXF installation owner (typically the `gpadmin` user) as follows:

```
export PXF_BASE=/path/to/dir
```

About the PXF Configuration Files

`$PXF_BASE/conf` includes these user-customizable configuration files:

- `pxf-application.properties` - PXF Service application configuration properties
- `pxf-env.sh` - PXF command and JVM-specific runtime configuration properties
- `pxf-log4j2.xml` - PXF logging configuration properties
- `pxf-profiles.xml` - Custom PXF profile definitions

pxf-application.properties

The `pxf-application.properties` file exposes these PXF Service application configuration properties:

Parameter	Description	Default Value
pxf.connection.timeout	The Tomcat server connection timeout for read operations (-1 for infinite timeout).	5m (5 minutes)
pxf.connection.upload-timeout	The Tomcat server connection timeout for write operations (-1 for infinite timeout).	5m (5 minutes)
pxf.max.threads	The maximum number of PXF tomcat threads.	200
pxf.task.pool.allow-core-thread-timeout	Identifies whether or not core streaming threads are allowed to time out.	false
pxf.task.pool.core-size	The number of core streaming threads.	8
pxf.task.pool.queue-capacity	The capacity of the core streaming thread pool queue.	0
pxf.task.pool.max-size	The maximum allowed number of core streaming threads.	pxf.max.threads if set, or 200
pxf.log.level	The log level for the PXF Service.	info
pxf.fragmenter-cache.expiration	The amount of time after which an entry expires and is removed from the fragment cache.	10s (10 seconds)

To change the value of a PXF Service application property, you may first need to add the property to, or uncomment the property in, the `pxf-application.properties` file before you can set the new value.

pxf-env.sh

The `pxf-env.sh` file exposes these PXF JVM configuration properties:

Parameter	Description	Default Value
JAVA_HOME	The path to the Java JRE home directory.	/usr/java/default
PXF_LOGDIR	The PXF log directory.	$PXF_BASE/logs
PXF_RUNDIR	The PXF run directory.	$PXF_BASE/run
PXF_JVM_OPTS	The default options for the PXF Java virtual machine.	-Xmx2g -Xms1g
PXF_OOM_KILL	Enable/disable PXF auto-kill on OutOfMemoryError (OOM).	true (enabled)

Parameter	Description	Default Value
PXF_OOM_DUMP_PATH	The absolute path to the dump file that PXF generates on OOM.	No dump file (empty)
PXF_LOADER_PATH	Additional directories and JARs for PXF to class-load.	(empty)
LD_LIBRARY_PATH	Additional directories and native libraries for PXF to load.	(empty)

To set a new value for a PXF JVM configuration property, you may first need to uncomment the property in the `pxf-env.sh` file before you set the new value.

pxf-log4j2.xml

The `pxf-log4j2.xml` file configures PXF and subcomponent logging. By default, PXF is configured to log at the `info` level, and logs at the `warn` or `error` levels for some third-party libraries to reduce verbosity.

The Logging advanced configuration topic describes how to enable more verbose client- and server-level logging for PXF.

pxf-profiles.xml

PXF defines its default profiles in the `pxf-profiles-default.xml` file. If you choose to add a custom profile, you configure the profile in `pxf-profiles.xml`.

Modifying the PXF Configuration

When you update a PXF configuration file, you must sychronize the changes to all hosts in the Greenplum Database cluster and then restart PXF for the changes to take effect.

Procedure:

1. Update the configuration file(s) of interest.
2. Synchronize the PXF configuration to all hosts in the Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

3. (Re)start PXF on all Greenplum hosts:

   ```
   gpadmin@gpmaster$ pxf cluster restart
   ```

Configuring PXF

Your Greenplum Database deployment consists of a master node, standby master, and multiple segment hosts. After you configure the Greenplum Platform Extension Framework (PXF), you start a single PXF JVM process (PXF Service) on each Greenplum Database host.

PXF provides connectors to Hadoop, Hive, HBase, object stores, network file systems, and external SQL data stores. You must configure PXF to support the connectors that you plan to use.

To configure PXF, you must:

1. Install Java 8 or 11 on each Greenplum Database host as described in Installing Java for PXF.

If your `JAVA_HOME` is different to `/usr/java/default`, you must inform PXF of the $JAVA_HOME setting by specifying its value in the `pxf-env.sh` configuration file.

- Edit the `$PXF_BASE/conf/pxf-env.sh` file on the Greenplum master node.

  ```
  gpadmin@gpmaster$ vi /usr/local/pxf-gp6/conf/pxf-env.sh
  ```

- Locate the `JAVA_HOME` setting in the `pxf-env.sh` file, uncomment if necessary, and set it to your `$JAVA_HOME` value. For example:

  ```
  export JAVA_HOME=/usr/lib/jvm/java-1.8.0-openjdk/jre/
  ```

2. Register the PXF extension with Greenplum Database (see pxf cluster register). Run this command after your first installation of a PXF version 6.x, and/or after you upgrade your Greenplum Database installation:

   ```
   gpadmin@gpmaster$ pxf cluster register
   ```

3. If you plan to use the Hadoop, Hive, or HBase PXF connectors, you must perform the configuration procedure described in Configuring PXF Hadoop Connectors.
4. If you plan to use the PXF connectors to access the Azure, Google Cloud Storage, Minio, or S3 object store(s), you must perform the configuration procedure described in Configuring Connectors to Azure, Google Cloud Storage, Minio, and S3 Object Stores.
5. If you plan to use the PXF JDBC Connector to access an external SQL database, perform the configuration procedure described in Configuring the JDBC Connector.
6. If you plan to use PXF to access a network file system, perform the configuration procedure described in Configuring a PXF Network File System Server.
7. After making any configuration changes, synchronize the PXF configuration to all hosts in the cluster.

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

8. After synchronizing PXF configuration changes, Start PXF.
9. Enable the PXF extension and grant access to users.

Configuring PXF Servers

This topic provides an overview of PXF server configuration. To configure a server, refer to the topic specific to the connector that you want to configure.

You read from or write data to an external data store via a PXF connector. To access an external data store, you must provide the server location. You may also be required to provide client access credentials and other external data store-specific properties. PXF simplifies configuring access to external data stores by:

- Supporting file-based connector and user configuration
- Providing connector-specific template configuration files

A PXF *Server* definition is simply a named configuration that provides access to a specific external data store. A PXF server name is the name of a directory residing in `$PXF_BASE/servers/`. The

information that you provide in a server configuration is connector-specific. For example, a PXF JDBC Connector server definition may include settings for the JDBC driver class name, URL, username, and password. You can also configure connection-specific and session-specific properties in a JDBC server definition.

PXF provides a server template file for each connector; this template identifies the typical set of properties that you must configure to use the connector.

You will configure a server definition for each external data store that Greenplum Database users need to access. For example, if you require access to two Hadoop clusters, you will create a PXF Hadoop server configuration for each cluster. If you require access to an Oracle and a MySQL database, you will create one or more PXF JDBC server configurations for each database.

A server configuration may include default settings for user access credentials and other properties for the external data store. You can allow Greenplum Database users to access the external data store using the default settings, or you can configure access and other properties on a per-user basis. This allows you to configure different Greenplum Database users with different external data store access credentials in a single PXF server definition.

About Server Template Files

The configuration information for a PXF server resides in one or more `<connector>-site.xml` files in `$PXF_BASE/servers/<server_name>/`.

PXF provides a template configuration file for each connector. These server template configuration files are located in the `$PXF_HOME/templates/` directory after you install PXF:

```
gpadmin@gpmaster$ ls $PXF_HOME/templates
adl-site.xml   hbase-site.xml  jdbc-site.xml    pxf-site.xml    yarn-site.xml
core-site.xml  hdfs-site.xml   mapred-site.xml  s3-site.xml
gs-site.xml    hive-site.xml   minio-site.xml   wasbs-site.xml
```

For example, the contents of the `s3-site.xml` template file follow:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
    <property>
        <name>fs.s3a.access.key</name>
        <value>YOUR_AWS_ACCESS_KEY_ID</value>
    </property>
    <property>
        <name>fs.s3a.secret.key</name>
        <value>YOUR_AWS_SECRET_ACCESS_KEY</value>
    </property>
    <property>
        <name>fs.s3a.fast.upload</name>
        <value>true</value>
    </property>
</configuration>
```

You specify credentials to PXF in clear text in configuration files.

Note: The template files for the Hadoop connectors are not intended to be modified and used for configuration, as they only provide an example of the information needed. Instead of modifying the Hadoop templates, you will copy several Hadoop `*-site.xml` files from the Hadoop cluster to your

PXF Hadoop server configuration.

About the Default Server

PXF defines a special server named `default`. The PXF installation creates a `$PXF_BASE/servers/default/` directory. This directory, initially empty, identifies the default PXF server configuration. You can configure and assign the default PXF server to any external data source. For example, you can assign the PXF default server to a Hadoop cluster, or to a MySQL database that your users frequently access.

PXF automatically uses the `default` server configuration if you omit the `SERVER=<server_name>` setting in the `CREATE EXTERNAL TABLE` command `LOCATION` clause.

Configuring a Server

When you configure a PXF connector to an external data store, you add a named PXF server configuration for the connector. Among the tasks that you perform, you may:

1. Determine if you are configuring the `default` PXF server, or choose a new name for the server configuration.
2. Create the directory `$PXF_BASE/servers/<server_name>`.
3. Copy template or other configuration files to the new server directory.
4. Fill in appropriate default values for the properties in the template file.
5. Add any additional configuration properties and values required for your environment.
6. Configure one or more users for the server configuration as described in About Configuring a PXF User.
7. Synchronize the server and user configuration to the Greenplum Database cluster.

Note: You must re-sync the PXF configuration to the Greenplum Database cluster after you add or update PXF server configuration.

After you configure a PXF server, you publish the server name to Greenplum Database users who need access to the data store. A user only needs to provide the server name when they create an external table that accesses the external data store. PXF obtains the external data source location and access credentials from server and user configuration files residing in the server configuration directory identified by the server name.

To configure a PXF server, refer to the connector configuration topic:

- To configure a PXF server for Hadoop, refer to Configuring PXF Hadoop Connectors .
- To configure a PXF server for an object store, refer to Configuring Connectors to Minio and S3 Object Stores and Configuring Connectors to Azure and Google Cloud Storage Object Stores.
- To configure a PXF JDBC server, refer to Configuring the JDBC Connector .
- Configuring a PXF Network File System Server describes the process of configuring a PXF server for network file system access.

About the pxf-site.xml Configuration File

PXF includes a template file named `pxf-site.xml` for PXF-specific configuration parameters. You can use the `pxf-site.xml` template file to specify Kerberos and/or user impersonation settings for server configurations, or to specify a base directory for file access.

The Kerberos and user impersonation settings in this file apply only to Hadoop and JDBC server configurations; they do not apply to file system or object store server configurations.

You configure properties in the `pxf-site.xml` file for a PXF server when one or more of the following conditions hold:

- The remote Hadoop system utilizes Kerberos authentication.
- You want to enable/disable user impersonation on the remote Hadoop or external database system.
- You will access a network file system with the server configuration.
- You will access a remote Hadoop or object store file system with the server configuration, and you want to allow a user to access only a specific directory and subdirectories.

`pxf-site.xml` includes the following properties:

Property	Description	Default Value
pxf.service.kerberos.principal	The Kerberos principal name.	gpadmin/_HOST@EXAMPLE.COM
pxf.service.kerberos.keytab	The file system path to the Kerberos keytab file.	$PXF_BASE/keytabs/pxf.service.keytab
pxf.service.user.impersonation	Enables/disables user impersonation on the remote system.	If the `pxf.service.user.impersonation` property is missing from `pxf-site.xml`, the default is `true` (enabled) for PXF Hadoop servers and `false` (disabled) for JDBC servers.
pxf.service.user.name	The log in user for the remote system.	This property is commented out by default. When the property is unset, the default value is the operating system user that starts the pxf process, typically `gpadmin`. When the property is set, the default value depends on the user impersonation setting and, if you are accessing Hadoop, whether or not you are accessing a Kerberos-secured cluster; see the Use Cases and Configuration Scenarios section in the *Configuring the Hadoop User, User Impersonation, and Proxying* topic.
pxf.fs.basePath	Identifies the base path or share point on the remote file system. This property is applicable when the server configuration is used with a profile that accesses a file.	None; this property is commented out by default.

Property	Description	Default Value
pxf.ppd.hive[1]	Specifies whether or not predicate pushdown is enabled for queries on external tables that specify the `hive`, `hive:rc`, or `hive:orc` profiles.	True; predicate pushdown is enabled.
pxf.sasl.connection.retries	Specifies the maximum number of times that PXF retries a SASL connection request after a refused connection returns a `GSS initiate failed` error.	5

[1] Should you need to, you can override this setting on a per-table basis by specifying the `&PPD=<boolean>` option in the `LOCATION` clause when you create the external table.

Refer to Configuring PXF Hadoop Connectors and Configuring the JDBC Connector for information about relevant `pxf-site.xml` property settings for Hadoop and JDBC server configurations, respectively. See Configuring a PXF Network File System Server for information about relevant `pxf-site.xml` property settings when you configure a PXF server to access a network file system.

About the pxf.fs.basePath Property

You can use the `pxf.fs.basePath` property to restrict a user's access to files in a specific remote directory. When set, this property applies to any profile that accesses a file, including `*:text`, `*:parquet`, `*:json`, etc.

When you configure the `pxf.fs.basePath` property for a server, PXF considers the file path specified in the `CREATE EXTERNAL TABLE LOCATION` clause to be relative to this base path setting, and constructs the remote path accordingly.

You must set `pxf.fs.basePath` when you configure a PXF server for access to a network file system with a `file:*` profile. This property is optional for a PXF server that accesses a file in Hadoop or in an object store.

Configuring a PXF User

You can configure access to an external data store on a per-server, per-Greenplum-user basis.

PXF per-server, per-user configuration provides the most benefit for JDBC servers.

You configure external data store user access credentials and properties for a specific Greenplum Database user by providing a `<greenplum_user_name>-user.xml` user configuration file in the PXF server configuration directory, `$PXF_BASE/servers/<server_name>/`. For example, you specify the properties for the Greenplum Database user named `bill` in the file `$PXF_BASE/servers/<server_name>/bill-user.xml`. You can configure zero, one, or more users in a

PXF server configuration.

The properties that you specify in a user configuration file are connector-specific. You can specify any configuration property supported by the PXF connector server in a `<greenplum_user_name>-user.xml` configuration file.

For example, suppose you have configured access to a PostgreSQL database in the PXF JDBC server configuration named `pgsrv1`. To allow the Greenplum Database user named `bill` to access this database as the PostgreSQL user named `pguser1`, password `changeme`, you create the user configuration file `$PXF_BASE/servers/pgsrv1/bill-user.xml` with the following properties:

```
<configuration>
    <property>
        <name>jdbc.user</name>
        <value>pguser1</value>
    </property>
    <property>
        <name>jdbc.password</name>
        <value>changeme</value>
    </property>
</configuration>
```

If you want to configure a specific search path and a larger read fetch size for `bill`, you would also add the following properties to the `bill-user.xml` user configuration file:

```
    <property>
        <name>jdbc.session.property.search_path</name>
        <value>bill_schema</value>
    </property>
    <property>
        <name>jdbc.statement.fetchSize</name>
        <value>2000</value>
    </property>
```

Procedure

For each PXF user that you want to configure, you will:

1. Identify the name of the Greenplum Database user.
2. Identify the PXF server definition for which you want to configure user access.
3. Identify the name and value of each property that you want to configure for the user.
4. Create/edit the file `$PXF_BASE/servers/<server_name>/<greenplum_user_name>-user.xml`, and add the outer configuration block:

   ```
   <configuration>
   </configuration>
   ```

5. Add each property/value pair that you identified in Step 3 within the configuration block in the `<greenplum_user_name>-user.xml` file.
6. If you are adding the PXF user configuration to previously configured PXF server definition, synchronize the user configuration to the Greenplum Database cluster.

About Configuration Property Precedence

A PXF server configuration may include default settings for user access credentials and other properties for accessing an external data store. Some PXF connectors, such as the S3 and JDBC connectors, allow you to directly specify certain server properties via custom options in the `CREATE EXTERNAL TABLE` command `LOCATION` clause. A `<greenplum_user_name>-user.xml` file specifies property settings for an external data store that are specific to a Greenplum Database user.

For a given Greenplum Database user, PXF uses the following precedence rules (highest to lowest) to obtain configuration property settings for the user:

1. A property that you configure in `<server_name>/<greenplum_user_name>-user.xml` overrides any setting of the property elsewhere.
2. A property that is specified via custom options in the `CREATE EXTERNAL TABLE` command `LOCATION` clause overrides any setting of the property in a PXF server configuration.
3. Properties that you configure in the `<server_name>` PXF server definition identify the default property values.

These precedence rules allow you create a single external table that can be accessed by multiple Greenplum Database users, each with their own unique external data store user credentials.

Using a Server Configuration

To access an external data store, the Greenplum Database user specifies the server name in the `CREATE EXTERNAL TABLE` command `LOCATION` clause `SERVER=<server_name>` option. The `<server_name>` that the user provides identifies the server configuration directory from which PXF obtains the configuration and credentials to access the external data store.

For example, the following command accesses an S3 object store using the server configuration defined in the `$PXF_BASE/servers/s3srvcfg/s3-site.xml` file:

```
CREATE EXTERNAL TABLE pxf_ext_tbl(name text, orders int)
  LOCATION ('pxf://BUCKET/dir/file.txt?PROFILE=s3:text&SERVER=s3srvcfg')
FORMAT 'TEXT' (delimiter=E',');
```

PXF automatically uses the `default` server configuration when no `SERVER=<server_name>` setting is provided.

For example, if the `default` server configuration identifies a Hadoop cluster, the following example command references the HDFS file located at `/path/to/file.txt`:

```
CREATE EXTERNAL TABLE pxf_ext_hdfs(location text, miles int)
  LOCATION ('pxf://path/to/file.txt?PROFILE=hdfs:text')
FORMAT 'TEXT' (delimiter=E',');
```

A Greenplum Database user who queries or writes to an external table accesses the external data store with the credentials configured for the `<server_name>` user. If no user-specific credentials are configured for `<server_name>`, the Greenplum user accesses the external data store with the default credentials configured for `<server_name>`.

Configuring PXF Hadoop Connectors (Optional)

PXF is compatible with Cloudera, Hortonworks Data Platform, MapR, and generic Apache Hadoop distributions. This topic describes how configure the PXF Hadoop, Hive, and HBase connectors.

If you do not want to use the Hadoop-related PXF connectors, then you do not need to perform this procedure.

Prerequisites

Configuring PXF Hadoop connectors involves copying configuration files from your Hadoop cluster to the Greenplum Database master host. If you are using the MapR Hadoop distribution, you must also copy certain JAR files to the master host. Before you configure the PXF Hadoop connectors, ensure that you can copy files from hosts in your Hadoop cluster to the Greenplum Database master.

Procedure

Perform the following procedure to configure the desired PXF Hadoop-related connectors on the Greenplum Database master host. After you configure the connectors, you will use the `pxf cluster sync` command to copy the PXF configuration to the Greenplum Database cluster.

In this procedure, you use the `default`, or create a new, PXF server configuration. You copy Hadoop configuration files to the server configuration directory on the Greenplum Database master host. You identify Kerberos and user impersonation settings required for access, if applicable. You may also copy libraries to `$PXF_BASE/lib` for MapR support. You then synchronize the PXF configuration on the master host to the standby master and segment hosts.

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Identify the name of your PXF Hadoop server configuration.

3. If you are not using the `default` PXF server, create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a Hadoop server configuration named `hdp3`:

   ```
   gpadmin@gpmaster$ mkdir $PXF_BASE/servers/hdp3
   ```

4. Change to the server directory. For example:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/default
   ```

   Or,

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/hdp3
   ```

5. PXF requires information from `core-site.xml` and other Hadoop configuration files. Copy the `core-site.xml`, `hdfs-site.xml`, `mapred-site.xml`, and `yarn-site.xml` Hadoop configuration files from your Hadoop cluster NameNode host to the current host using your tool of choice. Your file paths may differ based on the Hadoop distribution in use. For example, these commands use `scp` to copy the files:

```
gpadmin@gpmaster$ scp hdfsuser@namenode:/etc/hadoop/conf/core-site.xml .
gpadmin@gpmaster$ scp hdfsuser@namenode:/etc/hadoop/conf/hdfs-site.xml .
gpadmin@gpmaster$ scp hdfsuser@namenode:/etc/hadoop/conf/mapred-site.xml .
gpadmin@gpmaster$ scp hdfsuser@namenode:/etc/hadoop/conf/yarn-site.xml .
```

6. If you plan to use the PXF Hive connector to access Hive table data, similarly copy the Hive configuration to the Greenplum Database master host. For example:

```
gpadmin@gpmaster$ scp hiveuser@hivehost:/etc/hive/conf/hive-site.xml .
```

7. If you plan to use the PXF HBase connector to access HBase table data, similarly copy the HBase configuration to the Greenplum Database master host. For example:

```
gpadmin@gpmaster$ scp hbaseuser@hbasehost:/etc/hbase/conf/hbase-site.xml .
```

8. If you are using PXF with the MapR Hadoop distribution, you must copy certain JAR files from your MapR cluster to the Greenplum Database master host. (Your file paths may differ based on the version of MapR in use.) For example, these commands use `scp` to copy the files:

```
gpadmin@gpmaster$ cd $PXF_BASE/lib
gpadmin@gpmaster$ scp mapruser@maprhost:/opt/mapr/hadoop/hadoop-2.7.0/share/had
oop/common/lib/maprfs-5.2.2-mapr.jar .
gpadmin@gpmaster$ scp mapruser@maprhost:/opt/mapr/hadoop/hadoop-2.7.0/share/had
oop/common/lib/hadoop-auth-2.7.0-mapr-1707.jar .
gpadmin@gpmaster$ scp mapruser@maprhost:/opt/mapr/hadoop/hadoop-2.7.0/share/had
oop/common/hadoop-common-2.7.0-mapr-1707.jar .
```

9. Synchronize the PXF configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

10. PXF accesses Hadoop services on behalf of Greenplum Database end users. By default, PXF tries to access HDFS, Hive, and HBase using the identity of the Greenplum Database user account that logs into Greenplum Database. In order to support this functionality, you must configure proxy settings for Hadoop, as well as for Hive and HBase if you intend to use those PXF connectors. Follow procedures in Configuring User Impersonation and Proxying to configure user impersonation and proxying for Hadoop services, or to turn off PXF user impersonation.

11. Grant read permission to the HDFS files and directories that will be accessed as external tables in Greenplum Database. If user impersonation is enabled (the default), you must grant this permission to each Greenplum Database user/role name that will use external tables that reference the HDFS files. If user impersonation is not enabled, you must grant this permission to the `gpadmin` user.

12. If your Hadoop cluster is secured with Kerberos, you must configure PXF and generate Kerberos principals and keytabs for each Greenplum Database host as described in Configuring PXF for Secure HDFS.

About Updating the Hadoop Configuration

If you update your Hadoop, Hive, or HBase configuration while the PXF Service is running, you must

copy the updated configuration to the `$PXF_BASE/servers/<server_name>` directory and re-sync the PXF configuration to your Greenplum Database cluster. For example:

```
gpadmin@gpmaster$ cd $PXF_BASE/servers/<server_name>
gpadmin@gpmaster$ scp hiveuser@hivehost:/etc/hive/conf/hive-site.xml .
gpadmin@gpmaster$ pxf cluster sync
```

Configuring the Hadoop User, User Impersonation, and Proxying

PXF accesses Hadoop services on behalf of Greenplum Database end users.

When user impersonation is enabled (the default), PXF accesses Hadoop services using the identity of the Greenplum Database user account that logs in to Greenplum and performs an operation that uses a PXF connector. Keep in mind that PXF uses only the *login* identity of the user when accessing Hadoop services. For example, if a user logs in to Greenplum Database as the user `jane` and then execute `SET ROLE` or `SET SESSION AUTHORIZATION` to assume a different user identity, all PXF requests still use the identity `jane` to access Hadoop services. When user impersonation is enabled, you must explicitly configure each Hadoop data source (HDFS, Hive, HBase) to allow PXF to act as a proxy for impersonating specific Hadoop users or groups.

When user impersonation is disabled, PXF executes all Hadoop service requests as the PXF process owner (usually `gpadmin`) or the Hadoop user identity that you specify. This behavior provides no means to control access to Hadoop services for different Greenplum Database users. It requires that this user have access to all files and directories in HDFS, and all tables in Hive and HBase that are referenced in PXF external table definitions.

You configure the Hadoop user and PXF user impersonation setting for a server via the `pxf-site.xml` server configuration file. Refer to About the pxf-site.xml Configuration File for more information about the configuration properties in this file.

Use Cases and Configuration Scenarios

User, user impersonation, and proxy configuration for Hadoop depends on how you use PXF to access Hadoop, and whether or not the Hadoop cluster is secured with Kerberos.

The following scenarios describe the use cases and configuration required when you use PXF to access *non-secured Hadoop*. If you are using PXF to access a *Kerberos-secured Hadoop cluster*, refer to the Use Cases and Configuration Scenarios section in the *Configuring PXF for Secure HDFS* topic.

Note: These scenarios assume that `gpadmin` is the PXF process owner.

Accessing Hadoop as the Greenplum User Proxied by gpadmin

This is the default configuration for PXF. The `gpadmin` user proxies Greenplum queries on behalf of Greenplum users. The effective user in Hadoop is the Greenplum user that runs the query.

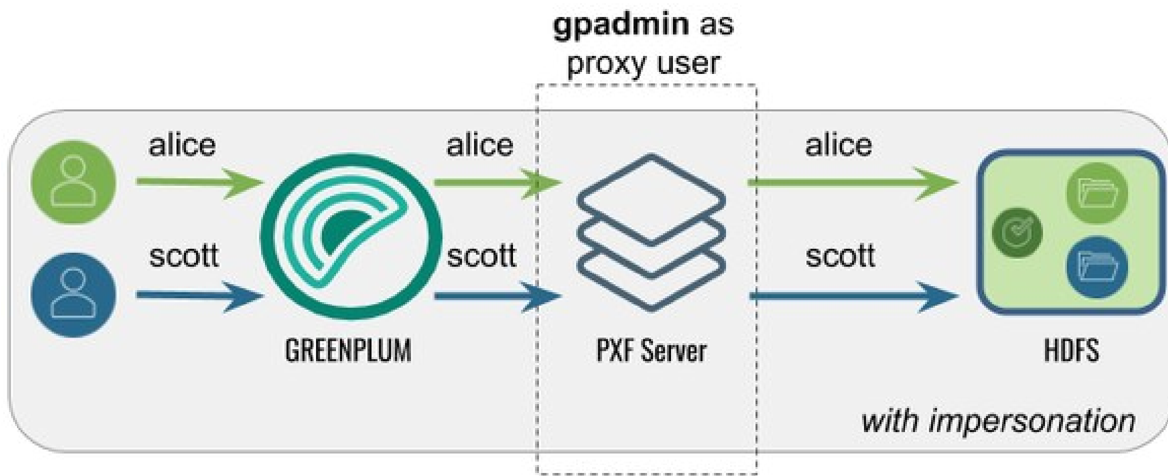

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
true	`gpadmin`	None; this is the default configuration.	Set the `gpadmin` user as the Hadoop proxy user as described in Configure Hadoop Proxying.

Accessing Hadoop as the Greenplum User Proxied by a <custom> User

In this configuration, PXF accesses Hadoop as the Greenplum user proxied by <custom> user. A query initiated by a Greenplum user appears on the Hadoop side as originating from the (<custom> user.

This configuration might be desirable when Hadoop is already configured with a proxy user, or when you want a user different than `gpadmin` to proxy Greenplum queries.

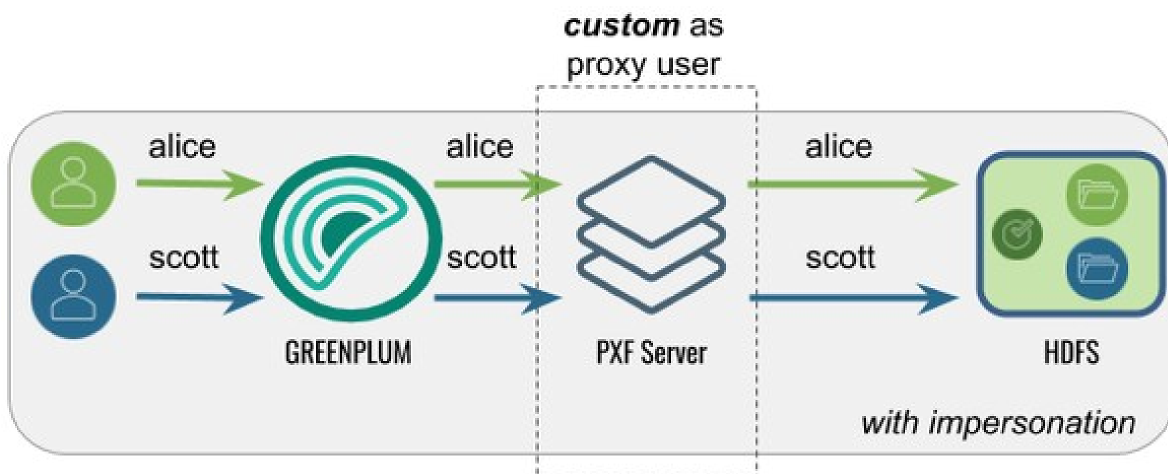

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
true	<custom>	Configure the Hadoop User to the <custom> user name.	Set the <custom> user as the Hadoop proxy user as described in Configure Hadoop Proxying.

Accessing Hadoop as the gpadmin User

In this configuration, PXF accesses Hadoop as the `gpadmin` user. A query initiated by any Greenplum user appears on the Hadoop side as originating from the `gpadmin` user.

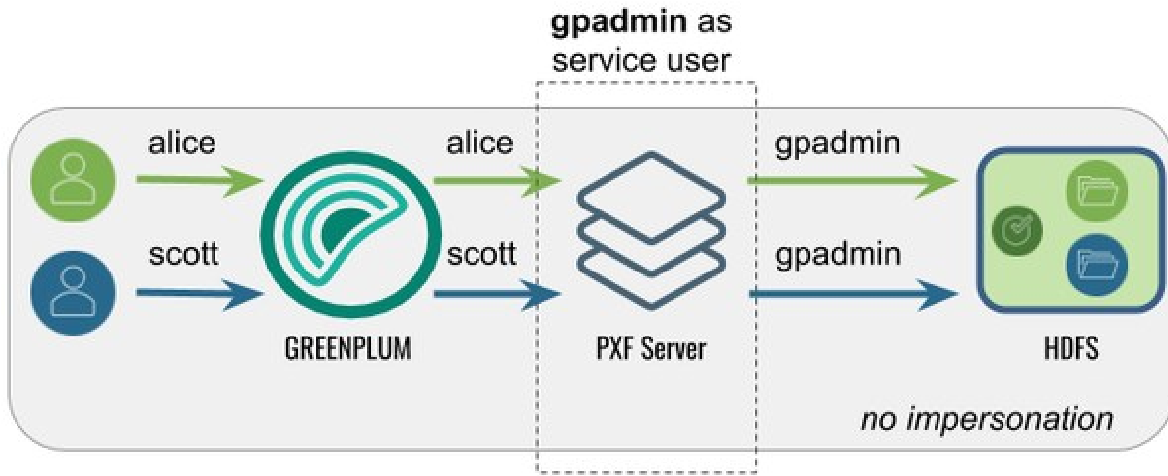

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
false	`gpadmin`	Turn off user impersonation as described in Configure PXF User Impersonation.	None required.

Accessing Hadoop as a <custom> User

In this configuration, PXF accesses Hadoop as a <custom> user. A query initiated by any Greenplum user appears on the Hadoop side as originating from the <custom> user.

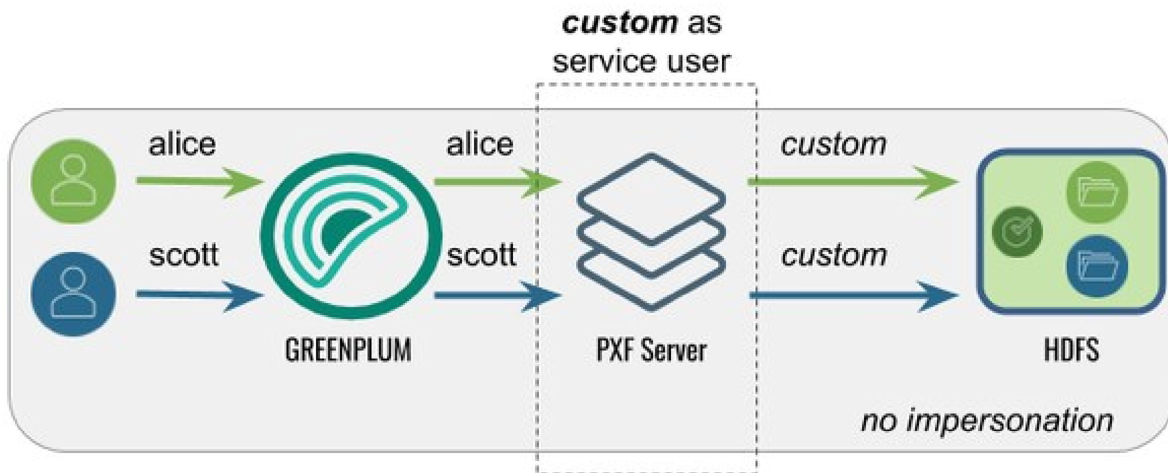

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
false	<custom>	Turn off user impersonation as described in Configure PXF User Impersonation and Configure the Hadoop User to the <custom> user name.	None required.

Configure the Hadoop User

By default, PXF accesses Hadoop using the identity of the Greenplum Database user, and you are required to set up a proxy Hadoop user. You can configure PXF to access Hadoop as a different user on a per-server basis.

Perform the following procedure to configure the Hadoop user:

1. Log in to your Greenplum Database master node as the administrative user:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Identify the name of the PXF Hadoop server configuration that you want to update.
3. Navigate to the server configuration directory. For example, if the server is named `hdp3`:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/hdp3
   ```

4. If the server configuration does not yet include a `pxf-site.xml` file, copy the template file to the directory. For example:

   ```
   gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml .
   ```

5. Open the `pxf-site.xml` file in the editor of your choice, and configure the Hadoop user name. When impersonation is disabled, this name identifies the Hadoop user identity that PXF will use to access the Hadoop system. When user impersonation is enabled, this name identifies the PXF proxy Hadoop user. For example, if you want to access Hadoop as the user `hdfsuser1`, uncomment the property and set it as follows:

   ```
   <property>
       <name>pxf.service.user.name</name>
       <value>hdfsuser1</value>
   </property>
   ```

6. Save the `pxf-site.xml` file and exit the editor.
7. Use the `pxf cluster sync` command to synchronize the PXF Hadoop server configuration to your Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

Configure PXF User Impersonation

PXF user impersonation is enabled by default for Hadoop servers. You can configure PXF user impersonation on a per-server basis. Perform the following procedure to turn PXF user impersonation on or off for the Hadoop server configuration:

1. Navigate to the server configuration directory. For example, if the server is named `hdp3`:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/hdp3
   ```

2. If the server configuration does not yet include a `pxf-site.xml` file, copy the template file to the directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml .
```

3. Open the `pxf-site.xml` file in the editor of your choice, and update the user impersonation property setting. For example, if you do not require user impersonation for this server configuration, set the `pxf.service.user.impersonation` property to `false`:

```
<property>
    <name>pxf.service.user.impersonation</name>
    <value>false</value>
</property>
```

   If you require user impersonation, turn it on:

```
<property>
    <name>pxf.service.user.impersonation</name>
    <value>true</value>
</property>
```

4. If you enabled user impersonation, you must configure Hadoop proxying as described in Configure Hadoop Proxying. You must also configure Hive User Impersonation and HBase User Impersonation if you plan to use those services.
5. Save the `pxf-site.xml` file and exit the editor.
6. Use the `pxf cluster sync` command to synchronize the PXF Hadoop server configuration to your Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Configure Hadoop Proxying

When PXF user impersonation is enabled for a Hadoop server configuration, you must configure Hadoop to permit PXF to proxy Greenplum users. This configuration involves setting certain `hadoop.proxyuser.*` properties. Follow these steps to set up PXF Hadoop proxy users:

1. Log in to your Hadoop cluster and open the `core-site.xml` configuration file using a text editor, or use Ambari or another Hadoop cluster manager to add or edit the Hadoop property values described in this procedure.
2. Set the property `hadoop.proxyuser.<name>.hosts` to specify the list of PXF host names from which proxy requests are permitted. Substitute the PXF proxy Hadoop user for `<name>`. The PXF proxy Hadoop user is the `pxf.service.user.name` that you configured in the procedure above, or, if you are using Kerberos authentication to Hadoop, the proxy user identity is the *primary* component of the Kerberos principal. If you have not explicitly configured `pxf.service.user.name`, the proxy user is the operating system user that started PXF. Provide multiple PXF host names in a comma-separated list. For example, if the PXF proxy user is named `hdfsuser2`:

```
<property>
    <name>hadoop.proxyuser.hdfsuser2.hosts</name>
    <value>pxfhost1,pxfhost2,pxfhost3</value>
</property>
```

3. Set the property `hadoop.proxyuser.<name>.groups` to specify the list of HDFS groups that PXF as Hadoop user `<name>` can impersonate. You should limit this list to only those groups that require access to HDFS data from PXF. For example:

```
<property>
    <name>hadoop.proxyuser.hdfsuser2.groups</name>
    <value>group1,group2</value>
</property>
```

4. You must restart Hadoop for your `core-site.xml` changes to take effect.

5. Copy the updated `core-site.xml` file to the PXF Hadoop server configuration directory `$PXF_BASE/servers/<server_name>` on the Greenplum Database master and synchronize the configuration to the standby master and each Greenplum Database segment host.

Hive User Impersonation

The PXF Hive connector uses the Hive MetaStore to determine the HDFS locations of Hive tables, and then accesses the underlying HDFS files directly. No specific impersonation configuration is required for Hive, because the Hadoop proxy configuration in `core-site.xml` also applies to Hive tables accessed in this manner.

HBase User Impersonation

In order for user impersonation to work with HBase, you must enable the `AccessController` coprocessor in the HBase configuration and restart the cluster. See 61.3 Server-side Configuration for Simple User Access Operation in the Apache HBase Reference Guide for the required `hbase-site.xml` configuration settings.

Configuring PXF for Secure HDFS

When Kerberos is enabled for your HDFS filesystem, the PXF Service, as an HDFS client, requires a principal and keytab file to authenticate access to HDFS. To read or write files on a secure HDFS, you must create and deploy Kerberos principals and keytabs for PXF, and ensure that Kerberos authentication is enabled and functioning.

PXF supports simultaneous access to multiple Kerberos-secured Hadoop clusters.

When Kerberos is enabled, you access Hadoop with the PXF principal and keytab. You can also choose to access Hadoop using the identity of the Greenplum Database user.

You configure the impersonation setting and the Kerberos principal and keytab for a Hadoop server via the `pxf-site.xml` server-specific configuration file. Refer to About the pxf-site.xml Configuration File for more information about the configuration properties in this file.

Configure the Kerberos principal and keytab using the following `pxf-site.xml` properties:

Property	Description	Default Value
pxf.service.kerberos.principal	The Kerberos principal name.	gpadmin/_HOST@EXAMPLE.COM
pxf.service.kerberos.keytab	The file system path to the Kerberos keytab file.	$PXF_BASE/keytabs/pxf.service.keytab

Prerequisites

Before you configure PXF for access to a secure HDFS filesystem, ensure that you have:

- Configured a PXF server for the Hadoop cluster, and can identify the server configuration name.
- Configured and started PXF as described in Configuring PXF.
- Verified that Kerberos is enabled for your Hadoop cluster.
- Verified that the HDFS configuration parameter `dfs.block.access.token.enable` is set to `true`. You can find this setting in the `hdfs-site.xml` configuration file on a host in your Hadoop cluster.
- Noted the host name or IP address of each Greenplum Database host (<gphost>) and the Kerberos Key Distribution Center (KDC) <kdc-server> host.
- Noted the name of the Kerberos <realm> in which your cluster resides.
- Installed the Kerberos client packages on **each** Greenplum Database host if they are not already installed. You must have superuser permissions to install operating system packages. For example:

  ```
  root@gphost$ rpm -qa | grep krb
  root@gphost$ yum install krb5-libs krb5-workstation
  ```

Use Cases and Configuration Scenarios

The following scenarios describe the use cases and configuration required when you use PXF to access a Kerberos-secured Hadoop cluster.

Note: These scenarios assume that `gpadmin` is the PXF process owner.

Accessing Hadoop as the Greenplum User Proxied by the Kerberos Principal

In this configuration, PXF accesses Hadoop as the Greenplum user proxied by the Kerberos principal. The Kerberos principal is the Hadoop proxy user and accesses Hadoop as the Greenplum user.

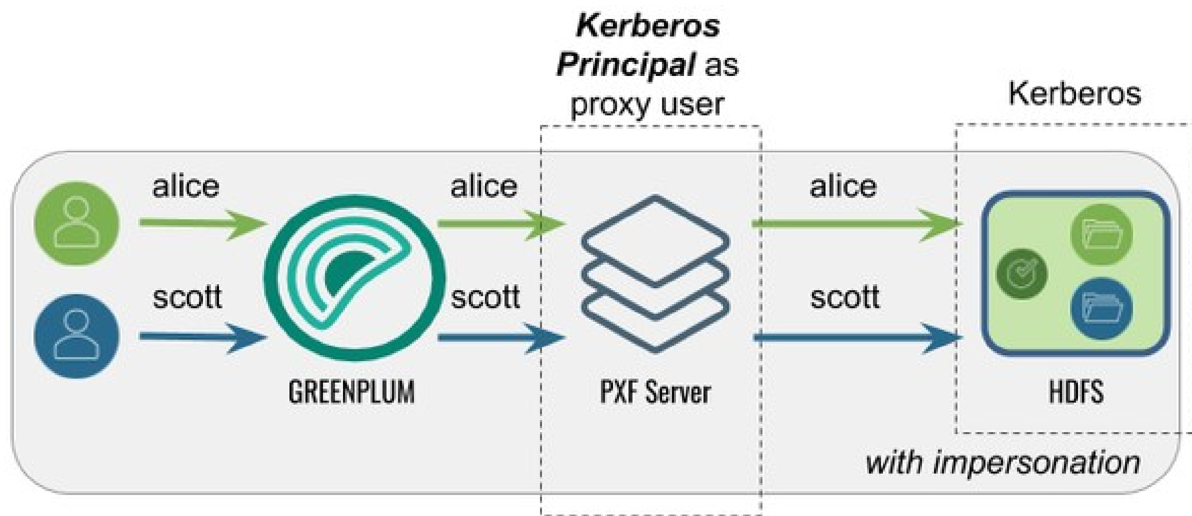

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
true	Kerberos principal	Perform the Configuration Procedure in this topic.	Set the Kerberos principal as the Hadoop proxy user as described in Configure Hadoop Proxying.

Accessing Hadoop as the Kerberos Principal

In this configuration, PXF accesses Hadoop as the Kerberos principal. A query initiated by any Greenplum user appears on the Hadoop side as originating from the Kerberos principal.

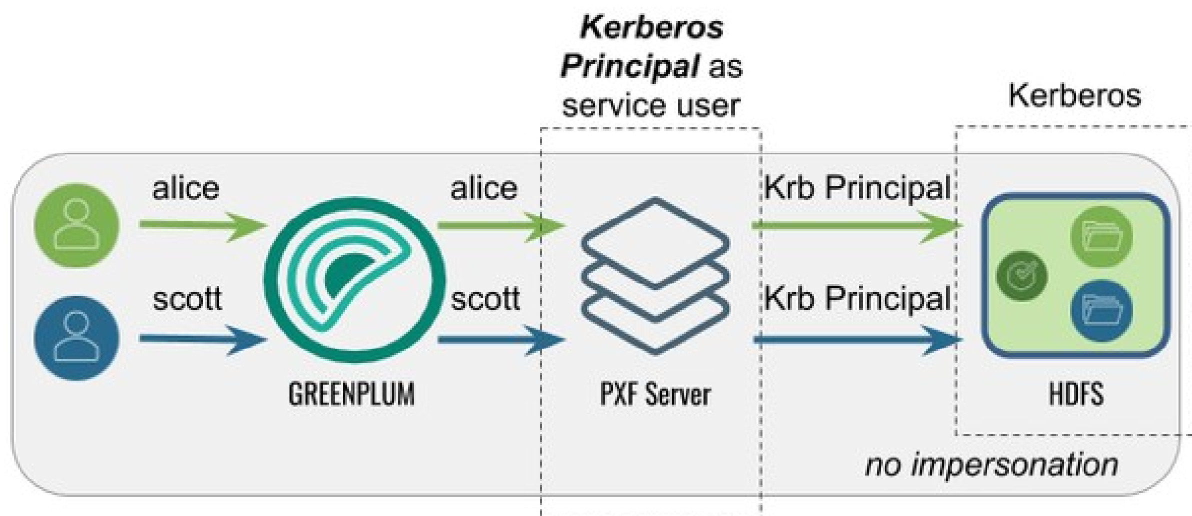

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
false	Kerberos principal	Perform the Configuration Procedure in this topic, and then turn off user impersonation as described in Configure PXF User Impersonation.	None required.

Accessing Hadoop as a <custom> User

In this configuration, PXF accesses Hadoop as a <custom> user (for example, `hive`). The Kerberos principal is the Hadoop proxy user. A query initiated by any Greenplum user appears on the Hadoop side as originating from the <custom> user.

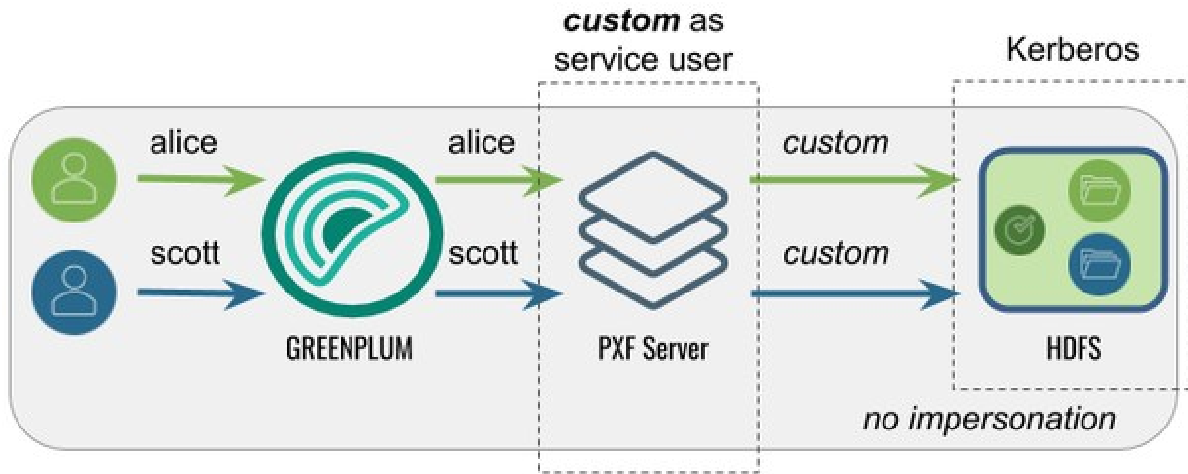

The following table identifies the `pxf.service.user.impersonation` and `pxf.service.user.name` settings, and the PXF and Hadoop configuration required for this use case:

Impersonation	Service User	PXF Configuration	Hadoop Configuration
false	<custom>	Perform the Configuration Procedure in this topic, turn off user impersonation as described in Configure PXF User Impersonation, and Configure the Hadoop User to the <custom> user name.	Set the Kerberos principal as the Hadoop proxy user as described in Configure Hadoop Proxying.

Note: PXF does not support accessing a Kerberos-secured Hadoop cluster with a <custom> user impersonating Greenplum users. PXF requires that you impersonate Greenplum users using the Kerberos principal.

Procedure

There are different procedures for configuring PXF for secure HDFS with a Microsoft Active Directory KDC Server vs. with an MIT Kerberos KDC Server.

Configuring PXF with a Microsoft Active Directory Kerberos KDC Server

When you configure PXF for secure HDFS using an AD Kerberos KDC server, you will perform tasks on both the KDC server host and the Greenplum Database master host.

Perform the following steps the Active Directory domain controller:

1. Start **Active Directory Users and Computers**.
2. Expand the forest domain and the top-level UNIX organizational unit that describes your Greenplum user domain.
3. Select **Service Accounts**, right-click, then select **New->User**.
4. Type a name, eg. `ServiceGreenplumPROD1`, and change the login name to `gpadmin`. Note that the login name should be in compliance with POSIX standard and match `hadoop.proxyuser.<name>.hosts/groups` in the Hadoop `core-site.xml` and the Kerberos principal.

5. Type and confirm the Active Directory service account password. Select the **User cannot change password** and **Password never expires** check boxes, then click **Next**. For security reasons, if you can't have **Password never expires** checked, you will need to generate new keytab file (step 7) every time you change the password of the service account.

6. Click **Finish** to complete the creation of the new user principal.

7. Open Powershell or a command prompt and run the `ktpass` command to generate the keytab file. For example:

   ```
   powershell#>ktpass -out pxf.service.keytab -princ gpadmin@EXAMPLE.COM -mapUser
   ServiceGreenplumPROD1 -pass ******* -crypto all -ptype KRB5_NT_PRINCIPAL
   ```

   With Active Directory, the principal and the keytab file are shared by all Greenplum Database hosts.

8. Copy the `pxf.service.keytab` file to the Greenplum master host.

Perform the following procedure on the Greenplum Database master host:

1. Log in to the Greenplum Database master host. For example:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Identify the name of the PXF Hadoop server configuration, and navigate to the server configuration directory. For example, if the server is named `hdp3`:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/hdp3
   ```

3. If the server configuration does not yet include a `pxf-site.xml` file, copy the template file to the directory. For example:

   ```
   gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml .
   ```

4. Open the `pxf-site.xml` file in the editor of your choice, and update the keytab and principal property settings, if required. Specify the location of the keytab file and the Kerberos principal, substituting your realm. For example:

   ```
   <property>
       <name>pxf.service.kerberos.principal</name>
       <value>gpadmin@EXAMPLE.COM</value>
   </property>
   <property>
       <name>pxf.service.kerberos.keytab</name>
       <value>${pxf.conf}/keytabs/pxf.service.keytab</value>
   </property>
   ```

5. Enable user impersonation as described in Configure PXF User Impersonation, and configure or verify Hadoop proxying for the *primary* component of the Kerberos principal as described in Configure Hadoop Proxying. For example, if your principal is `gpadmin@EXAMPLE.COM`, configure proxying for the Hadoop user `gpadmin`.

6. Save the file and exit the editor.

7. Synchronize the PXF configuration to your Greenplum Database cluster and restart PXF:

```
gpadmin@master$ pxf cluster sync
gpadmin@master$ pxf cluster restart
```

8. Step 7 does not synchronize the keytabs in `$PXF_BASE`. You must distribute the keytab file to `$PXF_BASE/keytabs/`. Locate the keytab file, copy the file to the `$PXF_BASE` runtime configuration directory, and set required permissions. For example:

```
gpadmin@gpmaster$ gpscp -f hostfile_all pxf.service.keytab =:$PXF_BASE/keytabs/
gpadmin@gpmaster$ gpssh -f hostfile_all chmod 400 $PXF_BASE/keytabs/pxf.service
.keytab
```

Configuring PXF with an MIT Kerberos KDC Server

When you configure PXF for secure HDFS using an MIT Kerberos KDC server, you will perform tasks on both the KDC server host and the Greenplum Database master host.

Perform the following steps on the MIT Kerberos KDC server host:

1. Log in to the Kerberos KDC server as the `root` user.

```
$ ssh root@<kdc-server>
root@kdc-server$
```

2. Distribute the `/etc/krb5.conf` Kerberos configuration file on the KDC server host to **each** host in your Greenplum Database cluster if not already present. For example:

```
root@kdc-server$ scp /etc/krb5.conf <gphost>:/etc/krb5.conf
```

3. Use the `kadmin.local` command to create a Kerberos PXF Service principal for **each** Greenplum Database host. The service principal should be of the form `gpadmin/<gphost>@<realm>` where <gphost> is the DNS resolvable, fully-qualified hostname of the host system (output of the `hostname -f` command).

   For example, these commands create Kerberos PXF Service principals for the hosts named host1.example.com, host2.example.com, and host3.example.com in the Kerberos realm named `EXAMPLE.COM`:

```
root@kdc-server$ kadmin.local -q "addprinc -randkey -pw changeme gpadmin/host1.
example.com@EXAMPLE.COM"
root@kdc-server$ kadmin.local -q "addprinc -randkey -pw changeme gpadmin/host2.
example.com@EXAMPLE.COM"
root@kdc-server$ kadmin.local -q "addprinc -randkey -pw changeme gpadmin/host3.
example.com@EXAMPLE.COM"
```

4. Generate a keytab file for each PXF Service principal that you created in the previous step. Save the keytab files in any convenient location (this example uses the directory `/etc/security/keytabs`). You will deploy the keytab files to their respective Greenplum Database host machines in a later step. For example:

```
root@kdc-server$ kadmin.local -q "xst -norandkey -k /etc/security/keytabs/pxf-h
ost1.service.keytab gpadmin/host1.example.com@EXAMPLE.COM"
root@kdc-server$ kadmin.local -q "xst -norandkey -k /etc/security/keytabs/pxf-h
ost2.service.keytab gpadmin/host2.example.com@EXAMPLE.COM"
root@kdc-server$ kadmin.local -q "xst -norandkey -k /etc/security/keytabs/pxf-h
```

```
ost3.service.keytab gpadmin/host3.example.com@EXAMPLE.COM"
```

Repeat the `xst` command as necessary to generate a keytab for each PXF Service principal that you created in the previous step.

5. List the principals. For example:

```
root@kdc-server$ kadmin.local -q "listprincs"
```

6. Copy the keytab file for each PXF Service principal to its respective host. For example, the following commands copy each principal generated in step 4 to the PXF default keytab directory on the host when `PXF_BASE=/usr/local/pxf-gp6`:

```
root@kdc-server$ scp /etc/security/keytabs/pxf-host1.service.keytab host1.examp
le.com:/usr/local/pxf-gp6/keytabs/pxf.service.keytab
root@kdc-server$ scp /etc/security/keytabs/pxf-host2.service.keytab host2.examp
le.com:/usr/local/pxf-gp6/keytabs/pxf.service.keytab
root@kdc-server$ scp /etc/security/keytabs/pxf-host3.service.keytab host3.examp
le.com:/usr/local/pxf-gp6/keytabs/pxf.service.keytab
```

Note the file system location of the keytab file on each PXF host; you will need this information for a later configuration step.

7. Change the ownership and permissions on the `pxf.service.keytab` files. The files must be owned and readable by only the `gpadmin` user. For example:

```
root@kdc-server$ ssh host1.example.com chown gpadmin:gpadmin /usr/local/pxf-gp6
/keytabs/pxf.service.keytab
root@kdc-server$ ssh host1.example.com chmod 400 /usr/local/pxf-gp6/keytabs/pxf
.service.keytab
root@kdc-server$ ssh host2.example.com chown gpadmin:gpadmin /usr/local/pxf-gp6
/keytabs/pxf.service.keytab
root@kdc-server$ ssh host2.example.com chmod 400 /usr/local/pxf-gp6/keytabs/pxf
.service.keytab
root@kdc-server$ ssh host3.example.com chown gpadmin:gpadmin /usr/local/pxf-gp6
/keytabs/pxf.service.keytab
root@kdc-server$ ssh host3.example.com chmod 400 /usr/local/pxf-gp6/keytabs/pxf
.service.keytab
```

Perform the following steps on the Greenplum Database master host:

1. Log in to the master host. For example:

```
$ ssh gpadmin@<gpmaster>
```

2. Identify the name of the PXF Hadoop server configuration that requires Kerberos access.
3. Navigate to the server configuration directory. For example, if the server is named `hdp3`:

```
gpadmin@gpmaster$ cd $PXF_BASE/servers/hdp3
```

4. If the server configuration does not yet include a `pxf-site.xml` file, copy the template file to the directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml .
```

5. Open the `pxf-site.xml` file in the editor of your choice, and update the keytab and principal property settings, if required. Specify the location of the keytab file and the Kerberos principal, substituting your realm. *The default values for these settings are identified below*:

```
<property>
    <name>pxf.service.kerberos.principal</name>
    <value>gpadmin/_HOST@EXAMPLE.COM</value>
</property>
<property>
    <name>pxf.service.kerberos.keytab</name>
    <value>${pxf.conf}/keytabs/pxf.service.keytab</value>
</property>
```

   PXF automatically replaces `_HOST` with the FQDN of the host.

6. If you want to access Hadoop as the Greenplum Database user:
   1. Enable user impersonation as described in Configure PXF User Impersonation.
   2. Configure Hadoop proxying for the *primary* component of the Kerberos principal as described in Configure Hadoop Proxying. For example, if your principal is `gpadmin/_HOST@EXAMPLE.COM`, configure proxying for the Hadoop user `gpadmin`.
7. If you want to access Hadoop using the identity of the Kerberos principal, disable user impersonation as described in Configure PXF User Impersonation.
8. If you want to access Hadoop as a custom user:
   1. Disable user impersonation as described in Configure PXF User Impersonation.
   2. Configure the custom user name as described in Configure the Hadoop User.
   3. Configure Hadoop proxying for the *primary* component of the Kerberos principal as described in Configure Hadoop Proxying. For example, if your principal is `gpadmin/_HOST@EXAMPLE.COM`, configure proxying for the Hadoop user `gpadmin`.
9. Save the file and exit the editor.
10. Synchronize the PXF configuration to your Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Configuring Connectors to Minio and S3 Object Stores (Optional)

You can use PXF to access S3-compatible object stores. This topic describes how to configure the PXF connectors to these external data sources.

If you do not plan to use these PXF object store connectors, then you do not need to perform this procedure.

About Object Store Configuration

To access data in an object store, you must provide a server location and client credentials. When you configure a PXF object store connector, you add at least one named PXF server configuration for the connector as described in Configuring PXF Servers.

PXF provides a configuration file template for each object store connector. These template files are located in the `$PXF_HOME/templates/` directory.

Minio Server Configuration

The template configuration file for Minio is `$PXF_HOME/templates/minio-site.xml`. When you configure a Minio server, you must provide the following server configuration properties and replace the template values with your credentials:

Property	Description	Value
fs.s3a.endpoint	The Minio S3 endpoint to which to connect.	Your endpoint.
fs.s3a.access.key	The Minio account access key ID.	Your access key.
fs.s3a.secret.key	The secret key associated with the Minio access key ID.	Your secret key.

S3 Server Configuration

The template configuration file for S3 is `$PXF_HOME/templates/s3-site.xml`. When you configure an S3 server, you must provide the following server configuration properties and replace the template values with your credentials:

Property	Description	Value
fs.s3a.access.key	The AWS account access key ID.	Your access key.
fs.s3a.secret.key	The secret key associated with the AWS access key ID.	Your secret key.

If required, fine-tune PXF S3 connectivity by specifying properties identified in the S3A section of the Hadoop-AWS module documentation in your `s3-site.xml` server configuration file.

You can override the credentials for an S3 server configuration by directly specifying the S3 access ID and secret key via custom options in the CREATE EXTERNAL TABLE command LOCATION clause. Refer to Overriding the S3 Server Configuration with DDL for additional information.

Configuring S3 Server-Side Encryption

PXF supports Amazon Web Service S3 Server-Side Encryption (SSE) for S3 files that you access with readable and writable Greenplum Database external tables that specify the `pxf` protocol and an `s3:*` profile. AWS S3 server-side encryption protects your data at rest; it encrypts your object data as it writes to disk, and transparently decrypts the data for you when you access it.

PXF supports the following AWS SSE encryption key management schemes:

- SSE with S3-Managed Keys (SSE-S3) - Amazon manages the data and master encryption keys.
- SSE with Key Management Service Managed Keys (SSE-KMS) - Amazon manages the data key, and you manage the encryption key in AWS KMS.
- SSE with Customer-Provided Keys (SSE-C) - You set and manage the encryption key.

Your S3 access key and secret key govern your access to all S3 bucket objects, whether the data is encrypted or not.

S3 transparently decrypts data during a read operation of an encrypted file that you access via a readable external table that is created by specifying the `pxf` protocol and an `s3:*` profile. No additional configuration is required.

To encrypt data that you write to S3 via this type of external table, you have two options:

- Configure the default SSE encryption key management scheme on a per-S3-bucket basis via the AWS console or command line tools (recommended).
- Configure SSE encryption options in your PXF S3 server `s3-site.xml` configuration file.

Configuring SSE via an S3 Bucket Policy (Recommended)

You can create S3 *Bucket Policy*(s) that identify the objects that you want to encrypt, the encryption key management scheme, and the write actions permitted on those objects. Refer to Protecting Data Using Server-Side Encryption in the AWS S3 documentation for more information about the SSE encryption key management schemes. How Do I Enable Default Encryption for an S3 Bucket? describes how to set default encryption bucket policies.

Specifying SSE Options in a PXF S3 Server Configuration

You must include certain properties in `s3-site.xml` to configure server-side encryption in a PXF S3 server configuration. The properties and values that you add to the file are dependent upon the SSE encryption key management scheme.

SSE-S3

To enable SSE-S3 on any file that you write to any S3 bucket, set the following encryption algorithm property and value in the `s3-site.xml` file:

```
<property>
  <name>fs.s3a.server-side-encryption-algorithm</name>
  <value>AES256</value>
</property>
```

To enable SSE-S3 for a specific S3 bucket, use the property name variant that includes the bucket name. For example:

```
<property>
  <name>fs.s3a.bucket.YOUR_BUCKET1_NAME.server-side-encryption-algorithm</name>
  <value>AES256</value>
</property>
```

Replace `YOUR_BUCKET1_NAME` with the name of the S3 bucket.

SSE-KMS

To enable SSE-KMS on any file that you write to any S3 bucket, set both the encryption algorithm and encryption key ID. To set these properties in the `s3-site.xml` file:

```
<property>
  <name>fs.s3a.server-side-encryption-algorithm</name>
  <value>SSE-KMS</value>
</property>
<property>
  <name>fs.s3a.server-side-encryption.key</name>
```

```
  <value>YOUR_AWS_SSE_KMS_KEY_ARN</value>
</property>
```

Substitute `YOUR_AWS_SSE_KMS_KEY_ARN` with your key resource name. If you do not specify an encryption key, the default key defined in the Amazon KMS is used. Example KMS key: `arn:aws:kms:us-west-2:123456789012:key/1a23b456-7890-12cc-d345-6ef7890g12f3`.

Note: Be sure to create the bucket and the key in the same Amazon Availability Zone.

To enable SSE-KMS for a specific S3 bucket, use property name variants that include the bucket name. For example:

```
<property>
  <name>fs.s3a.bucket.YOUR_BUCKET2_NAME.server-side-encryption-algorithm</name>
  <value>SSE-KMS</value>
</property>
<property>
  <name>fs.s3a.bucket.YOUR_BUCKET2_NAME.server-side-encryption.key</name>
  <value>YOUR_AWS_SSE_KMS_KEY_ARN</value>
</property>
```

Replace `YOUR_BUCKET2_NAME` with the name of the S3 bucket.

SSE-C

To enable SSE-C on any file that you write to any S3 bucket, set both the encryption algorithm and the encryption key (base-64 encoded). All clients must share the same key.

To set these properties in the `s3-site.xml` file:

```
<property>
  <name>fs.s3a.server-side-encryption-algorithm</name>
  <value>SSE-C</value>
</property>
<property>
  <name>fs.s3a.server-side-encryption.key</name>
  <value>YOUR_BASE64-ENCODED_ENCRYPTION_KEY</value>
</property>
```

To enable SSE-C for a specific S3 bucket, use the property name variants that include the bucket name as described in the SSE-KMS example.

Example Server Configuration Procedure

In this procedure, you name and add a PXF server configuration in the `$PXF_BASE/servers` directory on the Greenplum Database master host for the S3 Cloud Storage connector. You then use the `pxf cluster sync` command to sync the server configuration(s) to the Greenplum Database cluster.

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Choose a name for the server. You will provide the name to end users that need to reference files in the object store.

3. Create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a server configuration for an S3 server named `s3srvcfg`:

```
gpadmin@gpmaster$ mkdir $PXF_BASE/servers/s3srvcfg
```

4. Copy the PXF template file for S3 to the server configuration directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/s3-site.xml $PXF_BASE/servers/s3srvcfg
/
```

5. Open the template server configuration file in the editor of your choice, and provide appropriate property values for your environment. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
    <property>
        <name>fs.s3a.access.key</name>
        <value>access_key_for_user1</value>
    </property>
    <property>
        <name>fs.s3a.secret.key</name>
        <value>secret_key_for_user1</value>
    </property>
    <property>
        <name>fs.s3a.fast.upload</name>
        <value>true</value>
    </property>
</configuration>
```

6. Save your changes and exit the editor.

7. Use the `pxf cluster sync` command to copy the new server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Configuring Connectors to Azure and Google Cloud Storage Object Stores (Optional)

You can use PXF to access Azure Data Lake, Azure Blob Storage, and Google Cloud Storage object stores. This topic describes how to configure the PXF connectors to these external data sources.

If you do not plan to use these PXF object store connectors, then you do not need to perform this procedure.

About Object Store Configuration

To access data in an object store, you must provide a server location and client credentials. When you configure a PXF object store connector, you add at least one named PXF server configuration for the connector as described in Configuring PXF Servers.

PXF provides a template configuration file for each object store connector. These template files are located in the `$PXF_HOME/templates/` directory.

Azure Blob Storage Server Configuration

The template configuration file for Azure Blob Storage is `$PXF_HOME/templates/wasbs-site.xml`. When you configure an Azure Blob Storage server, you must provide the following server configuration properties and replace the template value with your account name:

Property	Description	Value
fs.adl.oauth2.access.token.provider.type	The token type.	Must specify `ClientCredential`.
fs.azure.account.key.<YOUR_AZURE_BLOB_STORAGE_ACCOUNT_NAME>.blob.core.windows.net	The Azure account key.	Replace with your account key.
fs.AbstractFileSystem.wasbs.impl	The file system class name.	Must specify `org.apache.hadoop.fs.azure.Wasbs`.

Azure Data Lake Server Configuration

The template configuration file for Azure Data Lake is `$PXF_HOME/templates/adl-site.xml`. When you configure an Azure Data Lake server, you must provide the following server configuration properties and replace the template values with your credentials:

Property	Description	Value
fs.adl.oauth2.access.token.provider.type	The type of token.	Must specify `ClientCredential`.
fs.adl.oauth2.refresh.url	The Azure endpoint to which to connect.	Your refresh URL.
fs.adl.oauth2.client.id	The Azure account client ID.	Your client ID (UUID).
fs.adl.oauth2.credential	The password for the Azure account client ID.	Your password.

Google Cloud Storage Server Configuration

The template configuration file for Google Cloud Storage is `$PXF_HOME/templates/gs-site.xml`. When you configure a Google Cloud Storage server, you must provide the following server configuration properties and replace the template values with your credentials:

Property	Description	Value
google.cloud.auth.service.account.enable	Enable service account authorization.	Must specify `true`.
google.cloud.auth.service.account.json.keyfile	The Google Storage key file.	Path to your key file.
fs.AbstractFileSystem.gs.impl	The file system class name.	Must specify `com.google.cloud.hadoop.fs.gcs.GoogleHadoopFS`.

Example Server Configuration Procedure

In this procedure, you name and add a PXF server configuration in the `$PXF_BASE/servers` directory

on the Greenplum Database master host for the Google Cloud Storate (GCS) connector. You then use the `pxf cluster sync` command to sync the server configuration(s) to the Greenplum Database cluster.

1. Log in to your Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Choose a name for the server. You will provide the name to end users that need to reference files in the object store.

3. Create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a server configuration for a Google Cloud Storage server named `gs_public`:

```
gpadmin@gpmaster$ mkdir $PXF_BASE/servers/gs_public
```

4. Copy the PXF template file for GCS to the server configuration directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/gs-site.xml $PXF_BASE/servers/gs_publi
c/
```

5. Open the template server configuration file in the editor of your choice, and provide appropriate property values for your environment. For example, if your Google Cloud Storage key file is located in `/home/gpadmin/keys/gcs-account.key.json`:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
    <property>
        <name>google.cloud.auth.service.account.enable</name>
        <value>true</value>
    </property>
    <property>
        <name>google.cloud.auth.service.account.json.keyfile</name>
        <value>/home/gpadmin/keys/gcs-account.key.json</value>
    </property>
    <property>
        <name>fs.AbstractFileSystem.gs.impl</name>
        <value>com.google.cloud.hadoop.fs.gcs.GoogleHadoopFS</value>
    </property>
</configuration>
```

6. Save your changes and exit the editor.

7. Use the `pxf cluster sync` command to copy the new server configurations to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Configuring the JDBC Connector (Optional)

You can use PXF to access an external SQL database including MySQL, ORACLE, Microsoft SQL Server, DB2, PostgreSQL, Hive, and Apache Ignite. This topic describes how to configure the PXF JDBC Connector to access these external data sources.

If you do not plan to use the PXF JDBC Connector, then you do not need to perform this procedure.

About JDBC Configuration

To access data in an external SQL database with the PXF JDBC Connector, you must:

- Register a compatible JDBC driver JAR file
- Specify the JDBC driver class name, database URL, and client credentials

In previous releases of Greenplum Database, you may have specified the JDBC driver class name, database URL, and client credentials via options in the `CREATE EXTERNAL TABLE` command. PXF now supports file-based server configuration for the JDBC Connector. This configuration, described below, allows you to specify these options and credentials in a file.

Note: PXF external tables that you previously created that directly specified the JDBC connection options will continue to work. If you want to move these tables to use JDBC file-based server configuration, you must create a server configuration, drop the external tables, and then recreate the tables specifying an appropriate `SERVER=<server_name>` clause.

JDBC Driver JAR Registration

PXF is bundled with the `postgresql-42.2.14.jar` JAR file. If you require a different JDBC driver, ensure that you install the JDBC driver JAR file for the external SQL database in the `$PXF_BASE/lib` directory on each Greenplum host. Be sure to install JDBC driver JAR files that are compatible with your JRE version. See Registering PXF Library Dependencies for additional information.

JDBC Server Configuration

When you configure the PXF JDBC Connector, you add at least one named PXF server configuration for the connector as described in Configuring PXF Servers. You can also configure one or more statically-defined queries to run against the remote SQL database.

PXF provides a template configuration file for the JDBC Connector. This server template configuration file, located in `$PXF_HOME/templates/jdbc-site.xml`, identifies properties that you can configure to establish a connection to the external SQL database. The template also includes optional properties that you can set before executing query or insert commands in the external database session.

The required properties in the `jdbc-site.xml` server template file follow:

Property	Description	Value
jdbc.driver	Class name of the JDBC driver.	The JDBC driver Java class name; for example `org.postgresql.Driver`.
jdbc.url	The URL that the JDBC driver uses to connect to the database.	The database connection URL (database-specific); for example `jdbc:postgresql://phost:pport/pdatabase`.
jdbc.user	The database user name.	The user name for connecting to the database.
jdbc.password	The password for `jdbc.user`.	The password for connecting to the database.

When you configure a PXF JDBC server, you specify the external database user credentials to PXF in clear text in a configuration file.

Connection-Level Properties

To set additional JDBC connection-level properties, add `jdbc.connection.property.<CPROP_NAME>` properties to `jdbc-site.xml`. PXF passes these properties to the JDBC driver when it establishes the connection to the external SQL database (`DriverManager.getConnection()`).

Replace `<CPROP_NAME>` with the connection property name and specify its value:

Property	Description	Value
jdbc.connection.property.<CPROP_NAME>	The name of a property (<CPROP_NAME>) to pass to the JDBC driver when PXF establishes the connection to the external SQL database.	The value of the <CPROP_NAME> property.

Example: To set the `createDatabaseIfNotExist` connection property on a JDBC connection to a PostgreSQL database, include the following property block in `jdbc-site.xml`:

```
<property>
    <name>jdbc.connection.property.createDatabaseIfNotExist</name>
    <value>true</value>
 </property>
```

Ensure that the JDBC driver for the external SQL database supports any connection-level property that you specify.

Connection Transaction Isolation Property

The SQL standard defines four transaction isolation levels. The level that you specify for a given connection to an external SQL database determines how and when the changes made by one transaction executed on the connection are visible to another.

The PXF JDBC Connector exposes an optional server configuration property named `jdbc.connection.transactionIsolation` that enables you to specify the transaction isolation level. PXF sets the level (`setTransactionIsolation()`) just after establishing the connection to the external SQL database.

The JDBC Connector supports the following `jdbc.connection.transactionIsolation` property values:

SQL Level	PXF Property Value
Read uncommitted	READ_UNCOMMITTED
Read committed	READ_COMMITTED
Repeatable Read	REPEATABLE_READ
Serializable	SERIALIZABLE

For example, to set the transaction isolation level to *Read uncommitted*, add the following property block to the `jdbc-site.xml` file:

```
<property>
    <name>jdbc.connection.transactionIsolation</name>
    <value>READ_UNCOMMITTED</value>
</property>
```

Different SQL databases support different transaction isolation levels. Ensure that the external database supports the level that you specify.

Statement-Level Properties

The PXF JDBC Connector executes a query or insert command on an external SQL database table in a *statement*. The Connector exposes properties that enable you to configure certain aspects of the statement before the command is executed in the external database. The Connector supports the following statement-level properties:

Property	Description	Value
jdbc.statement.batchSize	The number of rows to write to the external database table in a batch.	The number of rows. The default write batch size is 100.
jdbc.statement.fetchSize	The number of rows to fetch/buffer when reading from the external database table.	The number of rows. The default read fetch size for MySQL is `-2147483648` (`Integer.MIN_VALUE`). The default read fetch size for all other databases is 1000.
jdbc.statement.queryTimeout	The amount of time (in seconds) the JDBC driver waits for a statement to execute. This timeout applies to statements created for both read and write operations.	The timeout duration in seconds. The default wait time is unlimited.

PXF uses the default value for any statement-level property that you do not explicitly configure.

Example: To set the read fetch size to 5000, add the following property block to `jdbc-site.xml`:

```
<property>
    <name>jdbc.statement.fetchSize</name>
    <value>5000</value>
</property>
```

Ensure that the JDBC driver for the external SQL database supports any statement-level property that you specify.

Session-Level Properties

To set session-level properties, add the `jdbc.session.property.<SPROP_NAME>` property to `jdbc-site.xml`. PXF will `SET` these properties in the external database before executing a query.

Replace `<SPROP_NAME>` with the session property name and specify its value:

Property	Description	Value
jdbc.session.property.<SPROP_NAME>	The name of a session property (<SPROP_NAME>) to set before query execution.	The value of the <SPROP_NAME> property.

Note: The PXF JDBC Connector passes both the session property name and property value to the external SQL database exactly as specified in the `jdbc-site.xml` server configuration file. To limit the potential threat of SQL injection, the Connector rejects any property name or value that contains the `;`, `\n`, `\b`, or `\0` characters.

The PXF JDBC Connector handles the session property `SET` syntax for all supported external SQL databases.

Example: To set the `search_path` parameter before running a query in a PostgreSQL database, add the following property block to `jdbc-site.xml`:

```
<property>
    <name>jdbc.session.property.search_path</name>
    <value>public</value>
</property>
```

Ensure that the JDBC driver for the external SQL database supports any property that you specify.

About JDBC Connection Pooling

The PXF JDBC Connector uses JDBC connection pooling implemented by HikariCP. When a user queries or writes to an external table, the Connector establishes a connection pool for the associated server configuration the first time that it encounters a unique combination of `jdbc.url`, `jdbc.user`, `jdbc.password`, connection property, and pool property settings. The Connector reuses connections in the pool subject to certain connection and timeout settings.

One or more connection pools may exist for a given server configuration, and user access to different external tables specifying the same server may share a connection pool.

Note: If you have enabled JDBC user impersonation in a server configuration, the JDBC Connector creates a separate connection pool for each Greenplum Database user that accesses any external table specifying that server configuration.

The `jdbc.pool.enabled` property governs JDBC connection pooling for a server configuration. Connection pooling is enabled by default. To disable JDBC connection pooling for a server configuration, set the property to false:

```
<property>
    <name>jdbc.pool.enabled</name>
    <value>false</value>
</property>
```

If you disable JDBC connection pooling for a server configuration, PXF does not reuse JDBC connections for that server. PXF creates a connection to the remote database for every partition of a query, and closes the connection when the query for that partition completes.

PXF exposes connection pooling properties that you can configure in a JDBC server definition. These properties are named with the `jdbc.pool.property.` prefix and *apply to each PXF JVM*. The JDBC Connector automatically sets the following connection pool properties and default values:

Property	Description	Default Value
jdbc.pool.property.maximumPoolSize	The maximum number of connections to the database backend.	15
jdbc.pool.property.connectionTimeout	The maximum amount of time, in milliseconds, to wait for a connection from the pool.	30000
jdbc.pool.property.idleTimeout	The maximum amount of time, in milliseconds, after which an inactive connection is considered idle.	30000
jdbc.pool.property.minimumIdle	The minimum number of idle connections maintained in the connection pool.	0

You can set other HikariCP-specific connection pooling properties for a server configuration by specifying `jdbc.pool.property.<HIKARICP_PROP_NAME>` and the desired value in the `jdbc-site.xml` configuration file for the server. Also note that the JDBC Connector passes along any property that you specify with a `jdbc.connection.property.` prefix when it requests a connection from the JDBC `DriverManager`. Refer to Connection-Level Properties above.

Tuning the Maximum Connection Pool Size

To not exceed the maximum number of connections allowed by the target database, and at the same time ensure that each PXF JVM services a fair share of the JDBC connections, determine the maximum value of `maximumPoolSize` based on the size of the Greenplum Database cluster as follows:

```
max_conns_allowed_by_remote_db / #_greenplum_segment_hosts
```

For example, if your Greenplum Database cluster has 16 segment hosts and the target database allows 160 concurrent connections, calculate `maximumPoolSize` as follows:

```
160 / 16 = 10
```

In practice, you may choose to set `maximumPoolSize` to a lower value, since the number of concurrent connections per JDBC query depends on the number of partitions used in the query. When a query uses no partitions, a single PXF JVM services the query. If a query uses 12 partitions, PXF establishes 12 concurrent JDBC connections to the remote database. Ideally, these connections are distributed equally among the PXF JVMs, but that is not guaranteed.

JDBC User Impersonation

The PXF JDBC Connector uses the `jdbc.user` setting or information in the `jdbc.url` to determine the identity of the user to connect to the external data store. When PXF JDBC user impersonation is disabled (the default), the behavior of the JDBC Connector is further dependent upon the external data store. For example, if you are using the JDBC Connector to access Hive, the Connector uses the settings of certain Hive authentication and impersonation properties to determine the user. You may be required to provide a `jdbc.user` setting, or add properties to the `jdbc.url` setting in the server `jdbc-site.xml` file. Refer to Configuring Hive Access via the JDBC Connector for more information on this procedure.

When you enable PXF JDBC user impersonation, the PXF JDBC Connector accesses the external data store on behalf of a Greenplum Database end user. The Connector uses the name of the Greenplum Database user that accesses the PXF external table to try to connect to the external data store.

When you enable JDBC user impersonation for a PXF server, PXF overrides the value of a `jdbc.user` property setting defined in either `jdbc-site.xml` or `<greenplum_user_name>-user.xml`, or specified in the external table DDL, with the Greenplum Database user name. For user impersonation to work effectively when the external data store requires passwords to authenticate connecting users, you must specify the `jdbc.password` setting for each user that can be impersonated in that user's `<greenplum_user_name>-user.xml` property override file. Refer to Configuring a PXF User for more information about per-server, per-Greenplum-user configuration.

The `pxf.service.user.impersonation` property in the `jdbc-site.xml` configuration file governs

JDBC user impersonation.

Example Configuration Procedure

By default, PXF JDBC user impersonation is disabled. Perform the following procedure to turn PXF user impersonation on or off for a JDBC server configuration.

1. Log in to your Greenplum Database master node as the administrative user:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Identify the name of the PXF JDBC server configuration that you want to update.
3. Navigate to the server configuration directory. For example, if the server is named `mysqldb`:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/mysqldb
   ```

4. Open the `jdbc-site.xml` file in the editor of your choice, and add or uncomment the user impersonation property and setting. For example, if you require user impersonation for this server configuration, set the `pxf.service.user.impersonation` property to `true`:

   ```
   <property>
       <name>pxf.service.user.impersonation</name>
       <value>true</value>
   </property>
   ```

5. Save the `jdbc-site.xml` file and exit the editor.
6. Use the `pxf cluster sync` command to synchronize the PXF JDBC server configuration to your Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

About Session Authorization

Certain SQL databases, including PostgreSQL and DB2, allow a privileged user to change the effective database user that runs commands in a session. You might take advantage of this feature if, for example, you connect to the remote database as a proxy user and want to switch session authorization after establishing the database connection.

In databases that support it, you can configure a session property to switch the effective user. For example, in DB2, you use the `SET SESSION_USER <username>` command to switch the effective DB2 user. If you configure the DB2 `session_user` variable via a PXF session-level property (`jdbc.session.property.<SPROP_NAME>`) in your `jdbc-site.xml` file, PXF runs this command for you.

For example, to switch the effective DB2 user to the user named `bill`, you configure your `jdbc-site.xml` as follows:

```
<property>
    <name>jdbc.session.property.session_user</name>
    <value>bill</value>
</property>
```

After establishing the database connection, PXF implicitly runs the following command to set the

`session_user` DB2 session variable to the value that you configured:

```
SET SESSION_USER = bill
```

PXF recognizes a synthetic property value, `${pxf.session.user}`, that identifies the Greenplum Database user name. You may choose to use this value when you configure a property that requires a value that changes based on the Greenplum user running the session.

A scenario where you might use `${pxf.session.user}` is when you authenticate to the remote SQL database with Kerberos, the primary component of the Kerberos principal identifies the Greenplum Database user name, and you want to run queries in the remote database using this effective user name. For example, if you are accessing DB2, you would configure your `jdbc-site.xml` to specify the Kerberos `securityMechanism` and `KerberosServerPrincipal`, and then set the `session_user` variable as follows:

```
<property>
    <name>jdbc.session.property.session_user</name>
    <value>${pxf.session.user}</value>
</property>
```

With this configuration, PXF `SET`s the DB2 `session_user` variable to the current Greenplum Database user name, and runs subsequent operations on the DB2 table as that user.

Session Authorization Considerations for Connection Pooling

When PXF performs session authorization on your behalf and JDBC connection pooling is enabled (the default), you may choose to set the `jdbc.pool.qualifier` property. Setting this property instructs PXF to include the property value in the criteria that it uses to create and reuse connection pools. In practice, you would not set this to a fixed value, but rather to a value that changes based on the user/session/transaction, etc. When you set this property to `${pxf.session.user}`, PXF includes the Greenplum Database user name in the criteria that it uses to create and re-use connection pools. The default setting is no qualifier.

To make use of this feature, add or uncomment the following property block in `jdbc-site.xml` to prompt PXF to include the Greenplum user name in connection pool creation/reuse criteria:

```
<property>
    <name>jdbc.pool.qualifier</name>
    <value>${pxf.session.user}</value>
</property>
```

JDBC Named Query Configuration

A PXF *named query* is a static query that you configure, and that PXF runs in the remote SQL database.

To configure and use a PXF JDBC named query:

1. You define the query in a text file.
2. You provide the query name to Greenplum Database users.
3. The Greenplum Database user references the query in a Greenplum Database external table definition.

PXF runs the query each time the user invokes a `SELECT` command on the Greenplum Database external table.

Defining a Named Query

You create a named query by adding the query statement to a text file that has the following naming format: `<query_name>.sql`. You can define one or more named queries for a JDBC server configuration. Each query must reside in a separate text file.

You must place a query text file in the PXF JDBC server configuration directory from which it will be accessed. If you want to make the query available to more than one JDBC server configuration, you must copy the query text file to the configuration directory for each JDBC server.

The query text file must contain a single query that you want to run in the remote SQL database. You must construct the query in accordance with the syntax supported by the database.

For example, if a MySQL database has a `customers` table and an `orders` table, you could include the following SQL statement in a query text file:

```
SELECT c.name, c.city, sum(o.amount) AS total, o.month
  FROM customers c JOIN orders o ON c.id = o.customer_id
  WHERE c.state = 'CO'
GROUP BY c.name, c.city, o.month
```

You may optionally provide the ending semicolon (`;`) for the SQL statement.

Query Naming

The Greenplum Database user references a named query by specifying the query file name without the extension. For example, if you define a query in a file named `report.sql`, the name of that query is `report`.

Named queries are associated with a specific JDBC server configuration. You will provide the available query names to the Greenplum Database users that you allow to create external tables using the server configuration.

Referencing a Named Query

The Greenplum Database user specifies `query:<query_name>` rather than the name of a remote SQL database table when they create the external table. For example, if the query is defined in the file `$PXF_BASE/servers/mydb/report.sql`, the `CREATE EXTERNAL TABLE LOCATION` clause would include the following components:

```
LOCATION ('pxf://query:report?PROFILE=jdbc&SERVER=mydb ...')
```

Refer to About Using Named Queries for information about using PXF JDBC named queries.

Overriding the JDBC Server Configuration

You can override the JDBC server configuration by directly specifying certain JDBC properties via custom options in the `CREATE EXTERNAL TABLE` command `LOCATION` clause. Refer to Overriding the JDBC Server Configuration via DDL for additional information.

Configuring Access to Hive

You can use the JDBC Connector to access Hive. Refer to Configuring the JDBC Connector for Hive Access for detailed information on this configuration procedure.

Example Configuration Procedure

In this procedure, you name and add a PXF JDBC server configuration for a PostgreSQL database and synchronize the server configuration(s) to the Greenplum Database cluster.

1. Log in to your Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Choose a name for the JDBC server. You will provide the name to Greenplum users that you choose to allow to reference tables in the external SQL database as the configured user.

   Note: The server name `default` is reserved.

3. Create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a JDBC server configuration named `pg_user1_testdb`:

```
gpadmin@gpmaster$ mkdir $PXF_BASE/servers/pg_user1_testdb
```

4. Copy the PXF JDBC server template file to the server configuration directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/jdbc-site.xml $PXF_BASE/servers/pg_use
r1_testdb/
```

5. Open the template server configuration file in the editor of your choice, and provide appropriate property values for your environment. For example, if you are configuring access to a PostgreSQL database named `testdb` on a PostgreSQL instance running on the host named `pgserverhost` for the user named `user1`:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
    <property>
        <name>jdbc.driver</name>
        <value>org.postgresql.Driver</value>
    </property>
    <property>
        <name>jdbc.url</name>
        <value>jdbc:postgresql://pgserverhost:5432/testdb</value>
    </property>
    <property>
        <name>jdbc.user</name>
        <value>user1</value>
    </property>
    <property>
        <name>jdbc.password</name>
        <value>changeme</value>
    </property>
</configuration>
```

6. Save your changes and exit the editor.

7. Use the `pxf cluster sync` command to copy the new server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Configuring the JDBC Connector for Hive Access (Optional)

You can use the PXF JDBC Connector to retrieve data from Hive. You can also use a JDBC named query to submit a custom SQL query to Hive and retrieve the results using the JDBC Connector.

This topic describes how to configure the PXF JDBC Connector to access Hive. When you configure Hive access with JDBC, you must take into account the Hive user impersonation setting, as well as whether or not the Hadoop cluster is secured with Kerberos.

If you do not plan to use the PXF JDBC Connector to access Hive, then you do not need to perform this procedure.

JDBC Server Configuration

The PXF JDBC Connector is installed with the JAR files required to access Hive via JDBC, `hive-jdbc-<version>.jar` and `hive-service-<version>.jar`, and automatically registers these JARs.

When you configure a PXF JDBC server for Hive access, you must specify the JDBC driver class name, database URL, and client credentials just as you would when configuring a client connection to an SQL database.

To access Hive via JDBC, you must specify the following properties and values in the `jdbc-site.xml` server configuration file:

Property	Value
jdbc.driver	org.apache.hive.jdbc.HiveDriver
jdbc.url	jdbc:hive2://<hiveserver2_host>:<hiveserver2_port>/<database>

The value of the HiveServer2 authentication (`hive.server2.authentication`) and impersonation (`hive.server2.enable.doAs`) properties, and whether or not the Hive service is utilizing Kerberos authentication, will inform the setting of other JDBC server configuration properties. These properties are defined in the `hive-site.xml` configuration file in the Hadoop cluster. You will need to obtain the values of these properties.

The following table enumerates the Hive2 authentication and impersonation combinations supported by the PXF JDBC Connector. It identifies the possible Hive user identities and the JDBC server configuration required for each.

Table heading key:

- *authentication* -> Hive `hive.server2.authentication` Setting
- *enable.doAs* -> Hive `hive.server2.enable.doAs` Setting
- *User Identity* -> Identity that HiveServer2 will use to access data
- *Configuration Required* -> PXF JDBC Connector or Hive configuration required for *User Identity*

authentication	enable.doAs	User Identity	Configuration Required
`NOSASL`	n/a	No authentication	Must set `jdbc.connection.property.auth` = `noSasl`.
`NONE`, or not specified	`TRUE`	User name that you provide	Set `jdbc.user`.
`NONE`, or not specified	`TRUE`	Greenplum user name	Set `pxf.service.user.impersonation` to `true` in `jdbc-site.xml`.
`NONE`, or not specified	`FALSE`	Name of the user who started Hive, typically `hive`	None
`KERBEROS`	`TRUE`	Identity provided in the PXF Kerberos principal, typically `gpadmin`	Must set `hadoop.security.authentication` to `kerberos` in `jdbc-site.xml`.
`KERBEROS`	`TRUE`	User name that you provide	Set `hive.server2.proxy.user` in `jdbc.url` and set `hadoop.security.authentication` to `kerberos` in `jdbc-site.xml`.
`KERBEROS`	`TRUE`	Greenplum user name	Set `pxf.service.user.impersonation` to `true` and `hadoop.security.authentication` to `kerberos` in `jdbc-site.xml`.
`KERBEROS`	`FALSE`	Identity provided in the `jdbc.url` `principal` parameter, typically `hive`	Must set `hadoop.security.authentication` to `kerberos` in `jdbc-site.xml`.

Note: There are additional configuration steps required when Hive utilizes Kerberos authentication.

Example Configuration Procedure

Perform the following procedure to configure a PXF JDBC server for Hive:

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Choose a name for the JDBC server.
3. Create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a JDBC server configuration named `hivejdbc1`:

   ```
   gpadmin@gpmaster$ mkdir $PXF_BASE/servers/hivejdbc1
   ```

4. Navigate to the server configuration directory. For example:

   ```
   gpadmin@gpmaster$ cd $PXF_BASE/servers/hivejdbc1
   ```

5. Copy the PXF JDBC server template file to the server configuration directory. For example:

   ```
   gpadmin@gpmaster$ cp $PXF_HOME/templates/jdbc-site.xml .
   ```

6. When you access Hive secured with Kerberos, you also need to specify configuration properties in the `pxf-site.xml` file. *If this file does not yet exist in your server configuration,*

copy the `pxf-site.xml` template file to the server config directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml .
```

7. Open the `jdbc-site.xml` file in the editor of your choice and set the `jdbc.driver` and `jdbc.url` properties. Be sure to specify your Hive host, port, and database name:

```
<property>
    <name>jdbc.driver</name>
    <value>org.apache.hive.jdbc.HiveDriver</value>
</property>
<property>
    <name>jdbc.url</name>
    <value>jdbc:hive2://<hiveserver2_host>:<hiveserver2_port>/<database></value
>
</property>
```

8. Obtain the `hive-site.xml` file from your Hadoop cluster and examine the file.

9. If the `hive.server2.authentication` property in `hive-site.xml` is set to `NOSASL`, HiveServer2 performs no authentication. Add the following connection-level property to `jdbc-site.xml`:

```
<property>
    <name>jdbc.connection.property.auth</name>
    <value>noSasl</value>
</property>
```

Alternatively, you may choose to add `;auth=noSasl` to the `jdbc.url`.

10. If the `hive.server2.authentication` property in `hive-site.xml` is set to `NONE`, or the property is not specified, you must set the `jdbc.user` property. The value to which you set the `jdbc.user` property is dependent upon the `hive.server2.enable.doAs` impersonation setting in `hive-site.xml`:

    1. If `hive.server2.enable.doAs` is set to `TRUE` (the default), Hive runs Hadoop operations on behalf of the user connecting to Hive. *Choose/perform one of the following options*:

       Set `jdbc.user` to specify the user that has read permission on all Hive data accessed by Greenplum Database. For example, to connect to Hive and run all requests as user `gpadmin`:

       ```
       <property>
           <name>jdbc.user</name>
           <value>gpadmin</value>
       </property>
       ```

       Or, turn on JDBC server-level user impersonation so that PXF automatically uses the Greenplum Database user name to connect to Hive; uncomment the `pxf.service.user.impersonation` property in `jdbc-site.xml` and set the value to `true:

       ```
       <property>
           <name>pxf.service.user.impersonation</name>
           <value>true</value>
       ```

```
</property>
```

If you enable JDBC impersonation in this manner, you must not specify a `jdbc.user` nor include the setting in the `jdbc.url`.

2. If required, create a PXF user configuration file as described in Configuring a PXF User to manage the password setting.

3. If `hive.server2.enable.doAs` is set to `FALSE`, Hive runs Hadoop operations as the user who started the HiveServer2 process, usually the user `hive`. PXF ignores the `jdbc.user` setting in this circumstance.

11. If the `hive.server2.authentication` property in `hive-site.xml` is set to `KERBEROS`:

    1. Identify the name of the server configuration.

    2. Ensure that you have configured Kerberos authentication for PXF as described in Configuring PXF for Secure HDFS, and that you have specified the Kerberos principal and keytab in the `pxf-site.xml` properties as described in the procedure.

    3. Comment out the `pxf.service.user.impersonation` property in the `pxf-site.xml` file. If you require user impersonation, you will uncomment and set the property in an upcoming step.)

    4. Uncomment the `hadoop.security.authentication` setting in `$PXF_BASE/servers/<name>/jdbc-site.xml`:

    ```
    <property>
        <name>hadoop.security.authentication</name>
        <value>kerberos</value>
    </property>
    ```

    5. Add the `saslQop` property to `jdbc.url`, and set it to match the `hive.server2.thrift.sasl.qop` property setting in `hive-site.xml`. For example, if the `hive-site.xml` file includes the following property setting:

    ```
    <property>
        <name>hive.server2.thrift.sasl.qop</name>
        <value>auth-conf</value>
    </property>
    ```

    You would add `;saslQop=auth-conf` to the `jdbc.url`.

    6. Add the HiverServer2 `principal` name to the `jdbc.url`. For example:

    ```
    jdbc:hive2://hs2server:10000/default;principal=hive/hs2server@REALM;saslQop=auth-conf
    ```

    7. If `hive.server2.enable.doAs` is set to `TRUE` (the default), Hive runs Hadoop operations on behalf of the user connecting to Hive. *Choose/perform one of the following options*:

    Do not specify any additional properties. In this case, PXF initiates all Hadoop access with the identity provided in the PXF Kerberos principal (usually `gpadmin`).

    Or, set the `hive.server2.proxy.user` property in the `jdbc.url` to specify the user

that has read permission on all Hive data. For example, to connect to Hive and run all requests as the user named `integration` use the following `jdbc.url`:

```
jdbc:hive2://hs2server:10000/default;principal=hive/hs2server@REALM;saslQ
op=auth-conf;hive.server2.proxy.user=integration
```

Or, enable PXF JDBC impersonation in the `pxf-site.xml` file so that PXF automatically uses the Greenplum Database user name to connect to Hive. Add or uncomment the `pxf.service.user.impersonation` property and set the value to `true`. For example:

```
<property>
    <name>pxf.service.user.impersonation</name>
    <value>true</value>
</property>
```

If you enable JDBC impersonation, you must not explicitly specify a `hive.server2.proxy.user` in the `jdbc.url`.

8. If required, create a PXF user configuration file to manage the password setting.
9. If `hive.server2.enable.doAs` is set to `FALSE`, Hive runs Hadoop operations with the identity provided by the PXF Kerberos principal (usually `gpadmin`).

12. Save your changes and exit the editor.
13. Use the `pxf cluster sync` command to copy the new server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Starting, Stopping, and Restarting PXF

PXF provides two management commands:

- `pxf cluster` - manage all PXF Service instances in the Greenplum Database cluster
- `pxf` - manage the PXF Service instance on a specific Greenplum Database host

The procedures in this topic assume that you have added the `$PXF_HOME/bin` directory to your `$PATH`.

Starting PXF

After configuring PXF, you must start PXF on each host in your Greenplum Database cluster. The PXF Service, once started, runs as the `gpadmin` user on default port 5888. Only the `gpadmin` user can start and stop the PXF Service.

If you want to change the default PXF configuration, you must update the configuration before you start PXF, or restart PXF if it is already running. See About the PXF Configuration Files for information about the user-customizable PXF configuration properties and the configuration update procedure.

Prerequisites

Before you start PXF in your Greenplum Database cluster, ensure that:

- Your Greenplum Database cluster is up and running.
- You have previously configured PXF.

Procedure

Perform the following procedure to start PXF on each host in your Greenplum Database cluster.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Run the `pxf cluster start` command to start PXF on each host:

```
gpadmin@gpmaster$ pxf cluster start
```

Stopping PXF

If you must stop PXF, for example if you are upgrading PXF, you must stop PXF on each host in your Greenplum Database cluster. Only the `gpadmin` user can stop the PXF Service.

Prerequisites

Before you stop PXF in your Greenplum Database cluster, ensure that your Greenplum Database cluster is up and running.

Procedure

Perform the following procedure to stop PXF on each host in your Greenplum Database cluster.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Run the `pxf cluster stop` command to stop PXF on each host:

```
gpadmin@gpmaster$ pxf cluster stop
```

Restarting PXF

If you must restart PXF, for example if you updated PXF user configuration files in `$PXF_BASE/conf`, you run `pxf cluster restart` to stop, and then start, PXF on all hosts in your Greenplum Database cluster.

Only the `gpadmin` user can restart the PXF Service.

Prerequisites

Before you restart PXF in your Greenplum Database cluster, ensure that your Greenplum Database cluster is up and running.

Procedure

Perform the following procedure to restart PXF in your Greenplum Database cluster.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Restart PXF:

```
gpadmin@gpmaster$ pxf cluster restart
```

Granting Users Access to PXF

The Greenplum Platform Extension Framework (PXF) implements a protocol named `pxf` that you can use to create an external table that references data in an external data store. The PXF protocol and Java service are packaged as a Greenplum Database extension.

You must enable the PXF extension in each database in which you plan to use the framework to access external data. You must also explicitly `GRANT` permission to the `pxf` protocol to those users/roles who require access.

Enabling PXF in a Database

You must explicitly register the PXF extension in each Greenplum Database in which you plan to use the extension. You must have Greenplum Database administrator privileges to register an extension.

Perform the following procedure for ***each*** database in which you want to use PXF:

1. Connect to the database as the `gpadmin` user:

```
gpadmin@gpmaster$ psql -d <dbname> -U gpadmin
```

2. Create the PXF extension. You must have Greenplum Database administrator privileges to create an extension. For example:

```
dbname=# CREATE EXTENSION pxf;
```

Creating the `pxf` extension registers the `pxf` protocol and the call handlers required for PXF to access external data.

Disabling PXF in a Database

When you no longer want to use PXF on a specific database, you must explicitly drop the PXF extension for that database. You must have Greenplum Database administrator privileges to drop an extension.

1. Connect to the database as the `gpadmin` user:

```
gpadmin@gpmaster$ psql -d <dbname> -U gpadmin
```

2. Drop the PXF extension:

```
dbname=# DROP EXTENSION pxf;
```

The `DROP` command fails if there are any currently defined external tables using the `pxf` protocol. Add the `CASCADE` option if you choose to forcibly remove these external tables.

Granting a Role Access to PXF

To read external data with PXF, you create an external table with the `CREATE EXTERNAL TABLE` command that specifies the `pxf` protocol. You must specifically grant `SELECT` permission to the `pxf` protocol to all non-`SUPERUSER` Greenplum Database roles that require such access.

To grant a specific role access to the `pxf` protocol, use the `GRANT` command. For example, to grant the role named `bill` read access to data referenced by an external table created with the `pxf` protocol:

```
GRANT SELECT ON PROTOCOL pxf TO bill;
```

To write data to an external data store with PXF, you create an external table with the `CREATE WRITABLE EXTERNAL TABLE` command that specifies the `pxf` protocol. You must specifically grant `INSERT` permission to the `pxf` protocol to all non-`SUPERUSER` Greenplum Database roles that require such access. For example:

```
GRANT INSERT ON PROTOCOL pxf TO bill;
```

Registering PXF Library Dependencies

You use PXF to access data stored on external systems. Depending upon the external data store, this access may require that you install and/or configure additional components or services for the external data store.

PXF depends on JAR files and other configuration information provided by these additional components. In most cases, PXF manages internal JAR dependencies as necessary based on the connectors that you use.

Should you need to register a JAR or native library dependency with PXF, you copy the library to a location known to PXF *or* you inform PXF of a custom location, and then you must synchronize and restart PXF.

Registering a JAR Dependency

PXF loads JAR dependencies from the following directories, in this order:

1. The directories that you specify in the `$PXF_BASE/conf/pxf-env.sh` configuration file, `PXF_LOADER_PATH` environment variable. The `pxf-env.sh` file includes this commented-out block:

   ```
   # Additional locations to be class-loaded by PXF
   # export PXF_LOADER_PATH=
   ```

   You would uncomment the `PXF_LOADER_PATH` setting and specify one or more colon-separated directory names.

2. The default PXF JAR directory `$PXF_BASE/lib`.

To add a JAR dependency for PXF, for example a MySQL driver JAR file, you must log in to the Greenplum Database master host, copy the JAR file to the PXF user configuration runtime library directory (`$PXF_BASE/lib`), sync the PXF configuration to the Greenplum Database cluster, and then restart PXF on each host. For example:

```
$ ssh gpadmin@<gpmaster>
gpadmin@gpmaster$ cp new_dependent_jar.jar $PXF_BASE/lib/
gpadmin@gpmaster$ pxf cluster sync
gpadmin@gpmaster$ pxf cluster restart
```

Alternatively, you could have identified the file system location of the JAR in the `pxf-env.sh` `PXF_LOADER_PATH` environment variable. If you choose this registration option, you must ensure that you copy the JAR file to the same location on the Greenplum Database standby and segment hosts before you synchronize and restart PXF.

Registering a Native Library Dependency

PXF loads native libraries from the following directories, in this order:

1. The directories that you specify in the `$PXF_BASE/conf/pxf-env.sh` configuration file, `LD_LIBRARY_PATH` environment variable. The `pxf-env.sh` file includes this commented-out block:

   ```
   # Additional native libraries to be loaded by PXF
   # export LD_LIBRARY_PATH=
   ```

   You would uncomment the `LD_LIBRARY_PATH` setting and specify one or more colon-sperated directory names.

2. The default PXF native library directory `$PXF_BASE/lib/native`.
3. The default Hadoop native library directory `/usr/lib/hadoop/lib/native`.

As such, you have three file location options when you register a native library with PXF:

- Copy the library to the default PXF native library directory, `$PXF_BASE/lib/native`, on only the Greenplum Database master host. When you next synchronize PXF, PXF copies the native library to all hosts in the Greenplum cluster.
- Copy the library to the default Hadoop native library directory, `/usr/lib/hadoop/lib/native`, on the Greenplum master, standby, and each segment host.
- Copy the library to the same, custom location on the Greenplum master, standby, and each segment host, and uncomment and add the directory path to the `pxf-env.sh` `LD_LIBRARY_PATH` environment variable.

Procedure

1. Copy the native library file to one of the following:
    - The `$PXF_BASE/lib/native` directory on the Greenplum Database master host. (You may need to create this directory.)

- The `/usr/lib/hadoop/lib/native` directory on all Greenplum Database hosts.
- A user-defined location on all Greenplum Database hosts; note the file system location of the native library.

2. If you copied the native library to a custom location:
    1. Open the `$PXF_BASE/conf/pxf-env.sh` file in the editor of your choice, and uncomment the `LD_LIBRARY_PATH` setting:

    ```
    # Additional native libraries to be loaded by PXF
    export LD_LIBRARY_PATH=
    ```

    2. Specify the custom location in the `LD_LIBRARY_PATH` environment variable. For example, if you copied a library named `dependent_native_lib.so` to `/usr/local/lib` on all Greenplum hosts, you would set `LD_LIBRARY_PATH` as follows:

    ```
    export LD_LIBRARY_PATH=/usr/local/lib
    ```

    3. Save the file and exit the editor.
3. Synchronize the PXF configuration from the Greenplum Database master host to the standby and segment hosts.

    ```
    gpadmin@gpmaster$ pxf cluster sync
    ```

    If you copied the native library to the `$PXF_BASE/lib/native` directory, this command copies the library to the same location on the Greenplum Database standby and segment hosts.

    If you updated the `pxf-env.sh` `LD_LIBRARY_PATH` environment variable, this command copies the configuration change to the Greenplum Database standby and segment hosts.
4. Restart PXF on all Greenplum hosts:

    ```
    gpadmin@gpmaster$ pxf cluster restart
    ```

Monitoring PXF

You can monitor the status of PXF from the command line.

PXF also provides additional information about the runtime status of the PXF Service by exposing HTTP endpoints that you can use to query the health, build information, and various metrics of the running process.

Viewing PXF Status on the Command Line

The `pxf cluster status` command displays the status of the PXF Service instance on all hosts in your Greenplum Database cluster. `pxf status` displays the status of the PXF Service instance on the local Greenplum host.

Only the `gpadmin` user can request the status of the PXF Service.

Perform the following procedure to request the PXF status of your Greenplum Database cluster.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Run the `pxf cluster status` command:

```
gpadmin@gpmaster$ pxf cluster status
```

About PXF Service Runtime Monitoring

PXF exposes the following HTTP endpoints that you can use to monitor a running PXF Service on the local host:

- `actuator/health` - Returns the status of the PXF Service.
- `actuator/info` - Returns build information for the PXF Service.
- `actuator/metrics` - Returns JVM, extended Tomcat, system, process, Log4j2, and PXF-specific metrics for the PXF Service.
- `actuator/prometheus` - Returns all metrics in a format that can be scraped by a Prometheus server.

Any user can access the HTTP endpoints and view the monitoring information that PXF returns.

You can view the data associated with a specific endpoint by viewing in a browser, or `curl`-ing, a URL of the following format:

```
http://localhost:5888/<endpoint>[/<name>]
```

For example, to view the build information for the PXF service running on `localhost`, query the `actuator/info` endpoint:

```
http://localhost:5888/actuator/info
```

Sample output:

```
{"build":{"version":"6.0.0","artifact":"pxf-service","name":"pxf-service","pxfApiVersi
on":"16","group":"org.greenplum.pxf","time":"2021-03-29T22:26:22.780Z"}}
```

To view the status of the PXF Service running on the local Greenplum Database host, query the `actuator/health` endpoint:

```
http://localhost:5888/actuator/health
```

Sample output:

```
{"status":"UP","groups":["liveness","readiness"]}
```

Examining PXF Metrics

PXF exposes JVM, extended Tomcat, and system metrics via its integration with Spring Boot. Refer to Supported Metrics in the Spring Boot documentation for more information about these metrics.

PXF also exposes metrics that are specific to its processing, including:

Metric Name	Description
pxf.fragments.sent	The number of fragments, and the total time that it took to send all fragments to Greenplum Database.
pxf.records.sent	The number of records that PXF sent to Greenplum Database.
pxf.records.received	The number of records that PXF received from Greenplum Database.
pxf.bytes.sent	The number of bytes that PXF sent to Greenplum Database.
pxf.bytes.received	The number of bytes that PXF received from Greenplum Database.
http.server.requests	Standard metric augmented with PXF tags.

The information that PXF returns when you query a metric is the aggregate data collected since the last (re)start of the PXF Service.

To view a list of all of the metrics (names) available from the PXF Service, query just the `metrics` endpoint:

```
http://localhost:5888/actuator/metrics
```

Filtering Metric Data

PXF tags all metrics that it returns with an `application` label; the value of this tag is always `pxf-service`.

PXF tags its specific metrics with the additional labels: `user`, `segment`, `profile`, and `server`. All of these tags are present for each PXF metric. PXF returns the tag value `unknown` when the value cannot be determined.

You can use the tags to filter the information returned for PXF-specific metrics. For example, to examine the `pxf.records.received` metric for the PXF server named `hadoop1` located on `segment` 1 on the local host:

```
http://localhost:5888/actuator/metrics/pxf.records.received?tag=segment:1&tag=server:hadoop1
```

Certain metrics, such as `pxf.fragments.sent`, include an aditional tag named `outcome`; you can examine its value (`success` or `error`) to determine if all data for the fragment was sent. You can also use this tag to filter the aggregated data.

PXF Service Host and Port

By default, a PXF Service started on a Greenplum host listens on port number `5888` on `localhost`. You can configure PXF to start on a different port number, or use a different hostname or IP address. To change the default configuration, you will set one or both of the environment variables identified below:

Environment Variable	Description
PXF_HOST	The name of the host or IP address. The default host name is `localhost`.

Environment Variable	Description
PXF_PORT	The port number on which the PXF Service listens for requests on the host. The default port number is `5888`.

Set the environment variables in the `gpadmin` user's `.bashrc` shell login file on each Greenplum host.

You must restart both Greenplum Database and PXF when you configure the service host and/or port in this manner. Consider performing this configuration during a scheduled down time.

Procedure

Perform the following procedure to configure the PXF Service host and/or port number on one or more Greenplum Database hosts:

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. For each Greenplum Database host:

   1. Identify the host name or IP address of the PXF Service.
   2. Identify the port number on which you want the PXF Service to run.
   3. Log in to the Greenplum Database host:

      ```
      $ ssh gpadmin@<seghost>
      ```

   4. Open the `~/.bashrc` file in the editor of your choice.
   5. Set the `PXF_HOST` and/or `PXF_PORT` environment variables. For example, to set the PXF Service port number to 5998, add the following to the `.bashrc` file:

      ```
      export PXF_PORT=5998
      ```

   6. Save the file and exit the editor.

3. Restart Greenplum Database as described in Restarting Greenplum Database in the Greenplum Documentation.
4. Restart PXF on each Greenplum Database host as described in Restarting PXF.

Logging

PXF provides two categories of message logging: service-level and client-level.

PXF manages its service-level logging, and supports the following log levels (more to less severe):

- fatal
- error
- warn
- info
- debug

- trace

The default configuration for the PXF Service logs at the `info` and more severe levels. For some third-party libraries, the PXF Service logs at the `warn` or `error` and more severe levels to reduce verbosity.

- PXF captures messages written to `stdout` and `stderr` and writes them to the `$PXF_LOGDIR/pxf-app.out` file. This file may contain service start-up messages that PXF logs before logging is fully configured. The file may also contain debug output.
- Messages that PXF logs after start-up are written to the `$PXF_LOGDIR/pxf-service.log` file.

You can change the PXF log directory if you choose.

Client-level logging is managed by the Greenplum Database client; this topic details configuring logging for a `psql` client.

Enabling more verbose service- or client-level logging for PXF may aid troubleshooting efforts.

Configuring the Log Directory

The default PXF logging configuration writes log messages to `$PXF_LOGDIR`, where the default log directory is `PXF_LOGDIR=$PXF_BASE/logs`.

To change the PXF log directory, you must update the `$PXF_LOGDIR` property in the `pxf-env.sh` configuration file, synchronize the configuration change to the Greenplum Database cluster, and then restart PXF:

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Use a text editor to uncomment the `export PXF_LOGDIR` line in `$PXF_BASE/conf/pxf-env.sh`, and replace the value with the new PXF log directory. For example:

   ```
   # Path to Log directory
   export PXF_LOGDIR="/new/log/dir"
   ```

3. Use the `pxf cluster sync` command to copy the updated `pxf-env.sh` file to all hosts in the Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

4. Restart PXF on each Greenplum Database host as described in Restarting PXF.

Configuring Service-Level Logging

PXF utilizes Apache Log4j 2 for service-level logging. PXF Service-related log messages are captured in `$PXF_LOGDIR/pxf-app.out` and `$PXF_LOGDIR/pxf-service.log`. The default configuration for the PXF Service logs at the `info` and more severe levels.

You can change the log level for the PXF Service on a single Greenplum Database host, or on all hosts in the Greenplum cluster.

PXF provides more detailed logging when the `debug` and `trace` log levels are enabled. Logging at these levels is quite verbose, and has both a performance and a storage impact. Be sure to turn it off after you have collected the desired information.

Configuring for a Specific Host

You can change the log level for the PXF Service running on a specific Greenplum Database host in two ways:

- Setting the `PXF_LOG_LEVEL` environment variable on the `pxf restart` command line.
- Setting the log level via a property update.

Procedure:

1. Log in to the Greenplum Database host:

   ```
   $ ssh gpadmin@<gphost>
   ```

2. Choose one of the following methods:

   - Set the log level on the `pxf restart` command line. For example, to change the log level from `info` (the default) to `debug`:

     ```
     gpadmin@gphost$ PXF_LOG_LEVEL=debug pxf restart
     ```

   - Set the log level in the `pxf-application.properties` file:

     1. Use a text editor to uncomment the following line in the `$PXF_BASE/conf/pxf-application.properties` file and set the desired log level. For example, to change the log level from `info` (the default) to `debug`:

        ```
        pxf.log.level=debug
        ```

     2. Restart PXF on the host:

        ```
        gpadmin@gphost$ pxf restart
        ```

3. `debug` logging is now enabled. Make note of the time; this will direct you to the relevant log messages in `$PXF_LOGDIR/pxf-service.log`.

   ```
   $ date
   Wed Oct  4 09:30:06 MDT 2017
   $ psql -d <dbname>
   ```

4. Perform operations that exercise the PXF Service.

5. Collect and examine the log messages in `pxf-service.log`.

6. Depending upon how you originally set the log level, reinstate `info`-level logging on the host:

   - Command line method:

     ```
     gpadmin@gphost$ pxf restart
     ```

- Properties file method: Comment out the line or set the property value back to `info`, and then restart PXF on the host.

Configuring for the Cluster

To change the log level for the PXF service running on every host in the Greenplum Database cluster:

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Use a text editor to uncomment the following line in the `$PXF_BASE/conf/pxf-application.properties` file and set the desired log level. For example, to change the log level from `info` (the default) to `debug`:

```
pxf.log.level=debug
```

3. Use the `pxf cluster sync` command to copy the updated `pxf-application.properties` file to all hosts in the Greenplum Database cluster. For example:

```
gpadmin@gpmaster$ pxf cluster sync
```

4. Restart PXF on each Greenplum Database host:

```
gpadmin@gpmaster$ pxf cluster restart
```

5. Perform operations that exercise the PXF Service, and then collect and examine the information in `$PXF_LOGDIR/pxf-service.log`.

6. Reinstate `info`-level logging - comment out the line or set the property value back to `info`, and then sync the update to and restart PXF on the cluster.

Configuring Client-Level Logging

Database-level client session logging may provide insight into internal PXF Service operations.

Enable Greenplum Database client debug message logging by setting the `client_min_messages` server configuration parameter to `DEBUG2` in your `psql` session. This logging configuration writes messages to `stdout`, and will apply to all operations that you perform in the session, including operations on PXF external tables. For example:

```
$ psql -d <dbname>
```

```
dbname=# SET client_min_messages=DEBUG2;
dbname=# SELECT * FROM hdfstest;
...
DEBUG2:  churl http header: cell #26: X-GP-URL-HOST: localhost  (seg0 slice1 127.0.0.1
:7002 pid=10659)
CONTEXT:  External table pxf_hdfs_textsimple, line 1 of file pxf://data/pxf_examples/p
xf_hdfs_simple.txt?PROFILE=hdfs:text
DEBUG2:  churl http header: cell #27: X-GP-URL-PORT: 5888  (seg0 slice1 127.0.0.1:7002
 pid=10659)
```

```
CONTEXT:  External table pxf_hdfs_textsimple, line 1 of file pxf://data/pxf_examples/p
xf_hdfs_simple.txt?PROFILE=hdfs:text
DEBUG2:  churl http header: cell #28: X-GP-DATA-DIR: data%2Fpxf_examples%2Fpxf_hdfs_si
mple.txt  (seg0 slice1 127.0.0.1:7002 pid=10659)
CONTEXT:  External table pxf_hdfs_textsimple, line 1 of file pxf://data/pxf_examples/p
xf_hdfs_simple.txt?PROFILE=hdfs:text
DEBUG2:  churl http header: cell #29: X-GP-TABLE-NAME: pxf_hdfs_textsimple  (seg0 slic
e1 127.0.0.1:7002 pid=10659)
CONTEXT:  External table pxf_hdfs_textsimple, line 1 of file pxf://data/pxf_examples/p
xf_hdfs_simple.txt?PROFILE=hdfs:text
...
```

Collect and examine the log messages written to `stdout`.

Note: `DEBUG2` database client session logging has a performance impact. Remember to turn off `DEBUG2` logging after you have collected the desired information.

```
dbname=# SET client_min_messages=NOTICE;
```

Memory and Threading

Because a single PXF Service (JVM) serves multiple segments on a segment host, the PXF heap size can be a limiting runtime factor. This becomes more evident under concurrent workloads or with queries against large files. You may run into situations where a query hangs or fails due to insufficient memory or the Java garbage collector impacting response times. To avert or remedy these situations, first try increasing the Java maximum heap size or decreasing the Tomcat maximum number of threads, depending upon what works best for your system configuration. You may also choose to configure PXF to perform specific actions when it detects an out of memory condition.

Increasing the JVM Memory for PXF

Each PXF Service running on a Greenplum Database host is configured with a default maximum Java heap size of 2GB and an initial heap size of 1GB. If the hosts in your Greenplum Database cluster have an ample amount of memory, try increasing the maximum heap size to a value between 3-4GB. Set the initial and maximum heap size to the same value if possible.

Perform the following procedure to increase the heap size for the PXF Service running on each host in your Greenplum Database cluster.

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Edit the `$PXF_BASE/conf/pxf-env.sh` file. For example:

   ```
   gpadmin@gpmaster$ vi $PXF_BASE/conf/pxf-env.sh
   ```

3. Locate the `PXF_JVM_OPTS` setting in the `pxf-env.sh` file, and update the `-Xmx` and/or `-Xms` options to the desired value. For example:

   ```
   PXF_JVM_OPTS="-Xmx3g -Xms3g"
   ```

4. Save the file and exit the editor.

5. Use the `pxf cluster sync` command to copy the updated `pxf-env.sh` file to the Greenplum Database cluster. For example:

```
gpadmin@gpmaster$ pxf cluster sync
```

6. Restart PXF on each Greenplum Database host as described in Restarting PXF.

Configuring Out of Memory Condition Actions

In an out of memory (OOM) situation, PXF returns the following error in response to a query:

```
java.lang.OutOfMemoryError: Java heap space
```

You can configure the PXF JVM to enable/disable the following actions when it detects an OOM condition:

- Auto-kill the PXF Service (enabled by default).
- Dump the Java heap (disabled by default).

Auto-Killing the PXF Server

By default, PXF is configured such that when the PXF JVM detects an out of memory condition on a Greenplum host, it automatically runs a script that kills the PXF Service running on the host. The `PXF_OOM_KILL` environment variable in the `$PXF_BASE/conf/pxf-env.sh` configuration file governs this auto-kill behavior.

When auto-kill is enabled and the PXF JVM detects an OOM condition and kills the PXF Service on the host:

- PXF logs the following messages to `$PXF_LOGDIR/pxf-oom.log` on the segment host:

  ```
  =====> <date> PXF Out of memory detected <======
  =====> <date> PXF shutdown scheduled <======
  =====> <date> Stopping PXF <======
  ```

- Any query that you run on a PXF external table will fail with the following error until you restart the PXF Service on the host:

  ```
  ... Failed to connect to <host> port 5888: Connection refused
  ```

When the PXF Service on a host is shut down in this manner, you must explicitly restart the PXF Service on the host. See the pxf reference page for more information on the `pxf start` command.

Refer to the configuration procedure below for the instructions to disable/enable this PXF configuration property.

Dumping the Java Heap

In an out of memory situation, it may be useful to capture the Java heap dump to help determine what factors contributed to the resource exhaustion. You can configure PXF to write the heap dump to a file when it detects an OOM condition by setting the `PXF_OOM_DUMP_PATH` environment variable in the `$PXF_BASE/conf/pxf-env.sh` configuration file. By default, PXF does not dump the Java heap on OOM.

If you choose to enable the heap dump on OOM, you must set PXF_OOM_DUMP_PATH to the absolute path to a file or directory:

- If you specify a directory, the PXF JVM writes the heap dump to the file <directory>/java_pid<pid>.hprof, where <pid> identifies the process ID of the PXF Service instance. The PXF JVM writes a new file to the directory every time the JVM goes OOM.
- If you specify a file and the file does not exist, the PXF JVM writes the heap dump to the file when it detects an OOM. If the file already exists, the JVM will not dump the heap.

Ensure that the gpadmin user has write access to the dump file or directory.

Note: Heap dump files are often rather large. If you enable heap dump on OOM for PXF and specify a directory for PXF_OOM_DUMP_PATH, multiple OOMs will generate multiple files in the directory and could potentially consume a large amount of disk space. If you specify a file for PXF_OOM_DUMP_PATH, disk usage is constant when the file name does not change. You must rename the dump file or configure a different PXF_OOM_DUMP_PATH to generate subsequent heap dumps.

Refer to the configuration procedure below for the instructions to enable/disable this PXF configuration property.

Procedure

Auto-kill of the PXF Service on OOM is enabled by default. Heap dump generation on OOM is disabled by default. To configure one or both of these properties, perform the following procedure:

1. Log in to your Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Edit the $PXF_BASE/conf/pxf-env.sh file. For example:

   ```
   gpadmin@gpmaster$ vi $PXF_BASE/conf/pxf-env.sh
   ```

3. If you want to configure (i.e. turn off, or turn back on) auto-kill of the PXF Service on OOM, locate the PXF_OOM_KILL property in the pxf-env.sh file. If the setting is commented out, uncomment it, and then update the value. For example, to turn off this behavior, set the value to false:

   ```
   export PXF_OOM_KILL=false
   ```

4. If you want to configure (i.e. turn on, or turn back off) automatic heap dumping when the PXF Service hits an OOM condition, locate the PXF_OOM_DUMP_PATH setting in the pxf-env.sh file.

   1. To turn this behavior on, set the PXF_OOM_DUMP_PATH property value to the file system location to which you want the PXF JVM to dump the Java heap. For example, to dump to a file named /home/gpadmin/pxfoom_segh1:

      ```
      export PXF_OOM_DUMP_PATH=/home/pxfoom_segh1
      ```

   2. To turn off heap dumping after you have turned it on, comment out the PXF_OOM_DUMP_PATH property setting:

```
#export PXF_OOM_DUMP_PATH=/home/pxfoom_segh1
```

5. Save the `pxf-env.sh` file and exit the editor.
6. Use the `pxf cluster sync` command to copy the updated `pxf-env.sh` file to the Greenplum Database cluster. For example:

```
gpadmin@gpmaster$ pxf cluster sync
```

7. Restart PXF on each Greenplum Database host as described in Restarting PXF.

Another Option for Resource-Constrained PXF Segment Hosts

If increasing the maximum heap size is not suitable for your Greenplum Database deployment, try decreasing the number of concurrent working threads configured for PXF's embedded Tomcat web server. A decrease in the number of running threads will prevent any PXF node from exhausting its memory, while ensuring that current queries run to completion (albeit a bit slower). Tomcat's default behavior is to queue requests until a thread is free, or the queue is exhausted.

The default maximum number of Tomcat threads for PXF is 200. The `pxf.max.threads` property in the `pxf-application.properties` configuration file controls this setting.

If you plan to run large workloads on a large number of files in an external Hive data store, or you are reading compressed ORC or Parquet data, consider specifying a lower `pxf.max.threads` value. Large workloads require more memory, and a lower thread count limits concurrency, and hence, memory consumption.

Note: Keep in mind that an increase in the thread count correlates with an increase in memory consumption.

Perform the following procedure to set the maximum number of Tomcat threads for the PXF Service running on each host in your Greenplum Database deployment.

1. Log in to your Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Edit the `$PXF_BASE/conf/pxf-application.properties` file. For example:

```
gpadmin@gpmaster$ vi $PXF_BASE/conf/pxf-application.properties
```

3. Locate the `pxf.max.threads` setting in the `pxf-application.properties` file. If the setting is commented out, uncomment it, and then update to the desired value. For example, to reduce the maximum number of Tomcat threads to 100:

```
pxf.max.threads=100
```

4. Save the file and exit the editor.
5. Use the `pxf cluster sync` command to copy the updated `pxf-application.properties` file to the Greenplum Database cluster. For example:

```
gpadmin@gpmaster$ pxf cluster sync
```

6. Restart PXF on each Greenplum Database host as described in Restarting PXF.

Accessing Hadoop with PXF

PXF is compatible with Cloudera, Hortonworks Data Platform, MapR, and generic Apache Hadoop distributions. PXF is installed with HDFS, Hive, and HBase connectors. You use these connectors to access varied formats of data from these Hadoop distributions.

Architecture

HDFS is the primary distributed storage mechanism used by Apache Hadoop. When a user or application performs a query on a PXF external table that references an HDFS file, the Greenplum Database master node dispatches the query to all segment instances. Each segment instance contacts the PXF Service running on its host. When it receives the request from a segment instance, the PXF Service:

1. Allocates a worker thread to serve the request from the segment instance.
2. Invokes the HDFS Java API to request metadata information for the HDFS file from the HDFS NameNode.

Figure: PXF-to-Hadoop Architecture

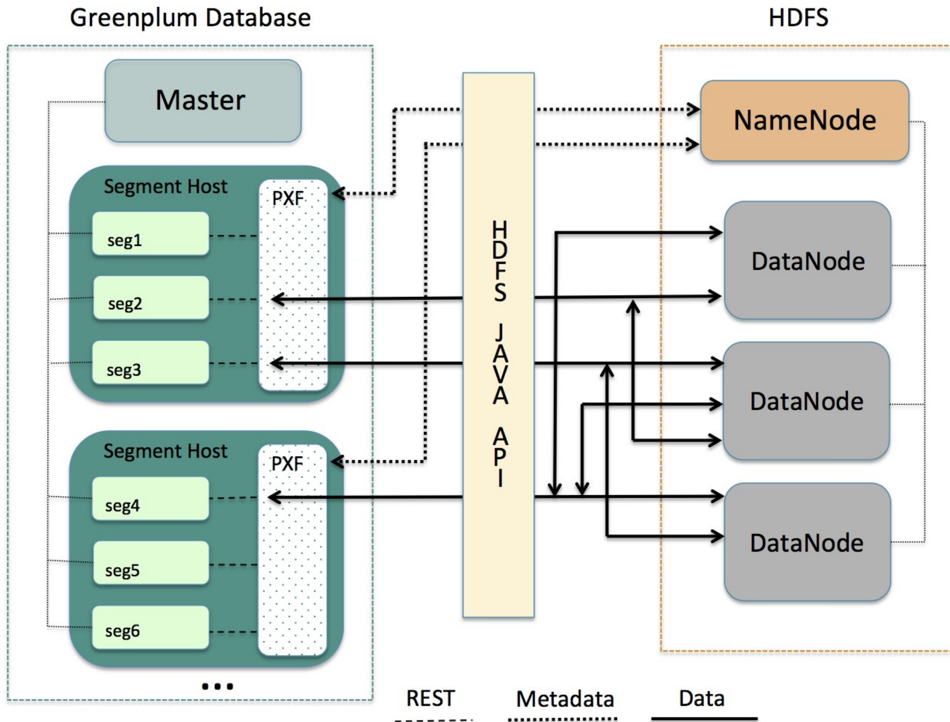

A PXF worker thread works on behalf of a segment instance. A worker thread uses its Greenplum Database `gp_segment_id` and the file block information described in the metadata to assign itself a specific portion of the query data. This data may reside on one or more HDFS DataNodes.

The PXF worker thread invokes the HDFS Java API to read the data and delivers it to the segment instance. The segment instance delivers its portion of the data to the Greenplum Database master node. This communication occurs across segment hosts and segment instances in parallel.

Prerequisites

Before working with Hadoop data using PXF, ensure that:

- You have configured PXF, and PXF is running on each Greenplum Database host. See Configuring PXF for additional information.
- You have configured the PXF Hadoop Connectors that you plan to use. Refer to Configuring PXF Hadoop Connectors for instructions. If you plan to access JSON-formatted data stored in a Cloudera Hadoop cluster, PXF requires a Cloudera version 5.8 or later Hadoop distribution.
- If user impersonation is enabled (the default), ensure that you have granted read (and write as appropriate) permission to the HDFS files and directories that will be accessed as external tables in Greenplum Database to each Greenplum Database user/role name that will access the HDFS files and directories. If user impersonation is not enabled, you must grant this permission to the `gpadmin` user.
- Time is synchronized between the Greenplum Database hosts and the external Hadoop systems.

HDFS Shell Command Primer

Examples in the PXF Hadoop topics access files on HDFS. You can choose to access files that already exist in your HDFS cluster. Or, you can follow the steps in the examples to create new files.

A Hadoop installation includes command-line tools that interact directly with your HDFS file system. These tools support typical file system operations that include copying and listing files, changing file permissions, and so forth. You run these tools on a system with a Hadoop client installation. By default, Greenplum Database hosts do not include a Hadoop client installation.

The HDFS file system command syntax is `hdfs dfs <options> [<file>]`. Invoked with no options, `hdfs dfs` lists the file system options supported by the tool.

The user invoking the `hdfs dfs` command must have read privileges on the HDFS data store to list and view directory and file contents, and write permission to create directories and files.

The `hdfs dfs` options used in the PXF Hadoop topics are:

Option	Description
`-cat`	Display file contents.
`-mkdir`	Create a directory in HDFS.
`-put`	Copy a file from the local file system to HDFS.

Examples:

Create a directory in HDFS:

```
$ hdfs dfs -mkdir -p /data/exampledir
```

Copy a text file from your local file system to HDFS:

```
$ hdfs dfs -put /tmp/example.txt /data/exampledir/
```

Display the contents of a text file located in HDFS:

```
$ hdfs dfs -cat /data/exampledir/example.txt
```

Connectors, Data Formats, and Profiles

The PXF Hadoop connectors provide built-in profiles to support the following data formats:

- Text
- CSV
- Avro
- JSON
- ORC
- Parquet
- RCFile

- SequenceFile
- AvroSequenceFile

The PXF Hadoop connectors expose the following profiles to read, and in many cases write, these supported data formats:

Data Source	Data Format	Profile Name(s)	Deprecated Profile Name	Supported Operations
HDFS	delimited single line text	hdfs:text	n/a	Read, Write
HDFS	delimited single line comma-separated values of text	hdfs:csv	n/a	Read, Write
HDFS	delimited text with quoted linefeeds	hdfs:text:multi	n/a	Read
HDFS	Avro	hdfs:avro	n/a	Read, Write
HDFS	JSON	hdfs:json	n/a	Read
HDFS	ORC	hdfs:orc	n/a	Read
HDFS	Parquet	hdfs:parquet	n/a	Read, Write
HDFS	AvroSequenceFile	hdfs:AvroSequenceFile	n/a	Read, Write
HDFS	SequenceFile	hdfs:SequenceFile	n/a	Read, Write
Hive	stored as TextFile	hive, hive:text	Hive, HiveText	Read
Hive	stored as SequenceFile	hive	Hive	Read
Hive	stored as RCFile	hive, hive:rc	Hive, HiveRC	Read
Hive	stored as ORC	hive, hive:orc	Hive, HiveORC, HiveVectorizedORC	Read
Hive	stored as Parquet	hive	Hive	Read
Hive	stored as Avro	hive	Hive	Read
HBase	Any	hbase	HBase	Read

Choosing the Profile

PXF provides more than one profile to access text and Parquet data on Hadoop. Here are some things to consider as you determine which profile to choose.

Choose the `hive` profile when:

- The data resides in a Hive table, and you do not know the underlying file type of the table up front.
- The data resides in a Hive table, and the Hive table is partitioned.

Choose the `hdfs:text`, `hdfs:csv` profiles when the file is text and you know the location of the file in the HDFS file system.

When accessing ORC-format data:

- Choose the `hdfs:orc` profile when the file is ORC, you know the location of the file in the HDFS file system, and the file is not managed by Hive or you do not want to use the Hive

Metastore.

- Choose the `hive:orc` profile when the table is ORC and the table is managed by Hive, and the data is partitioned or the data includes complex types.

Choose the `hdfs:parquet` profile when the file is Parquet, you know the location of the file in the HDFS file system, and you want to take advantage of extended filter pushdown support for additional data types and operators.

Specifying the Profile

You must provide the profile name when you specify the `pxf` protocol in a `CREATE EXTERNAL TABLE` command to create a Greenplum Database external table that references a Hadoop file or directory, HBase table, or Hive table. For example, the following command creates an external table that uses the default server and specifies the profile named `hdfs:text` to access the HDFS file `/data/pxf_examples/pxf_hdfs_simple.txt`:

```
CREATE EXTERNAL TABLE pxf_hdfs_text(location text, month text, num_orders int, total_sales float8)
   LOCATION ('pxf://data/pxf_examples/pxf_hdfs_simple.txt?PROFILE=hdfs:text')
FORMAT 'TEXT' (delimiter=E',');
```

Reading and Writing HDFS Text Data

The PXF HDFS Connector supports plain delimited and comma-separated value form text data. This section describes how to use PXF to access HDFS text data, including how to create, query, and insert data into an external table that references files in the HDFS data store.

PXF supports reading or writing text files compressed with the `default`, `bzip2`, and `gzip` codecs.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from or write data to HDFS.

Reading Text Data

Use the `hdfs:text` profile when you read plain text delimited, and `hdfs:csv` when reading .csv data where each row is a single record. The following syntax creates a Greenplum Database readable external table that references such a text file on HDFS:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-file>?PROFILE=hdfs:text|csv[&SERVER=<server_name>][&IGNORE_MISSING_PATH=<boolean>][&SKIP_HEADER_COUNT=<numlines>]')
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value

<path-to-hdfs-file>	The path to the directory or file in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	Use `PROFILE hdfs:text` when <path-to-hdfs-file> references plain text delimited data. Use `PROFILE hdfs:csv` when <path-to-hdfs-file> references comma-separated value data.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-hdfs-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
SKIP_HEADER_COUNT=<numlines>	Specify the number of header lines that PXF should skip in the first split of each <hdfs-file> before reading the data. The default value is 0, do not skip any lines.
FORMAT	Use `FORMAT 'TEXT'` when <path-to-hdfs-file> references plain text delimited data. Use `FORMAT 'CSV'` when <path-to-hdfs-file> references comma-separated value data.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.

Note: PXF does not support the `(HEADER)` formatter option in the `CREATE EXTERNAL TABLE` command. If your text file includes header line(s), use `SKIP_HEADER_COUNT` to specify the number of lines that PXF should skip at the beginning of the first split of each file.

Example: Reading Text Data on HDFS

Perform the following procedure to create a sample text file, copy the file to HDFS, and use the `hdfs:text` and `hdfs:csv` profiles and the default PXF server to create two PXF external tables to query the data:

1. Create an HDFS directory for PXF example data files. For example:

   ```
   $ hdfs dfs -mkdir -p /data/pxf_examples
   ```

2. Create a delimited plain text data file named `pxf_hdfs_simple.txt`:

   ```
   $ echo 'Prague,Jan,101,4875.33
   Rome,Mar,87,1557.39
   Bangalore,May,317,8936.99
   Beijing,Jul,411,11600.67' > /tmp/pxf_hdfs_simple.txt
   ```

   Note the use of the comma `,` to separate the four data fields.

3. Add the data file to HDFS:

   ```
   $ hdfs dfs -put /tmp/pxf_hdfs_simple.txt /data/pxf_examples/
   ```

4. Display the contents of the `pxf_hdfs_simple.txt` file stored in HDFS:

   ```
   $ hdfs dfs -cat /data/pxf_examples/pxf_hdfs_simple.txt
   ```

5. Start the `psql` subsystem:

```
$ psql -d postgres
```

6. Use the PXF `hdfs:text` profile to create a Greenplum Database external table that references the `pxf_hdfs_simple.txt` file that you just created and added to HDFS:

```
postgres=# CREATE EXTERNAL TABLE pxf_hdfs_textsimple(location text, month text,
 num_orders int, total_sales float8)
            LOCATION ('pxf://data/pxf_examples/pxf_hdfs_simple.txt?PROFILE=hdfs
:text')
          FORMAT 'TEXT' (delimiter=E',');
```

7. Query the external table:

```
postgres=# SELECT * FROM pxf_hdfs_textsimple;
```

```
   location    | month | num_orders | total_sales
---------------+-------+------------+-------------
 Prague        | Jan   |        101 |     4875.33
 Rome          | Mar   |         87 |     1557.39
 Bangalore     | May   |        317 |     8936.99
 Beijing       | Jul   |        411 |    11600.67
(4 rows)
```

8. Create a second external table that references `pxf_hdfs_simple.txt`, this time specifying the `hdfs:csv` `PROFILE` and the `CSV` `FORMAT`:

```
postgres=# CREATE EXTERNAL TABLE pxf_hdfs_textsimple_csv(location text, month t
ext, num_orders int, total_sales float8)
            LOCATION ('pxf://data/pxf_examples/pxf_hdfs_simple.txt?PROFILE=hdfs
:csv')
          FORMAT 'CSV';
postgres=# SELECT * FROM pxf_hdfs_textsimple_csv;
```

When you specify `FORMAT 'CSV'` for comma-separated value data, no `delimiter` formatter option is required because comma is the default delimiter value.

Reading Text Data with Quoted Linefeeds

Use the `hdfs:text:multi` profile to read plain text data with delimited single- or multi- line records that include embedded (quoted) linefeed characters. The following syntax creates a Greenplum Database readable external table that references such a text file on HDFS:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-file>?PROFILE=hdfs:text:multi[&SERVER=<server_name>][&I
GNORE_MISSING_PATH=<boolean>][&SKIP_HEADER_COUNT=<numlines>]')
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
```

The specific keywords and values used in the CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-hdfs-file>	The path to the directory or file in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:text:multi`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-hdfs-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
SKIP_HEADER_COUNT=<numlines>	Specify the number of header lines that PXF should skip in the first split of each <hdfs-file> before reading the data. The default value is 0, do not skip any lines.
FORMAT	Use `FORMAT 'TEXT'` when <path-to-hdfs-file> references plain text delimited data. Use `FORMAT 'CSV'` when <path-to-hdfs-file> references comma-separated value data.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.

Note: PXF does not support the `(HEADER)` formatter option in the `CREATE EXTERNAL TABLE` command. If your text file includes header line(s), use `SKIP_HEADER_COUNT` to specify the number of lines that PXF should skip at the beginning of the first split of each file.

Example: Reading Multi-Line Text Data on HDFS

Perform the following steps to create a sample text file, copy the file to HDFS, and use the PXF `hdfs:text:multi` profile and the default PXF server to create a Greenplum Database readable external table to query the data:

1. Create a second delimited plain text file:

```
$ vi /tmp/pxf_hdfs_multi.txt
```

2. Copy/paste the following data into `pxf_hdfs_multi.txt`:

```
"4627 Star Rd.
San Francisco, CA  94107":Sept:2017
"113 Moon St.
San Diego, CA  92093":Jan:2018
"51 Belt Ct.
Denver, CO  90123":Dec:2016
"93114 Radial Rd.
Chicago, IL  60605":Jul:2017
"7301 Brookview Ave.
Columbus, OH  43213":Dec:2018
```

Notice the use of the colon `:` to separate the three fields. Also notice the quotes around the first (address) field. This field includes an embedded line feed separating the street address from the city and state.

3. Copy the text file to HDFS:

```
$ hdfs dfs -put /tmp/pxf_hdfs_multi.txt /data/pxf_examples/
```

4. Use the `hdfs:text:multi` profile to create an external table that references the `pxf_hdfs_multi.txt` HDFS file, making sure to identify the `:` (colon) as the field separator:

```
postgres=# CREATE EXTERNAL TABLE pxf_hdfs_textmulti(address text, month text, y
ear int)
            LOCATION ('pxf://data/pxf_examples/pxf_hdfs_multi.txt?PROFILE=hdfs:
text:multi')
          FORMAT 'CSV' (delimiter ':');
```

   Notice the alternate syntax for specifying the `delimiter`.

5. Query the `pxf_hdfs_textmulti` table:

```
postgres=# SELECT * FROM pxf_hdfs_textmulti;
```

```
         address          | month | year
--------------------------+-------+------
 4627 Star Rd.            | Sept  | 2017
 San Francisco, CA  94107
 113 Moon St.             | Jan   | 2018
 San Diego, CA  92093
 51 Belt Ct.              | Dec   | 2016
 Denver, CO  90123
 93114 Radial Rd.         | Jul   | 2017
 Chicago, IL  60605
 7301 Brookview Ave.      | Dec   | 2018
 Columbus, OH  43213
(5 rows)
```

Writing Text Data to HDFS

The PXF HDFS connector profiles `hdfs:text` and `hdfs:csv` support writing single line plain text data to HDFS. When you create a writable external table with the PXF HDFS connector, you specify the name of a directory on HDFS. When you insert records into a writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specified.

Note: External tables that you create with a writable profile can only be used for `INSERT` operations. If you want to query the data that you inserted, you must create a separate readable external table that references the HDFS directory.

Use the following syntax to create a Greenplum Database writable external table that references an HDFS directory:

```
CREATE WRITABLE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-dir>
    ?PROFILE=hdfs:text|csv[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the CREATE EXTERNAL TABLE command are described

in the table below.

Keyword	Value
<path-to-hdfs-dir>	The path to the directory in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-dir> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-dir> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	Use `PROFILE hdfs:text` to write plain, delimited text to <path-to-hdfs-file>. Use `PROFILE hdfs:csv` to write comma-separated value text to <path-to-hdfs-dir>.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are described below.
FORMAT	Use `FORMAT 'TEXT'` to write plain, delimited text to <path-to-hdfs-dir>. Use `FORMAT 'CSV'` to write comma-separated value text to <path-to-hdfs-dir>.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

Writable external tables that you create using the `hdfs:text` or the `hdfs:csv` profiles can optionally use record or block compression. You specify the compression type and codec via custom options in the `CREATE EXTERNAL TABLE LOCATION` clause. The `hdfs:text` and `hdfs:csv` profiles support the following custom write options:

Option	Value Description
COMPRESSION_CODEC	The compression codec alias. Supported compression codecs for writing text data include: `default`, `bzip2`, `gzip`, and `uncompressed`. If this option is not provided, Greenplum Database performs no data compression.
COMPRESSION_TYPE	The compression type to employ; supported values are `RECORD` (the default) or `BLOCK`.

Example: Writing Text Data to HDFS

This example utilizes the data schema introduced in Example: Reading Text Data on HDFS.

Column Name	Data Type
location	text
month	text
number_of_orders	int
total_sales	float8

This example also optionally uses the Greenplum Database external table named `pxf_hdfs_textsimple` that you created in that exercise.

Procedure

Perform the following procedure to create Greenplum Database writable external tables utilizing the same data schema as described above, one of which will employ compression. You will use the PXF `hdfs:text` profile and the default PXF server to write data to the underlying HDFS directory. You will also create a separate, readable external table to read the data that you wrote to the HDFS directory.

1. Create a Greenplum Database writable external table utilizing the data schema described above. Write to the HDFS directory `/data/pxf_examples/pxfwritable_hdfs_textsimple1`. Create the table specifying a comma `,` as the delimiter:

    ```
    postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_hdfs_writabletbl_1(location text,
     month text, num_orders int, total_sales float8)
                LOCATION ('pxf://data/pxf_examples/pxfwritable_hdfs_textsimple1?PRO
    FILE=hdfs:text')
              FORMAT 'TEXT' (delimiter=',');
    ```

    You specify the `FORMAT` subclause `delimiter` value as the single ascii comma character `,`.

2. Write a few individual records to the `pxfwritable_hdfs_textsimple1` HDFS directory by invoking the SQL `INSERT` command on `pxf_hdfs_writabletbl_1`:

    ```
    postgres=# INSERT INTO pxf_hdfs_writabletbl_1 VALUES ( 'Frankfurt', 'Mar', 777,
     3956.98 );
    postgres=# INSERT INTO pxf_hdfs_writabletbl_1 VALUES ( 'Cleveland', 'Oct', 3812
    , 96645.37 );
    ```

3. (Optional) Insert the data from the `pxf_hdfs_textsimple` table that you created in Example: Reading Text Data on HDFS into `pxf_hdfs_writabletbl_1`:

    ```
    postgres=# INSERT INTO pxf_hdfs_writabletbl_1 SELECT * FROM pxf_hdfs_textsimple
    ;
    ```

4. In another terminal window, display the data that you just added to HDFS:

    ```
    $ hdfs dfs -cat /data/pxf_examples/pxfwritable_hdfs_textsimple1/*
    Frankfurt,Mar,777,3956.98
    Cleveland,Oct,3812,96645.37
    Prague,Jan,101,4875.33
    Rome,Mar,87,1557.39
    Bangalore,May,317,8936.99
    Beijing,Jul,411,11600.67
    ```

    Because you specified comma `,` as the delimiter when you created the writable external table, this character is the field separator used in each record of the HDFS data.

5. Greenplum Database does not support directly querying a writable external table. To query the data that you just added to HDFS, you must create a readable external Greenplum Database table that references the HDFS directory:

    ```
    postgres=# CREATE EXTERNAL TABLE pxf_hdfs_textsimple_r1(location text, month te
    xt, num_orders int, total_sales float8)
                LOCATION ('pxf://data/pxf_examples/pxfwritable_hdfs_textsimple1?PRO
    FILE=hdfs:text')
                FORMAT 'CSV';
    ```

    You specify the `'CSV'` `FORMAT` when you create the readable external table because you

created the writable table with a comma `,` as the delimiter character, the default delimiter for `'CSV' FORMAT`.

6. Query the readable external table:

```
postgres=# SELECT * FROM pxf_hdfs_textsimple_r1 ORDER BY total_sales;
```

```
 location  | month | num_orders | total_sales
-----------+-------+------------+-------------
 Rome      | Mar   |         87 |     1557.39
 Frankfurt | Mar   |        777 |     3956.98
 Prague    | Jan   |        101 |     4875.33
 Bangalore | May   |        317 |     8936.99
 Beijing   | Jul   |        411 |    11600.67
 Cleveland | Oct   |       3812 |    96645.37
(6 rows)
```

The `pxf_hdfs_textsimple_r1` table includes the records you individually inserted, as well as the full contents of the `pxf_hdfs_textsimple` table if you performed the optional step.

7. Create a second Greenplum Database writable external table, this time using Gzip compression and employing a colon `:` as the delimiter:

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_hdfs_writabletbl_2 (location text
, month text, num_orders int, total_sales float8)
            LOCATION ('pxf://data/pxf_examples/pxfwritable_hdfs_textsimple2?PRO
FILE=hdfs:text&COMPRESSION_CODEC=gzip')
          FORMAT 'TEXT' (delimiter=':');
```

8. Write a few records to the `pxfwritable_hdfs_textsimple2` HDFS directory by inserting directly into the `pxf_hdfs_writabletbl_2` table:

```
gpadmin=# INSERT INTO pxf_hdfs_writabletbl_2 VALUES ( 'Frankfurt', 'Mar', 777,
3956.98 );
gpadmin=# INSERT INTO pxf_hdfs_writabletbl_2 VALUES ( 'Cleveland', 'Oct', 3812,
 96645.37 );
```

9. In another terminal window, display the contents of the data that you added to HDFS; use the `-text` option to `hdfs dfs` to view the compressed data as text:

```
$ hdfs dfs -text /data/pxf_examples/pxfwritable_hdfs_textsimple2/*
Frankfurt:Mar:777:3956.98
Cleveland:Oct:3812:96645.3
```

Notice that the colon `:` is the field separator in this HDFS data.

To query data from the newly-created HDFS directory named `pxfwritable_hdfs_textsimple2`, you can create a readable external Greenplum Database table as described above that references this HDFS directory and specifies `FORMAT 'CSV' (delimiter=':')`.

Reading and Writing HDFS Avro Data

Use the PXF HDFS Connector to read and write Avro-format data. This section describes how to use

PXF to read and write Avro data in HDFS, including how to create, query, and insert into an external table that references an Avro file in the HDFS data store.

PXF supports reading or writing Avro files compressed with these codecs: `bzip2`, `xz`, `snappy`, and `deflate`.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from HDFS.

Working with Avro Data

Apache Avro is a data serialization framework where the data is serialized in a compact binary format. Avro specifies that data types be defined in JSON. Avro format data has an independent schema, also defined in JSON. An Avro schema, together with its data, is fully self-describing.

Data Type Mapping

The Avro specification defines primitive, complex, and logical types.

To represent Avro primitive data types and Avro arrays of primitive types in Greenplum Database, map data values to Greenplum Database columns of the same type.

Avro supports other complex data types including arrays of non-primitive types, maps, records, enumerations, and fixed types. Map top-level fields of these complex data types to the Greenplum Database `text` type. While PXF does not natively support reading these types, you can create Greenplum Database functions or application code to extract or further process subcomponents of these complex data types.

Avro supports logical data types including decimal, date, time, and duration types. You must similarly map these data types to the Greenplum Database `text` type.

Read Mapping

PXF uses the following data type mapping when reading Avro data:

Avro Data Type	PXF/Greenplum Data Type
boolean	boolean
bytes	bytea
double	double
float	real
int	int
long	bigint
string	text
Complex type: Array (any dimension) of type: boolean, bytes, double, float, int, long, string	array (any dimension) of type: boolean, bytea, double, real, bigint, text

Avro Data Type	PXF/Greenplum Data Type
Complex type: Array of other types (*Avro schema is provided*)	text[]
Complex type: Map, Record, or Enum	text, with delimiters inserted between collection items, mapped key-value pairs, and record data.
Complex type: Fixed	bytea (supported for read operations only).
Union	Follows the above conventions for primitive or complex data types, depending on the union; must contain 2 elements, one of which must be null.
Logical type: date	int
Logical type: time-millis, timestamp-millis, or local-timestamp-millis	int
Logical type: time-micros, timestamp-micros, or local-timestamp-micros	long
Logical type: duration	bytea

Write Mapping

PXF supports writing Avro primitive types and arrays of Avro primitive types. PXF supports writing other complex types to Avro as string.

PXF uses the following data type mapping when writing Avro data:

PXF/Greenplum Data Type	Avro Data Type
bigint	long
boolean	boolean
bytea	bytes
double	double
char[1]	string
enum	string
int	int
real	float
smallint[2]	int
text	string
varchar	string
numeric, date, time, timestamp, timestamptz (*no Avro schema is provided*)	string
array (any dimension) of type: bigint, boolean, bytea, double, int, real, text (*Avro schema is provided*)	Array (any dimension) of type: long, boolean, bytes, double, int, float, string

PXF/Greenplum Data Type	Avro Data Type
bigint[], boolean[], bytea[], double[], int[], real[], text[] (*no Avro schema is provided*)	long[], boolean[], bytes[], double[], int[], float[], string[] (one-dimensional array)
numeric[], date[], time[], timestamp[], timestamptz[] (*Avro is schema is provided*)	string[]
enum, record	string

[1] PXF right-pads `char[n]` types to length *n*, if required, with white space.

[2] PXF converts Greenplum `smallint` types to `int` before it writes the Avro data. Be sure to read the field into an `int`.

Avro Schemas and Data

Avro schemas are defined using JSON, and composed of the same primitive and complex types identified in the data type mapping section above. Avro schema files typically have a `.avsc` suffix.

Fields in an Avro schema file are defined via an array of objects, each of which is specified by a name and a type.

An Avro data file contains the schema and a compact binary representation of the data. Avro data files typically have the `.avro` suffix.

You can specify an Avro schema on both read and write operations to HDFS. You can provide either a binary `*.avro` file or a JSON-format `*.avsc` file for the schema file:

External Table Type	Schema Specified?	Description
readable	yes	PXF uses the specified schema; this overrides the schema embedded in the Avro data file.
readable	no	PXF uses the schema embedded in the Avro data file.
writable	yes	PXF uses the specified schema.
writable	no	PXF creates the Avro schema based on the external table definition.

When you provide the Avro schema file to PXF, the file must reside in the same location on each Greenplum Database host **or** the file may reside on the Hadoop file system. PXF first searches for an absolute file path on the Greenplum hosts. If PXF does not find the schema file there, it searches for the file relative to the PXF classpath. If PXF cannot find the schema file locally, it searches for the file on HDFS.

The `$PXF_BASE/conf` directory is in the PXF classpath. PXF can locate an Avro schema file that you add to this directory on every Greenplum Database host.

See Writing Avro Data for additional schema considerations when writing Avro data to HDFS.

Creating the External Table

Use the `hdfs:avro` profile to read or write Avro-format data in HDFS. The following syntax creates a Greenplum Database readable external table that references such a file:

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-file>?PROFILE=hdfs:avro[&SERVER=<server_name>][&<custom
-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export');
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-hdfs-file>	The path to the directory or file in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:avro`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are discussed below.
FORMAT 'CUSTOM'	Use `FORMAT` '`CUSTOM`' with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

For complex types, the PXF `hdfs:avro` profile inserts default delimiters between collection items and values before display. You can use non-default delimiter characters by identifying values for specific `hdfs:avro` custom options in the `CREATE EXTERNAL TABLE` command.

The `hdfs:avro` profile supports the following <custom-option>s:

Option Keyword	Description
COLLECTION_DELIM	The delimiter character(s) placed between entries in a top-level array, map, or record field when PXF maps an Avro complex data type to a text column. The default is the comma `,` character. (Read)
MAPKEY_DELIM	The delimiter character(s) placed between the key and value of a map entry when PXF maps an Avro complex data type to a text column. The default is the colon `:` character. (Read)
RECORDKEY_DELIM	The delimiter character(s) placed between the field name and value of a record entry when PXF maps an Avro complex data type to a text column. The default is the colon `:` character. (Read)
SCHEMA	The absolute path to the Avro schema file on the Greenplum host or on HDFS, or the relative path to the schema file on the host. (Read and Write)
IGNORE_MISSING_PATH	A Boolean value that specifies the action to take when <path-to-hdfs-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment. (Read)

The PXF `hdfs:avro` profile supports encoding- and compression-related write options. You specify these write options in the `CREATE WRITABLE EXTERNAL TABLE LOCATION` clause. The `hdfs:avro` profile

supports the following custom write options:

Write Option	Value Description
COMPRESSION_CODEC	The compression codec alias. Supported compression codecs for writing Avro data include: `bzip2`, `xz`, `snappy`, `deflate`, and `uncompressed` . If this option is not provided, PXF compresses the data using `deflate` compression.
CODEC_LEVEL	The compression level (applicable to the `deflate` and `xz` codecs only). This level controls the trade-off between speed and compression. Valid values are 1 (fastest) to 9 (most compressed). The default compression level is 6.

Example: Reading Avro Data

The examples in this section will operate on Avro data with the following field name and data type record schema:

- id - long
- username - string
- followers - array of string (string[])
- fmap - map of long
- relationship - enumerated type
- address - record comprised of street number (int), street name (string), and city (string)

You create an Avro schema and data file, and then create a readable external table to read the data.

Create Schema

Perform the following operations to create an Avro schema to represent the example schema described above.

1. Create a file named `avro_schema.avsc`:

```
$ vi /tmp/avro_schema.avsc
```

2. Copy and paste the following text into `avro_schema.avsc`:

```
{
"type" : "record",
  "name" : "example_schema",
  "namespace" : "com.example",
  "fields" : [ {
    "name" : "id",
    "type" : "long",
    "doc" : "Id of the user account"
  }, {
    "name" : "username",
    "type" : "string",
    "doc" : "Name of the user account"
  }, {
    "name" : "followers",
    "type" : {"type": "array", "items": "string"},
    "doc" : "Users followers"
  }, {
```

```
    "name": "fmap",
    "type": {"type": "map", "values": "long"}
  }, {
    "name": "relationship",
    "type": {
        "type": "enum",
        "name": "relationshipEnum",
        "symbols": ["MARRIED","LOVE","FRIEND","COLLEAGUE","STRANGER","ENEMY"]
    }
  }, {
    "name": "address",
    "type": {
        "type": "record",
        "name": "addressRecord",
        "fields": [
            {"name":"number", "type":"int"},
            {"name":"street", "type":"string"},
            {"name":"city", "type":"string"}]
    }
  } ],
  "doc:" : "A basic schema for storing messages"
}
```

Create Avro Data File (JSON)

Perform the following steps to create a sample Avro data file conforming to the above schema.

1. Create a text file named `pxf_avro.txt`:

   ```
   $ vi /tmp/pxf_avro.txt
   ```

2. Enter the following data into `pxf_avro.txt`:

   ```
   {"id":1, "username":"john","followers":["kate", "santosh"], "relationship": "FRIEND", "fmap": {"kate":10,"santosh":4}, "address":{"number":1, "street":"renaissance drive", "city":"san jose"}}

   {"id":2, "username":"jim","followers":["john", "pam"], "relationship": "COLLEAGUE", "fmap": {"john":3,"pam":3}, "address":{"number":9, "street":"deer creek", "city":"palo alto"}}
   ```

   The sample data uses a comma `,` to separate top level records and a colon `:` to separate map/key values and record field name/values.

3. Convert the text file to Avro format. There are various ways to perform the conversion, both programmatically and via the command line. In this example, we use the Java Avro tools; the jar `avro-tools-1.9.1.jar` file resides in the current directory:

   ```
   $ java -jar ./avro-tools-1.9.1.jar fromjson --schema-file /tmp/avro_schema.avsc /tmp/pxf_avro.txt > /tmp/pxf_avro.avro
   ```

   The generated Avro binary data file is written to `/tmp/pxf_avro.avro`.

4. Copy the generated Avro file to HDFS:

   ```
   $ hdfs dfs -put /tmp/pxf_avro.avro /data/pxf_examples/
   ```

Reading Avro Data

Perform the following operations to create and query an external table that references the `pxf_avro.avro` file that you added to HDFS in the previous section. When creating the table:

- Use the PXF default server.
- Map the top-level primitive fields, `id` (type long) and `username` (type string), to their equivalent Greenplum Database types (bigint and text).
- Map the `followers` field to a text array (text[]).
- Map the remaining complex fields to type text.
- Explicitly set the record, map, and collection delimiters using the `hdfs:avro` profile custom options.

1. Use the `hdfs:avro` profile to create a queryable external table from the `pxf_avro.avro` file:

```
postgres=# CREATE EXTERNAL TABLE pxf_hdfs_avro(id bigint, username text, followers text[], fmap text, relationship text, address text)
            LOCATION ('pxf://data/pxf_examples/pxf_avro.avro?PROFILE=hdfs:avro&COLLECTION_DELIM=,&MAPKEY_DELIM=:&RECORDKEY_DELIM=:')
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

2. Perform a simple query of the `pxf_hdfs_avro` table:

```
postgres=# SELECT * FROM pxf_hdfs_avro;
```

```
 id | username |   followers    |        fmap         | relationship |                      address
----+----------+----------------+--------------------+--------------+---------------------------------------------------
  1 | john     | {kate,santosh} | {kate:10,santosh:4} | FRIEND       | {number:1,street:renaissance drive,city:san jose}
  2 | jim      | {john,pam}     | {pam:3,john:3}      | COLLEAGUE    | {number:9,street:deer creek,city:palo alto}
(2 rows)
```

The simple query of the external table shows the components of the complex type data separated with the delimiters specified in the `CREATE EXTERNAL TABLE` call.

3. Query the table, displaying the `id` and the first element of the `followers` text array:

```
postgres=# SELECT id, followers[1] FROM pxf_hdfs_avro;
 id | followers
----+-----------
  1 | kate
  2 | john
```

Writing Avro Data

The PXF HDFS connector `hdfs:avro` profile supports writing Avro data to HDFS. When you create a writable external table to write Avro data, you specify the name of a directory on HDFS. When you insert records into the writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specify.

When you create a writable external table to write data to an Avro file, each table row is an Avro record and each table column is an Avro field.

If you do not specify a `SCHEMA` file, PXF generates a schema for the Avro file based on the Greenplum Database external table definition. PXF assigns the name of the external table column to the Avro field name. Because Avro has a `null` type and Greenplum external tables do not support the `NOT NULL` column qualifier, PXF wraps each data type in an Avro `union` of the mapped type and `null`. For example, for a writable external table column that you define with the Greenplum Database `text` data type, PXF generates the following schema element:

```
["string", "null"]
```

PXF returns an error if you provide a schema that does not include a `union` of the field data type with `null`, and PXF encounters a NULL data field.

PXF supports writing only Avro primitive data types and Avro Arrays of the types identified in Data Type Write Mapping. PXF does not support writing complex types to Avro:

- When you specify a `SCHEMA` file in the `LOCATION`, the schema must include only primitive data types.
- When PXF generates the schema, it writes any complex type that you specify in the writable external table column definition to the Avro file as a single Avro `string` type. For example, if you write an array of the Greenplum `numeric` type, PXF converts the array to a `string`, and you must read this data with a Greenplum `text`-type column.

Example: Writing Avro Data

In this example, you create an external table that writes to an Avro file on HDFS, letting PXF generate the Avro schema. After you insert some data into the file, you create a readable external table to query the Avro data.

The Avro file that you create and read in this example includes the following fields:

- id: `int`
- username: `text`
- followers: `text[]`

Example procedure:

1. Create the writable external table:

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_avrowrite(id int, username text,
followers text[])
            LOCATION ('pxf://data/pxf_examples/pxfwrite.avro?PROFILE=hdfs:avro'
)
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

1. Insert some data into the `pxf_avrowrite` table:

```
postgres=# INSERT INTO pxf_avrowrite VALUES (33, 'oliver', ARRAY['alex','frank'
]);
postgres=# INSERT INTO pxf_avrowrite VALUES (77, 'lisa', ARRAY['tom','mary']);
```

PXF uses the external table definition to generate the Avro schema.

2. Create an external table to read the Avro data that you just inserted into the table:

```
postgres=# CREATE EXTERNAL TABLE read_pxfwrite(id int, username text, followers
 text[])
            LOCATION ('pxf://data/pxf_examples/pxfwrite.avro?PROFILE=hdfs:avro'
)
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

3. Read the Avro data by querying the `read_pxfwrite` table:

```
postgres=# SELECT id, followers, followers[1], followers[2] FROM read_pxfwrite
ORDER BY id;
```

```
 id |   followers    | followers | followers
----+----------------+-----------+-----------
 33 | {alex,frank} | alex      | frank
 77 | {tom,mary}   | tom       | mary
(2 rows)
```

Reading JSON Data from HDFS

Use the PXF HDFS Connector to read JSON-format data. This section describes how to use PXF to access JSON data in HDFS, including how to create and query an external table that references a JSON file in the HDFS data store.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from HDFS.

Working with JSON Data

JSON is a text-based data-interchange format. JSON data is typically stored in a file with a `.json` suffix.

A `.json` file will contain a collection of objects. A JSON object is a collection of unordered name/value pairs. A value can be a string, a number, true, false, null, or an object or an array. You can define nested JSON objects and arrays.

Sample JSON data file content:

```
  {
    "created_at":"MonSep3004:04:53+00002013",
    "id_str":"384529256681725952",
    "user": {
      "id":31424214,
      "location":"COLUMBUS"
    },
    "coordinates":{
      "type":"Point",
      "values":[
         13,
```

```
          99
        ]
      }
    }
```

In the sample above, `user` is an object composed of fields named `id` and `location`. To specify the nested fields in the `user` object as Greenplum Database external table columns, use `.` projection:

```
user.id
user.location
```

`coordinates` is an object composed of a text field named `type` and an array of integers named `values`.

To retrieve all of the elements of the `values` array in a single column, define the corresponding Greenplum Database external table column as type `TEXT[]`.

```
"coordinates.values" TEXT[]
```

Refer to Introducing JSON for detailed information on JSON syntax.

JSON to Greenplum Database Data Type Mapping

To represent JSON data in Greenplum Database, map data values that use a primitive data type to Greenplum Database columns of the same type. JSON supports complex data types including projections and arrays. Use N-level projection to map members of nested objects and arrays to primitive data types.

The following table summarizes external mapping rules for JSON data.

Table 1. JSON Mapping

JSON Data Type	PXF/Greenplum Data Type
Primitive type (integer, float, string, boolean, null)	Use the corresponding Greenplum Database built-in data type; see Greenplum Database Data Types.
Array	Use `TEXT[]` to retrieve the JSON array as a Greenplum text array.
Object	Use dot `.` notation to specify each level of projection (nesting) to a member of a primitive or Array type.

JSON Data Read Modes

PXF supports two data read modes. The default mode expects one full JSON record per line. PXF also supports a read mode operating on JSON records that span multiple lines.

In upcoming examples, you will use both read modes to operate on a sample data set. The schema of the sample data set defines objects with the following member names and value data types:

- “created_at” - text
- “id_str” - text
- “user” - object
    - “id” - integer
    - “location” - text

- “coordinates” - object (optional)
  - “type” - text
  - “values” - array
    - [0] - integer
    - [1] - integer

The single-JSON-record-per-line data set follows:

```
{"created_at":"FriJun0722:45:03+00002013","id_str":"343136551322136576","user":{"id":3
95504494,"location":"NearCornwall"},"coordinates":{"type":"Point","values": [ 6, 50 ]}
},
{"created_at":"FriJun0722:45:02+00002013","id_str":"343136547115253761","user":{"id":2
6643566,"location":"Austin,Texas"}, "coordinates": null},
{"created_at":"FriJun0722:45:02+00002013","id_str":"343136547136233472","user":{"id":2
87819058,"location":""}, "coordinates": null}
```

This is the data set for the multi-line JSON record data set:

```
{
  "root":[
    {
      "record_obj":{
        "created_at":"MonSep3004:04:53+00002013",
        "id_str":"384529256681725952",
        "user":{
          "id":31424214,
          "location":"COLUMBUS"
        },
        "coordinates":null
      },
      "record_obj":{
        "created_at":"MonSep3004:04:54+00002013",
        "id_str":"384529260872228864",
        "user":{
          "id":67600981,
          "location":"KryberWorld"
        },
        "coordinates":{
          "type":"Point",
          "values":[
             8,
             52
          ]
        }
      }
    }
  ]
}
```

You will create JSON files for the sample data sets and add them to HDFS in the next section.

Loading the Sample JSON Data to HDFS

The PXF HDFS connector reads native JSON stored in HDFS. Before you can use Greenplum Database to query JSON format data, the data must reside in your HDFS data store.

Copy and paste the single line JSON record sample data set above to a file named `singleline.json`. Similarly, copy and paste the multi-line JSON record data set to a file named `multiline.json`.

Note: Ensure that there are **no** blank lines in your JSON files.

Copy the JSON data files that you just created to your HDFS data store. Create the `/data/pxf_examples` directory if you did not do so in a previous exercise. For example:

```
$ hdfs dfs -mkdir /data/pxf_examples
$ hdfs dfs -put singleline.json /data/pxf_examples
$ hdfs dfs -put multiline.json /data/pxf_examples
```

Once the data is loaded to HDFS, you can use Greenplum Database and PXF to query and analyze the JSON data.

Creating the External Table

Use the `hdfs:json` profile to read JSON-format files from HDFS. The following syntax creates a Greenplum Database readable external table that references such a file:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-file>?PROFILE=hdfs:json[&SERVER=<server_name>][&<custom
-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-hdfs-file>	The path to the directory or file in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:json`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are discussed below.
FORMAT 'CUSTOM'	Use `FORMAT 'CUSTOM'` with the `hdfs:json` profile. The `CUSTOM FORMAT` requires that you specify `(FORMATTER='pxfwritable_import')`.

PXF supports single- and multi- line JSON records. When you want to read multi-line JSON records, you must provide an `IDENTIFIER` <custom-option> and value. Use this <custom-option> to identify the name of a field whose parent JSON object you want to be returned as individual tuples.

The `hdfs:json` profile supports the following <custom-option>s:

Option Keyword	Syntax, Example(s)	Description

IDENTIFIER	`&IDENTIFIER=<value>` `&IDENTIFIER=created_at`	You must include the `IDENTIFIER` keyword and <value> in the `LOCATION` string only when you are accessing JSON data comprised of multi-line records. Use the <value> to identify the name of the field whose parent JSON object you want to be returned as individual tuples.
IGNORE_MISSING_PATH	`&IGNORE_MISSING_PATH=<boolean>`	Specify the action to take when <path-to-hdfs-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.

When a nested object in a multi-line record JSON file includes a field with the same name as that of a parent object field *and* the field name is also specified as the `IDENTIFIER`, there is a possibility that PXF could return incorrect results. Should you need to, you can work around this edge case by compressing the JSON file, and having PXF read the compressed file.

Example: Reading a JSON File with Single Line Records

Use the following CREATE EXTERNAL TABLE SQL command to create a readable external table that references the single-line-per-record JSON data file and uses the PXF default server.

```
CREATE EXTERNAL TABLE singleline_json_tbl(
  created_at TEXT,
  id_str TEXT,
  "user.id" INTEGER,
  "user.location" TEXT,
  "coordinates.values" TEXT[]
)
LOCATION('pxf://data/pxf_examples/singleline.json?PROFILE=hdfs:json')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Notice the use of `.` projection to access the nested fields in the `user` and `coordinates` objects.

To query the JSON data in the external table:

```
SELECT * FROM singleline_json_tbl;
```

To access specific elements of the `coordinates.values` array, you can specify the array subscript number in square brackets:

```
SELECT "coordinates.values"[1], "coordinates.values"[2] FROM singleline_json_tbl;
```

To access the array elements as some type other than `TEXT`, you can either cast the whole column:

```
SELECT "coordinates.values"::int[] FROM singleline_json_tbl;
```

or cast specific elements:

```
SELECT "coordinates.values"[1]::int, "coordinates.values"[2]::float FROM singleline_js
on_tbl;
```

Example: Reading a JSON file with Multi-Line Records

The SQL command to create a readable external table from the multi-line-per-record JSON file is

very similar to that of the single line data set above. You must additionally specify the `LOCATION` clause `IDENTIFIER` keyword and an associated value when you want to read multi-line JSON records. For example:

```
CREATE EXTERNAL TABLE multiline_json_tbl(
  created_at TEXT,
  id_str TEXT,
  "user.id" INTEGER,
  "user.location" TEXT,
  "coordinates.values" TEXT[]
)
LOCATION('pxf://data/pxf_examples/multiline.json?PROFILE=hdfs:json&IDENTIFIER=created_
at')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

`created_at` identifies the member name of the first field in the JSON record `record_obj` in the sample data schema.

To query the JSON data in this external table:

```
SELECT * FROM multiline_json_tbl;
```

Other Methods to Read a JSON Array

Starting in version 6.2.0, PXF supports reading a JSON array into a `TEXT[]` column. PXF still supports the old methods of using array element projection or a single text-type column to read a JSON array. These access methods are described here.

Using Array Element Projection

PXF supports accessing specific elements of a JSON array using the syntax `[n]` in the table definition to identify the specific element.

```
CREATE EXTERNAL TABLE singleline_json_tbl_aep(
  created_at TEXT,
  id_str TEXT,
  "user.id" INTEGER,
  "user.location" TEXT,
  "coordinates.values[0]" INTEGER,
  "coordinates.values[1]" INTEGER
)
LOCATION('pxf://data/pxf_examples/singleline.json?PROFILE=hdfs:json')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Note: When you use this method to identify specific array elements, PXF provides *only* those values to Greenplum Database, not the whole JSON array.

If your existing external table definition uses array element projection and you want to read the array into a `TEXT[]` column, you can use the `ALTER EXTERNAL TABLE` command to update the table definition. For example:

```
ALTER EXTERNAL TABLE singleline_json_tbl_aep DROP COLUMN "coordinates.values[0]", DROP
 COLUMN "coordinates.values[1]", ADD COLUMN "coordinates.values" TEXT[];
```

If you choose to alter the external table definition in this manner, be sure to update any existing queries on the external table to account for the changes to column name and type.

Specifying a Single Text-type Column

PXF supports accessing all of the elements within an array as a single string containing the serialized JSON array by defining the corresponding Greenplum table column with one of the following data types: `TEXT`, `VARCHAR`, or `BPCHAR`.

```
CREATE EXTERNAL TABLE singleline_json_tbl_stc(
  created_at TEXT,
  id_str TEXT,
  "user.id" INTEGER,
  "user.location" TEXT,
  "coordinates.values" TEXT
)
LOCATION('pxf://data/pxf_examples/singleline.json?PROFILE=hdfs:json')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

If you retrieve the JSON array in a single text-type column and wish to convert the JSON array serialized as `TEXT` back into a native Greenplum array type, you can use the example query below:

```
SELECT user.id,
       ARRAY(SELECT json_array_elements_text(coordinates.values::json))::int[] AS coor
ds
FROM singleline_json_tbl_stc;
```

Note: This conversion is possible only when you are using PXF with Greenplum Database 6.x; the function `json_array_elements_text()` is not available in Greenplum 5.x.

If your external table definition uses a single text-type column for a JSON array and you want to read the array into a `TEXT[]` column, you can use the `ALTER EXTERNAL TABLE` command to update the table definition. For example:

```
ALTER EXTERNAL TABLE singleline_json_tbl_stc ALTER COLUMN "coordinates.values" TYPE TE
XT[];
```

If you choose to alter the external table definition in this manner, be sure to update any existing queries on the external table to account for the change in column type.

Reading ORC Data

Use the PXF HDFS connector `hdfs:orc` profile to read ORC-format data when the data resides in a Hadoop file system. This section describes how to read HDFS files that are stored in ORC format, including how to create and query an external table that references these files in the HDFS data store.

The `hdfs:orc` profile:

- Reads 1024 rows of data at a time.
- Supports column projection.
- Supports filter pushdown based on file-level, stripe-level, and row-level ORC statistics.
- Supports the compound list type for a subset of ORC scalar types.

- Does not support the map, union or struct compound types.

The `hdfs:orc` profile currently supports reading scalar data types and lists of certain scalar types from ORC files. If the data resides in a Hive table, and you want to read complex types or the Hive table is partitioned, use the `hive:orc` profile.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from HDFS.

About the ORC Data Format

The Optimized Row Columnar (ORC) file format is a columnar file format that provides a highly efficient way to both store and access HDFS data. ORC format offers improvements over text and RCFile formats in terms of both compression and performance. PXF supports ORC file versions v0 and v1.

ORC is type-aware and specifically designed for Hadoop workloads. ORC files store both the type of, and encoding information for, the data in the file. All columns within a single group of row data (also known as stripe) are stored together on disk in ORC format files. The columnar nature of the ORC format type enables read projection, helping avoid accessing unnecessary columns during a query.

ORC also supports predicate pushdown with built-in indexes at the file, stripe, and row levels, moving the filter operation to the data loading phase.

Refer to the Apache orc documentation for detailed information about the ORC file format.

Data Type Mapping

To read ORC scalar data types in Greenplum Database, map ORC data values to Greenplum Database columns of the same type. PXF uses the following data type mapping when it reads ORC data:

ORC Physical Type	ORC Logical Type	PXF/Greenplum Data Type
binary	decimal	Numeric
binary	timestamp	Timestamp
byte[]	string	Text
byte[]	char	Bpchar
byte[]	varchar	Varchar
byte[]	binary	Bytea
Double	float	Real
Double	double	Float8
Integer	boolean (1 bit)	Boolean
Integer	tinyint (8 bit)	Smallint

ORC Physical Type	ORC Logical Type	PXF/Greenplum Data Type
Integer	smallint (16 bit)	Smallint
Integer	int (32 bit)	Integer
Integer	bigint (64 bit)	Bigint
Integer	date	Date

PXF supports only the list ORC compound type, and only for a subset of the ORC scalar types. The supported mappings follow:

ORC Compound Type	PXF/Greenplum Data Type
array<string>	Text[]
array<char>	Bpchar[]
array<varchar>	Varchar[]
array<binary>	Bytea[]
array<float>	Real[]
array<double>	Float8[]
array<boolean>	Boolean[]
array<tinyint>	Smallint[]
array<smallint>	Smallint[]
array<int>	Integer[]
array<bigint>	Bigint[]

Creating the External Table

The PXF HDFS connector `hdfs:orc` profile supports reading ORC-format HDFS files. Use the following syntax to create a Greenplum Database external table that references a file or directory:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-file>
    ?PROFILE=hdfs:orc[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import')
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described below.

Keyword	Value
<path-to-hdfs-file>	The path to the file or directory in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:orc`.

Keyword	Value
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are described below.
FORMAT	Use `FORMAT 'CUSTOM'`; the `CUSTOM` format requires the built-in `pxfwritable_import formatter`.

The PXF `hdfs:orc` profile supports the following read options. You specify this option in the `LOCATION` clause:

Read Option	Value Description
IGNORE_MISSING_PATH	A Boolean value that specifies the action to take when <path-to-hdfs-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
MAP_BY_POSITION	A Boolean value that, when set to `true`, specifies that PXF should map an ORC column to a Greenplum Database column by position. The default value is `false`, PXF maps an ORC column to a Greenplum column by name.

Example: Reading an ORC File on HDFS

This example operates on a simple data set that models a retail sales operation. The data includes fields with the following names and types:

Column Name	Data Type
location	text
month	text
num_orders	integer
total_sales	numeric(10,2)
items_sold	text[]

In this example, you:

- Create a sample data set in JSON format, use the `orc-tools` JAR utilities to convert the JSON file into an ORC-format file, and then copy the ORC file to HDFS.
- Create a Greenplum Database readable external table that references the ORC file and that specifies the `hdfs:orc` profile.
- Query the external table.

You must have administrative privileges to both a Hadoop cluster and a Greenplum Database cluster to run the example. You must also have configured a PXF server to access Hadoop.

Procedure:

1. Create a JSON file named `sampledata.json` in the `/tmp` directory:

```
hdfsclient$ echo '{"location": "Prague", "month": "Jan","num_orders": 101, "tot
al_sales": 4875.33, "items_sold": ["boots", "hats"]}
{"location": "Rome", "month": "Mar","num_orders": 87, "total_sales": 1557.39, "
items_sold": ["coats"]}
```

```
{"location": "Bangalore", "month": "May","num_orders": 317, "total_sales": 8936
.99, "items_sold": ["winter socks", "long-sleeved shirts", "boots"]}
{"location": "Beijing", "month": "Jul","num_orders": 411, "total_sales": 11600.
67, "items_sold": ["hoodies/sweaters", "pants"]}
{"location": "Los Angeles", "month": "Dec","num_orders": 0, "total_sales": 0.00
, "items_sold": null}
' > /tmp/sampledata.json
```

2. Download the `orc-tools` JAR.
3. Run the `orc-tools convert` command to convert `sampledata.json` to the ORC file `/tmp/sampledata.orc`; provide the schema to the command:

```
hdfsclient$ java -jar orc-tools-1.6.2-uber.jar convert /tmp/sampledata.json \
  --schema 'struct<location:string,month:string,num_orders:int,total_sales:deci
mal(10,2),items_sold:array<string>>' \
  -o /tmp/sampledata.orc
```

4. Copy the ORC file to HDFS. The following command copies the file to the `/data/pxf_examples` directory:

```
hdfsclient$ hdfs dfs -put /tmp/sampledata.orc /data/pxf_examples/
```

5. Log in to the Greenplum Database master host and connect to a database. This command connects to the database named `testdb` as the `gpadmin` user:

```
gpadmin@gpmaster$ psql -d testdb
```

6. Create an external table named `sample_orc` that references the `/data/pxf_examples/sampledata.orc` file on HDFS. This command creates the table with the column names specified in the ORC schema, and uses the `default` PXF server:

```
testdb=# CREATE EXTERNAL TABLE sample_orc(location text, month text, num_orders
 int, total_sales numeric(10,2), items_sold text[])
           LOCATION ('pxf://data/pxf_examples/sampledata.orc?PROFILE=hdfs:orc')
         FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

7. Read the data in the file by querying the `sample_orc` table:

```
testdb=# SELECT * FROM sample_orc;
```

```
  location    | month | num_orders | total_sales |                 items_sold
-------------+-------+------------+-------------+------------------------------
----------------
 Prague      | Jan   |        101 |     4875.33 | {boots,hats}
 Rome        | Mar   |         87 |     1557.39 | {coats}
 Bangalore   | May   |        317 |     8936.99 | {"winter socks","long-sleeved
 shirts",boots}
 Beijing     | Jul   |        411 |    11600.67 | {hoodies/sweaters,pants}
 Los Angeles | Dec   |          0 |        0.00 |
(5 rows)
```

8. You can query the data on any column, including the `items_sold` array column. For example, this query returns the rows where the items sold include `boots` and/or `pants`:

```
testdb=# SELECT * FROM sample_orc WHERE items_sold && '{"boots", "pants"}';
```

```
 location  | month | num_orders | total_sales |                items_sold
-----------+-------+------------+-------------+--------------------------------
--------------
 Prague    | Jan   |        101 |     4875.33 | {boots,hats}
 Bangalore | May   |        317 |     8936.99 | {"winter socks","long-sleeved s
hirts",boots}
 Beijing   | Jul   |        411 |    11600.67 | {hoodies/sweaters,pants}
(3 rows)
```

9. This query returns the rows where the first item sold is `boots`:

```
testdb=# SELECT * FROM sample_orc WHERE items_sold[0] = 'boots';
```

```
 location  | month | num_orders | total_sales |                items_sold
-----------+-------+------------+-------------+--------------------------------
--------------
 Prague    | Jan   |        101 |     4875.33 | {boots,hats}
(1 row)
```

Reading and Writing HDFS Parquet Data

Use the PXF HDFS connector to read and write Parquet-format data. This section describes how to read and write HDFS files that are stored in Parquet format, including how to create, query, and insert into external tables that reference files in the HDFS data store.

PXF supports reading or writing Parquet files compressed with these codecs: `snappy`, `gzip`, and `lzo`.

PXF currently supports reading and writing primitive Parquet data types only.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from or write data to HDFS.

Data Type Mapping

To read and write Parquet primitive data types in Greenplum Database, map Parquet data values to Greenplum Database columns of the same type.

Parquet supports a small set of primitive data types, and uses metadata annotations to extend the data types that it supports. These annotations specify how to interpret the primitive type. For example, Parquet stores both `INTEGER` and `DATE` types as the `INT32` primitive type. An annotation identifies the original type as a `DATE`.

Read Mapping

PXF uses the following data type mapping when reading Parquet data:

Parquet Data Type	Original Type	PXF/Greenplum Data Type

binary (byte_array)	Date	Date
binary (byte_array)	Timestamp_millis	Timestamp
binary (byte_array)	all others	Text
binary (byte_array)	–	Bytea
boolean	–	Boolean
double	–	Float8
fixed_len_byte_array	–	Numeric
float	–	Real
int32	Date	Date
int32	Decimal	Numeric
int32	int_8	Smallint
int32	int_16	Smallint
int32	–	Integer
int64	Decimal	Numeric
int64	–	Bigint
int96	–	Timestamp

Note: PXF supports filter predicate pushdown on all parquet data types listed above, *except* the `fixed_len_byte_array` and `int96` types.

Write Mapping

PXF uses the following data type mapping when writing Parquet data:

PXF/Greenplum Data Type	Original Type	Parquet Data Type
Boolean	–	boolean
Bytea	–	binary
Bigint	–	int64
SmallInt	int_16	int32
Integer	–	int32
Real	–	float
Float8	–	double
Numeric/Decimal	Decimal	fixed_len_byte_array
Timestamp[1]	–	int96
Timestamptz[2]	–	int96
Date	date	int32

PXF/Greenplum Data Type	Original Type	Parquet Data Type
Time	utf8	binary
Varchar	utf8	binary
Text	utf8	binary
OTHERS	–	UNSUPPORTED

[1] PXF localizes a `Timestamp` to the current system timezone and converts it to universal time (UTC) before finally converting to `int96`.

[2] PXF converts a `Timestamptz` to a UTC `timestamp` and then converts to `int96`. PXF loses the time zone information during this conversion.

Creating the External Table

The PXF HDFS connector `hdfs:parquet` profile supports reading and writing HDFS data in Parquet-format. When you insert records into a writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specified.

Use the following syntax to create a Greenplum Database external table that references an HDFS directory:

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-dir>
    ?PROFILE=hdfs:parquet[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export');
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-hdfs-file>	The path to the directory in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:parquet`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are described below.
FORMAT 'CUSTOM'	Use `FORMAT` '`CUSTOM`' with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

The PXF `hdfs:parquet` profile supports the following read option. You specify this option in the

CREATE EXTERNAL TABLE LOCATION clause:

Read Option	Value Description
IGNORE_MISSING_PATH	A Boolean value that specifies the action to take when <path-to-hdfs-file> is missing or invalid. The default value is false, PXF returns an error in this situation. When the value is true, PXF ignores missing path errors and returns an empty fragment.

The PXF hdfs:parquet profile supports encoding- and compression-related write options. You specify these write options in the CREATE WRITABLE EXTERNAL TABLE LOCATION clause. The hdfs:parquet profile supports the following custom write options:

Write Option	Value Description
COMPRESSION_CODEC	The compression codec alias. Supported compression codecs for writing Parquet data include: snappy, gzip, lzo, and uncompressed . If this option is not provided, PXF compresses the data using snappy compression.
ROWGROUP_SIZE	A Parquet file consists of one or more row groups, a logical partitioning of the data into rows. ROWGROUP_SIZE identifies the size (in bytes) of the row group. The default row group size is 8 * 1024 * 1024 bytes.
PAGE_SIZE	A row group consists of column chunks that are divided up into pages. PAGE_SIZE is the size (in bytes) of such a page. The default page size is 1 * 1024 * 1024 bytes.
ENABLE_DICTIONARY	A boolean value that specifies whether or not to enable dictionary encoding. The default value is true; dictionary encoding is enabled when PXF writes Parquet files.
DICTIONARY_PAGE_SIZE	When dictionary encoding is enabled, there is a single dictionary page per column, per row group. DICTIONARY_PAGE_SIZE is similar to PAGE_SIZE, but for the dictionary. The default dictionary page size is 1 * 1024 * 1024 bytes.
PARQUET_VERSION	The Parquet version; PXF supports the values v1 and v2 for this option. The default Parquet version is v1.
SCHEMA	The location of the Parquet schema file on the file system of the specified SERVER.

Note: You must explicitly specify uncompressed if you do not want PXF to compress the data.

Parquet files that you write to HDFS with PXF have the following naming format: <file>.<compress_extension>.parquet, for example 1547061635-0000004417_0.gz.parquet.

Example

This example utilizes the data schema introduced in Example: Reading Text Data on HDFS.

Column Name	Data Type
location	text
month	text
number_of_orders	int
total_sales	float8

In this example, you create a Parquet-format writable external table that uses the default PXF server to reference Parquet-format data in HDFS, insert some data into the table, and then create a readable external table to read the data.

1. Use the `hdfs:parquet` profile to create a writable external table. For example:

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_tbl_parquet (location text, month
 text, number_of_orders int, total_sales double precision)
    LOCATION ('pxf://data/pxf_examples/pxf_parquet?PROFILE=hdfs:parquet')
  FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

2. Write a few records to the `pxf_parquet` HDFS directory by inserting directly into the `pxf_tbl_parquet` table. For example:

```
postgres=# INSERT INTO pxf_tbl_parquet VALUES ( 'Frankfurt', 'Mar', 777, 3956.9
8 );
postgres=# INSERT INTO pxf_tbl_parquet VALUES ( 'Cleveland', 'Oct', 3812, 96645
.37 );
```

3. Recall that Greenplum Database does not support directly querying a writable external table. To read the data in `pxf_parquet`, create a readable external Greenplum Database referencing this HDFS directory:

```
postgres=# CREATE EXTERNAL TABLE read_pxf_parquet(location text, month text, nu
mber_of_orders int, total_sales double precision)
    LOCATION ('pxf://data/pxf_examples/pxf_parquet?PROFILE=hdfs:parquet')
    FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

4. Query the readable external table `read_pxf_parquet`:

```
postgres=# SELECT * FROM read_pxf_parquet ORDER BY total_sales;
```

```
 location  | month | number_of_orders | total_sales
-----------+-------+------------------+-------------
 Frankfurt | Mar   |              777 |     3956.98
 Cleveland | Oct   |             3812 |     96645.4
(2 rows)
```

Reading and Writing HDFS SequenceFile Data

The PXF HDFS connector supports SequenceFile format binary data. This section describes how to use PXF to read and write HDFS SequenceFile data, including how to create, insert, and query data in external tables that reference files in the HDFS data store.

PXF supports reading or writing SequenceFile files compressed with the `default`, `bzip2`, and `gzip` codecs.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read data from or write data to HDFS.

Creating the External Table

The PXF HDFS connector `hdfs:SequenceFile` profile supports reading and writing HDFS data in SequenceFile binary format. When you insert records into a writable external table, the block(s) of

data that you insert are written to one or more files in the directory that you specified.

Note: External tables that you create with a writable profile can only be used for INSERT operations. If you want to query the data that you inserted, you must create a separate readable external table that references the HDFS directory.

Use the following syntax to create a Greenplum Database external table that references an HDFS directory:

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-hdfs-dir>
    ?PROFILE=hdfs:SequenceFile[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (<formatting-properties>)
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-hdfs-dir>	The path to the directory in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-dir> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-hdfs-dir> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:SequenceFile`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
<custom-option>	<custom-option>s are described below.
FORMAT	Use `FORMAT` '`CUSTOM`' with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

SequenceFile format data can optionally employ record or block compression and a specific compression codec.

When you use the `hdfs:SequenceFile` profile to write SequenceFile format data, you must provide the name of the Java class to use for serializing/deserializing the binary data. This class must provide read and write methods for each data type referenced in the data schema.

You specify the compression type and codec, and the Java serialization/deserialization class, via custom options to the `CREATE EXTERNAL TABLE LOCATION` clause. The `hdfs:SequenceFile` profile supports the following custom options:

Option	Value Description
COMPRESSION_CODEC	The compression codec alias. Supported compression codecs include: `default`, `bzip2`, `gzip`, and `uncompressed`. If this option is not provided, Greenplum Database performs no data compression.
COMPRESSION_TYPE	The compression type to employ; supported values are `RECORD` (the default) or `BLOCK`.

Option	Value Description
DATA-SCHEMA	The name of the writer serialization/deserialization class. The jar file in which this class resides must be in the PXF classpath. This option is required for the `hdfs:SequenceFile` profile and has no default value.
IGNORE_MISSING_PATH	A Boolean value that specifies the action to take when <path-to-hdfs-dir> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.

Reading and Writing Binary Data

Use the HDFS connector `hdfs:SequenceFile` profile when you want to read or write SequenceFile format data to HDFS. Files of this type consist of binary key/value pairs. SequenceFile format is a common data transfer format between MapReduce jobs.

Example: Writing Binary Data to HDFS

In this example, you create a Java class named `PxfExample_CustomWritable` that will serialize/deserialize the fields in the sample schema used in previous examples. You will then use this class to access a writable external table that you create with the `hdfs:SequenceFile` profile and that uses the default PXF server.

Perform the following procedure to create the Java class and writable table.

1. Prepare to create the sample Java class:

```
$ mkdir -p pxfex/com/example/pxf/hdfs/writable/dataschema
$ cd pxfex/com/example/pxf/hdfs/writable/dataschema
$ vi PxfExample_CustomWritable.java
```

2. Copy and paste the following text into the `PxfExample_CustomWritable.java` file:

```
package com.example.pxf.hdfs.writable.dataschema;

import org.apache.hadoop.io.*;
import java.io.DataInput;
import java.io.DataOutput;
import java.io.IOException;
import java.lang.reflect.Field;

/**
* PxfExample_CustomWritable class - used to serialize and deserialize data with
* text, int, and float data types
*/
public class PxfExample_CustomWritable implements Writable {

public String st1, st2;
public int int1;
public float ft;

public PxfExample_CustomWritable() {
    st1 = new String("");
    st2 = new String("");
    int1 = 0;
    ft = 0.f;
}
```

```
public PxfExample_CustomWritable(int i1, int i2, int i3) {

    st1 = new String("short_string___" + i1);
    st2 = new String("short_string___" + i1);
    int1 = i2;
    ft = i1 * 10.f * 2.3f;

}

String GetSt1() {
    return st1;
}

String GetSt2() {
    return st2;
}

int GetInt1() {
    return int1;
}

float GetFt() {
    return ft;
}

@Override
public void write(DataOutput out) throws IOException {

    Text txt = new Text();
    txt.set(st1);
    txt.write(out);
    txt.set(st2);
    txt.write(out);

    IntWritable intw = new IntWritable();
    intw.set(int1);
    intw.write(out);

    FloatWritable fw = new FloatWritable();
    fw.set(ft);
    fw.write(out);
}

@Override
public void readFields(DataInput in) throws IOException {

    Text txt = new Text();
    txt.readFields(in);
    st1 = txt.toString();
    txt.readFields(in);
    st2 = txt.toString();

    IntWritable intw = new IntWritable();
    intw.readFields(in);
    int1 = intw.get();

    FloatWritable fw = new FloatWritable();
    fw.readFields(in);
    ft = fw.get();
```

```
}

public void printFieldTypes() {
    Class myClass = this.getClass();
    Field[] fields = myClass.getDeclaredFields();

    for (int i = 0; i < fields.length; i++) {
        System.out.println(fields[i].getType().getName());
    }
}
}
```

3. Compile and create a Java class JAR file for `PxfExample_CustomWritable`. Provide a classpath that includes the `hadoop-common.jar` file for your Hadoop distribution. For example, if you installed the Hortonworks Data Platform Hadoop client:

```
$ javac -classpath /usr/hdp/current/hadoop-client/hadoop-common.jar  PxfExample
_CustomWritable.java
$ cd ../../../../../../
$ jar cf pxfex-customwritable.jar com
$ cp pxfex-customwritable.jar /tmp/
```

   (Your Hadoop library classpath may differ.)

4. Copy the `pxfex-customwritable.jar` file to the Greenplum Database master node. For example:

```
$ scp pxfex-customwritable.jar gpadmin@gpmaster:/home/gpadmin
```

5. Log in to your Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

6. Copy the `pxfex-customwritable.jar` JAR file to the user runtime library directory, and note the location. For example, if `PXF_BASE=/usr/local/pxf-gp6`:

```
gpadmin@gpmaster$ cp /home/gpadmin/pxfex-customwritable.jar /usr/local/pxf-gp6/
lib/pxfex-customwritable.jar
```

7. Synchronize the PXF configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

8. Restart PXF on each Greenplum Database host as described in Restarting PXF.

9. Use the PXF `hdfs:SequenceFile` profile to create a Greenplum Database writable external table. Identify the serialization/deserialization Java class you created above in the `DATA-SCHEMA` <custom-option>. Use `BLOCK` mode compression with `bzip2` when you create the writable table.

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_tbl_seqfile (location text, month
 text, number_of_orders integer, total_sales real)
            LOCATION ('pxf://data/pxf_examples/pxf_seqfile?PROFILE=hdfs:Sequenc
eFile&DATA-SCHEMA=com.example.pxf.hdfs.writable.dataschema.PxfExample_CustomWri
table&COMPRESSION_TYPE=BLOCK&COMPRESSION_CODEC=bzip2')
```

```
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

Notice that the `'CUSTOM' FORMAT` <formatting-properties> specifies the built-in `pxfwritable_export` formatter.

10. Write a few records to the `pxf_seqfile` HDFS directory by inserting directly into the `pxf_tbl_seqfile` table. For example:

```
postgres=# INSERT INTO pxf_tbl_seqfile VALUES ( 'Frankfurt', 'Mar', 777, 3956.9
8 );
postgres=# INSERT INTO pxf_tbl_seqfile VALUES ( 'Cleveland', 'Oct', 3812, 96645
.37 );
```

11. Recall that Greenplum Database does not support directly querying a writable external table. To read the data in `pxf_seqfile`, create a readable external Greenplum Database referencing this HDFS directory:

```
postgres=# CREATE EXTERNAL TABLE read_pxf_tbl_seqfile (location text, month tex
t, number_of_orders integer, total_sales real)
            LOCATION ('pxf://data/pxf_examples/pxf_seqfile?PROFILE=hdfs:Sequenc
eFile&DATA-SCHEMA=com.example.pxf.hdfs.writable.dataschema.PxfExample_CustomWri
table')
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

You must specify the `DATA-SCHEMA` <custom-option> when you read HDFS data via the `hdfs:SequenceFile` profile. You need not provide compression-related options.

12. Query the readable external table `read_pxf_tbl_seqfile`:

```
gpadmin=# SELECT * FROM read_pxf_tbl_seqfile ORDER BY total_sales;
```

```
 location  | month | number_of_orders | total_sales
-----------+-------+------------------+-------------
 Frankfurt | Mar   |              777 |     3956.98
 Cleveland | Oct   |             3812 |     96645.4
(2 rows)
```

Reading the Record Key

When a Greenplum Database external table references SequenceFile or another data format that stores rows in a key-value format, you can access the key values in Greenplum queries by using the `recordkey` keyword as a field name.

The field type of `recordkey` must correspond to the key type, much as the other fields must match the HDFS data.

You can define `recordkey` to be any of the following Hadoop types:

- BooleanWritable
- ByteWritable
- DoubleWritable
- FloatWritable

- IntWritable
- LongWritable
- Text

If no record key is defined for a row, Greenplum Database returns the id of the segment that processed the row.

Example: Using Record Keys

Create an external readable table to access the record keys from the writable table `pxf_tbl_seqfile` that you created in Example: Writing Binary Data to HDFS. Define the `recordkey` in this example to be of type `int8`.

```
postgres=# CREATE EXTERNAL TABLE read_pxf_tbl_seqfile_recordkey(recordkey int8, locati
on text, month text, number_of_orders integer, total_sales real)
                LOCATION ('pxf://data/pxf_examples/pxf_seqfile?PROFILE=hdfs:SequenceFi
le&DATA-SCHEMA=com.example.pxf.hdfs.writable.dataschema.PxfExample_CustomWritable')
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
gpadmin=# SELECT * FROM read_pxf_tbl_seqfile_recordkey;
```

```
 recordkey |  location   | month | number_of_orders | total_sales
-----------+-------------+-------+------------------+-------------
         2 | Frankfurt   | Mar   |              777 |     3956.98
         1 | Cleveland   | Oct   |             3812 |     96645.4
(2 rows)
```

You did not define a record key when you inserted the rows into the writable table, so the `recordkey` identifies the segment on which the row data was processed.

Reading a Multi-Line Text File into a Single Table Row

You can use the PXF HDFS connector to read one or more multi-line text files in HDFS each as a single table row. This may be useful when you want to read multiple files into the same Greenplum Database external table, for example when individual JSON files each contain a separate record.

PXF supports reading only text and JSON files in this manner.

Note: Refer to the Reading JSON Data from HDFS topic if you want to use PXF to read JSON files that include more than one record.

Prerequisites

Ensure that you have met the PXF Hadoop Prerequisites before you attempt to read files from HDFS.

Reading Multi-Line Text and JSON Files

You can read single- and multi-line files into a single table row, including files with embedded linefeeds. If you are reading multiple JSON files, each file must be a complete record, and each file must contain the same record type.

PXF reads the complete file data into a single row and column. When you create the external table

to read multiple files, you must ensure that all of the files that you want to read are of the same (text or JSON) type. You must also specify a single `text` or `json` column, depending upon the file type.

The following syntax creates a Greenplum Database readable external table that references one or more text or JSON files on HDFS:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> text|json | LIKE <other_table> )
  LOCATION ('pxf://<path-to-files>?PROFILE=hdfs:text:multi[&SERVER=<server_name>][&IGN
ORE_MISSING_PATH=<boolean>]&FILE_AS_ROW=true')
FORMAT 'CSV');
```

The keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-files>	The path to the directory or files in the HDFS data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-hdfs-files> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-files> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword must specify `hdfs:text:multi`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
FILE_AS_ROW=true	The required option that instructs PXF to read each file into a single table row.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-files> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
FORMAT	The `FORMAT` must specify `'CSV'`.

Note: The `hdfs:text:multi` profile does not support additional custom or format options when you specify the `FILE_AS_ROW=true` option.

For example, if `/data/pxf_examples/jdir` identifies an HDFS directory that contains a number of JSON files, the following statement creates a Greenplum Database external table that references all of the files in that directory:

```
CREATE EXTERNAL TABLE pxf_readjfiles(j1 json)
  LOCATION ('pxf://data/pxf_examples/jdir?PROFILE=hdfs:text:multi&FILE_AS_ROW=true')
FORMAT 'CSV';
```

When you query the `pxf_readjfiles` table with a `SELECT` statement, PXF returns the contents of each JSON file in `jdir/` as a separate row in the external table.

When you read JSON files, you can use the JSON functions provided in Greenplum Database to access individual data fields in the JSON record. For example, if the `pxf_readjfiles` external table above reads a JSON file that contains this JSON record:

```
{
  "root":[
    {
      "record_obj":{
```

```
        "created_at":"MonSep3004:04:53+00002013",
        "id_str":"384529256681725952",
        "user":{
          "id":31424214,
          "location":"COLUMBUS"
        },
        "coordinates":null
      }
    }
  ]
}
```

You can use the `json_array_elements()` function to extract specific JSON fields from the table row. For example, the following command displays the `user->id` field:

```
SELECT json_array_elements(j1->'root')->'record_obj'->'user'->'id'
  AS userid FROM pxf_readjfiles;

  userid
----------
 31424214
(1 rows)
```

Refer to Working with JSON Data in the Greenplum Documentation for specific information on manipulating JSON data in Greenplum.

Example: Reading an HDFS Text File into a Single Table Row

Perform the following procedure to create 3 sample text files in an HDFS directory, and use the PXF `hdfs:text:multi` profile and the default PXF server to read all of these text files in a single external table query.

1. Create an HDFS directory for the text files. For example:

   ```
   $ hdfs dfs -mkdir -p /data/pxf_examples/tdir
   ```

2. Create a text data file named `file1.txt`:

   ```
   $ echo 'text file with only one line' > /tmp/file1.txt
   ```

3. Create a second text data file named `file2.txt`:

   ```
   $ echo 'Prague,Jan,101,4875.33
   Rome,Mar,87,1557.39
   Bangalore,May,317,8936.99
   Beijing,Jul,411,11600.67' > /tmp/file2.txt
   ```

   This file has multiple lines.

4. Create a third text file named `/tmp/file3.txt`:

   ```
   $ echo '"4627 Star Rd.
   San Francisco, CA  94107":Sept:2017
   "113 Moon St.
   San Diego, CA  92093":Jan:2018
   "51 Belt Ct.
   Denver, CO  90123":Dec:2016
   ```

```
"93114 Radial Rd.
Chicago, IL  60605":Jul:2017
"7301 Brookview Ave.
Columbus, OH  43213":Dec:2018' > /tmp/file3.txt
```

This file includes embedded line feeds.

5. Save the file and exit the editor.

6. Copy the text files to HDFS:

```
$ hdfs dfs -put /tmp/file1.txt /data/pxf_examples/tdir
$ hdfs dfs -put /tmp/file2.txt /data/pxf_examples/tdir
$ hdfs dfs -put /tmp/file3.txt /data/pxf_examples/tdir
```

7. Log in to a Greenplum Database system and start the `psql` subsystem.

8. Use the `hdfs:text:multi` profile to create an external table that references the `tdir` HDFS directory. For example:

```
CREATE EXTERNAL TABLE pxf_readfileasrow(c1 text)
  LOCATION ('pxf://data/pxf_examples/tdir?PROFILE=hdfs:text:multi&FILE_AS_ROW=t
rue')
FORMAT 'CSV';
```

9. Turn on expanded display and query the `pxf_readfileasrow` table:

```
postgres=# \x on
postgres=# SELECT * FROM pxf_readfileasrow;
```

```
-[ RECORD 1 ]---------------------------
c1 | Prague,Jan,101,4875.33
   | Rome,Mar,87,1557.39
   | Bangalore,May,317,8936.99
   | Beijing,Jul,411,11600.67
-[ RECORD 2 ]---------------------------
c1 | text file with only one line
-[ RECORD 3 ]---------------------------
c1 | "4627 Star Rd.
   | San Francisco, CA  94107":Sept:2017
   | "113 Moon St.
   | San Diego, CA  92093":Jan:2018
   | "51 Belt Ct.
   | Denver, CO  90123":Dec:2016
   | "93114 Radial Rd.
   | Chicago, IL  60605":Jul:2017
   | "7301 Brookview Ave.
   | Columbus, OH  43213":Dec:2018
```

Reading Hive Table Data

Apache Hive is a distributed data warehousing infrastructure. Hive facilitates managing large data sets supporting multiple data formats, including comma-separated value (.csv) TextFile, RCFile, ORC, and Parquet.

The PXF Hive connector reads data stored in a Hive table. This section describes how to use the

PXF Hive connector.

When accessing Hive 3, the PXF Hive connector supports using the `hive[:*]` profiles described below to access Hive 3 external tables only. The Connector does not support using the `hive[:*]` profiles to access Hive 3 managed (CRUD and insert-only transactional, and temporary) tables. Use the PXF JDBC Connector to access Hive 3 managed tables instead.

Prerequisites

Before working with Hive table data using PXF, ensure that you have met the PXF Hadoop Prerequisites.

If you plan to use PXF filter pushdown with Hive integral types, ensure that the configuration parameter `hive.metastore.integral.jdo.pushdown` exists and is set to `true` in the `hive-site.xml` file in both your Hadoop cluster **and** `$PXF_BASE/servers/default/hive-site.xml`. Refer to About Updating Hadoop Configuration for more information.

Hive Data Formats

The PXF Hive connector supports several data formats, and has defined the following profiles for accessing these formats:

File Format	Description	Profile
TextFile	Flat file with data in comma-, tab-, or space-separated value format or JSON notation.	hive, hive:text
SequenceFile	Flat file consisting of binary key/value pairs.	hive
RCFile	Record columnar data consisting of binary key/value pairs; high row compression rate.	hive, hive:rc
ORC	Optimized row columnar data with stripe, footer, and postscript sections; reduces data size.	hive, hive:orc
Parquet	Compressed columnar data representation.	hive

Note: The `hive` profile supports all file storage formats. It will use the optimal `hive[:*]` profile for the underlying file format type.

Data Type Mapping

The PXF Hive connector supports primitive and complex data types.

Primitive Data Types

To represent Hive data in Greenplum Database, map data values that use a primitive data type to Greenplum Database columns of the same type.

The following table summarizes external mapping rules for Hive primitive types.

Hive Data Type	Greenplum Data Type
boolean	bool

Hive Data Type	Greenplum Data Type
int	int4
smallint	int2
tinyint	int2
bigint	int8
float	float4
double	float8
string	text
binary	bytea
timestamp	timestamp

Note: The `hive:orc` profile does not support the timestamp data type when you specify vectorized query execution (`VECTORIZE=true`).

Complex Data Types

Hive supports complex data types including array, struct, map, and union. PXF maps each of these complex types to `text`. You can create Greenplum Database functions or application code to extract subcomponents of these complex data types.

Examples using complex data types with the `hive` and `hive:orc` profiles are provided later in this topic.

Note: The `hive:orc` profile does not support complex types when you specify vectorized query execution (`VECTORIZE=true`).

Sample Data Set

Examples presented in this topic operate on a common data set. This simple data set models a retail sales operation and includes fields with the following names and data types:

Column Name	Data Type
location	text
month	text
number_of_orders	integer
total_sales	double

Prepare the sample data set for use:

1. First, create a text file:

```
$ vi /tmp/pxf_hive_datafile.txt
```

2. Add the following data to `pxf_hive_datafile.txt`; notice the use of the comma `,` to separate the four field values:

```
Prague,Jan,101,4875.33
Rome,Mar,87,1557.39
Bangalore,May,317,8936.99
Beijing,Jul,411,11600.67
San Francisco,Sept,156,6846.34
Paris,Nov,159,7134.56
San Francisco,Jan,113,5397.89
Prague,Dec,333,9894.77
Bangalore,Jul,271,8320.55
Beijing,Dec,100,4248.41
```

Make note of the path to `pxf_hive_datafile.txt`; you will use it in later exercises.

Hive Command Line

The Hive command line is a subsystem similar to that of `psql`. To start the Hive command line:

```
$ HADOOP_USER_NAME=hdfs hive
```

The default Hive database is named `default`.

Example: Creating a Hive Table

Create a Hive table to expose the sample data set.

1. Create a Hive table named `sales_info` in the `default` database:

   ```
   hive> CREATE TABLE sales_info (location string, month string,
           number_of_orders int, total_sales double)
           ROW FORMAT DELIMITED FIELDS TERMINATED BY ','
           STORED AS textfile;
   ```

   Notice that:

   - The `STORED AS textfile` subclause instructs Hive to create the table in Textfile (the default) format. Hive Textfile format supports comma-, tab-, and space-separated values, as well as data specified in JSON notation.
   - The `DELIMITED FIELDS TERMINATED BY` subclause identifies the field delimiter within a data record (line). The `sales_info` table field delimiter is a comma (`,`).

2. Load the `pxf_hive_datafile.txt` sample data file into the `sales_info` table that you just created:

   ```
   hive> LOAD DATA LOCAL INPATH '/tmp/pxf_hive_datafile.txt'
           INTO TABLE sales_info;
   ```

   In examples later in this section, you will access the `sales_info` Hive table directly via PXF. You will also insert `sales_info` data into tables of other Hive file format types, and use PXF to access those directly as well.

3. Perform a query on `sales_info` to verify that you loaded the data successfully:

   ```
   hive> SELECT * FROM sales_info;
   ```

Determining the HDFS Location of a Hive Table

Should you need to identify the HDFS file location of a Hive managed table, reference it using its HDFS file path. You can determine a Hive table's location in HDFS using the `DESCRIBE` command. For example:

```
hive> DESCRIBE EXTENDED sales_info;
Detailed Table Information
...
location:hdfs://<namenode>:<port>/apps/hive/warehouse/sales_info
...
```

Querying External Hive Data

You can create a Greenplum Database external table to access Hive table data. As described previously, the PXF Hive connector defines specific profiles to support different file formats. These profiles are named `hive`, `hive:text`, `hive:rc`, and `hive:orc`.

The `hive:text` and `hive:rc` profiles are specifically optimized for text and RCFile formats, respectively. The `hive:orc` profile is optimized for ORC file formats. The `hive` profile is optimized for all file storage types; you can use the `hive` profile when the underlying Hive table is composed of multiple partitions with differing file formats.

PXF uses column projection to increase query performance when you access a Hive table using the `hive`, `hive:rc`, or `hive:orc` profiles.

Use the following syntax to create a Greenplum Database external table that references a Hive table:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<hive-db-name>.<hive-table-name>
    ?PROFILE=<profile_name>[&SERVER=<server_name>][&PPD=<boolean>][&VECTORIZE=<boolean
>]')
FORMAT 'CUSTOM|TEXT' (FORMATTER='pxfwritable_import' | delimiter='<delim>')
```

Hive connector-specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE call are described below.

Keyword	Value
<hive-db-name>	The name of the Hive database. If omitted, defaults to the Hive database named `default`.
<hive-table-name>	The name of the Hive table.
PROFILE=<profile_name>	`<profile_name>` must specify one of the values `hive`, `hive:text`, `hive:rc`, or `hive:orc`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
PPD=<boolean>	Enable or disable predicate pushdown for all queries on this table; this option applies only to the `hive`, `hive:orc`, and `hive:rc` profiles, and overrides a `pxf.ppd.hive` property setting in the <server_name> configuration.
VECTORIZE=<boolean>	When `PROFILE=hive:orc`, a Boolean value that specifies whether or not PXF uses vectorized query execution when accessing the underlying ORC files. The default value is `false`, does not use vectorized query execution.

Keyword	Value
FORMAT (`hive` and `hive:orc` profiles)	The `FORMAT` clause must specify `'CUSTOM'`. The `CUSTOM` format requires the built-in `pxfwritable_import formatter`.
FORMAT (`hive:text` and `hive:rc` profiles)	The `FORMAT` clause must specify `TEXT`. Specify the single ascii character field delimiter in the `delimiter='<delim>'` formatting option.

Because Hive tables can be backed by one or more files and each file can have a unique layout or schema, PXF requires that the column names that you specify when you create the external table match the column names defined for the Hive table. This allows you to:

- Create the PXF external table with columns in a different order than the Hive table.
- Create a PXF external table that reads a subset of the columns in the Hive table.
- Read a Hive table where the files backing the table have a different number of columns.

Accessing TextFile-Format Hive Tables

You can use the `hive` and `hive:text` profiles to access Hive table data stored in TextFile format.

Example: Using the hive Profile

Use the `hive` profile to create a readable Greenplum Database external table that references the Hive `sales_info` textfile format table that you created earlier.

1. Create the external table:

```
postgres=# CREATE EXTERNAL TABLE salesinfo_hiveprofile(location text, month tex
t, number_of_orders int, total_sales float8)
            LOCATION ('pxf://default.sales_info?PROFILE=hive')
          FORMAT 'custom' (FORMATTER='pxfwritable_import');
```

2. Query the table:

```
postgres=# SELECT * FROM salesinfo_hiveprofile;
```

```
   location     | month | number_of_orders | total_sales
---------------+-------+------------------+-------------
 Prague         | Jan   |              101 |     4875.33
 Rome           | Mar   |               87 |     1557.39
 Bangalore      | May   |              317 |     8936.99
 ...
```

Example: Using the hive:text Profile

Use the PXF `hive:text` profile to create a readable Greenplum Database external table from the Hive `sales_info` textfile format table that you created earlier.

1. Create the external table:

```
postgres=# CREATE EXTERNAL TABLE salesinfo_hivetextprofile(location text, month
 text, number_of_orders int, total_sales float8)
             LOCATION ('pxf://default.sales_info?PROFILE=hive:text')
```

```
            FORMAT 'TEXT' (delimiter=E',');
```

Notice that the `FORMAT` subclause `delimiter` value is specified as the single ascii comma character `','`. `E` escapes the character.

2. Query the external table:

```
postgres=# SELECT * FROM salesinfo_hivetextprofile WHERE location='Beijing';
```

```
 location | month | number_of_orders | total_sales
----------+-------+------------------+-------------
 Beijing  | Jul   |              411 |    11600.67
 Beijing  | Dec   |              100 |     4248.41
(2 rows)
```

Accessing RCFile-Format Hive Tables

The RCFile Hive table format is used for row columnar formatted data. The PXF `hive:rc` profile provides access to RCFile data.

Example: Using the hive:rc Profile

Use the `hive:rc` profile to query RCFile-formatted data in a Hive table.

1. Start the `hive` command line and create a Hive table stored in RCFile format:

```
$ HADOOP_USER_NAME=hdfs hive
```

```
hive> CREATE TABLE sales_info_rcfile (location string, month string,
        number_of_orders int, total_sales double)
      ROW FORMAT DELIMITED FIELDS TERMINATED BY ','
      STORED AS rcfile;
```

2. Insert the data from the `sales_info` table into `sales_info_rcfile`:

```
hive> INSERT INTO TABLE sales_info_rcfile SELECT * FROM sales_info;
```

A copy of the sample data set is now stored in RCFile format in the Hive `sales_info_rcfile` table.

3. Query the `sales_info_rcfile` Hive table to verify that the data was loaded correctly:

```
hive> SELECT * FROM sales_info_rcfile;
```

4. Use the PXF `hive:rc` profile to create a readable Greenplum Database external table that references the Hive `sales_info_rcfile` table that you created in the previous steps. For example:

```
postgres=# CREATE EXTERNAL TABLE salesinfo_hivercprofile(location text, month t
ext, number_of_orders int, total_sales float8)
             LOCATION ('pxf://default.sales_info_rcfile?PROFILE=hive:rc')
           FORMAT 'TEXT' (delimiter=E',');
```

5. Query the external table:

```
postgres=# SELECT location, total_sales FROM salesinfo_hivercprofile;
```

```
   location    | total_sales
---------------+-------------
 Prague        |     4875.33
 Rome          |     1557.39
 Bangalore     |     8936.99
 Beijing       |    11600.67
 ...
```

Accessing ORC-Format Hive Tables

The Optimized Row Columnar (ORC) file format is a columnar file format that provides a highly efficient way to both store and access HDFS data. ORC format offers improvements over text and RCFile formats in terms of both compression and performance. PXF supports ORC version 1.2.1.

ORC is type-aware and specifically designed for Hadoop workloads. ORC files store both the type of and encoding information for the data in the file. All columns within a single group of row data (also known as stripe) are stored together on disk in ORC format files. The columnar nature of the ORC format type enables read projection, helping avoid accessing unnecessary columns during a query.

ORC also supports predicate pushdown with built-in indexes at the file, stripe, and row levels, moving the filter operation to the data loading phase.

Refer to the Apache orc and the Apache Hive LanguageManual ORC websites for detailed information about the ORC file format.

Profiles Supporting the ORC File Format

When choosing an ORC-supporting profile, consider the following:

- The `hive:orc` profile:
    - Reads a single row of data at a time.
    - Supports column projection.
    - Supports complex types. You can access Hive tables composed of array, map, struct, and union data types. PXF serializes each of these complex types to `text`.
- The `hive:orc` profile with `VECTORIZE=true`:
    - Reads up to 1024 rows of data at once.
    - Supports column projection.
    - Does not support complex types or the timestamp data type.

Example: Using the hive:orc Profile

In the following example, you will create a Hive table stored in ORC format and use the `hive:orc` profile to query this Hive table.

1. Create a Hive table with ORC file format:

```
$ HADOOP_USER_NAME=hdfs hive
```

```
hive> CREATE TABLE sales_info_ORC (location string, month string,
        number_of_orders int, total_sales double)
      STORED AS ORC;
```

2. Insert the data from the `sales_info` table into `sales_info_ORC`:

```
hive> INSERT INTO TABLE sales_info_ORC SELECT * FROM sales_info;
```

   A copy of the sample data set is now stored in ORC format in `sales_info_ORC`.

3. Perform a Hive query on `sales_info_ORC` to verify that the data was loaded successfully:

```
hive> SELECT * FROM sales_info_ORC;
```

4. Start the `psql` subsystem and turn on timing:

```
$ psql -d postgres
```

```
postgres=> \timing
Timing is on.
```

5. Use the PXF `hive:orc` profile to create a Greenplum Database external table that references the Hive table named `sales_info_ORC` you created in Step 1. The `FORMAT` clause must specify `'CUSTOM'`. The `hive:orc` `CUSTOM` format supports only the built-in `'pxfwritable_import'` `formatter`.

```
postgres=> CREATE EXTERNAL TABLE salesinfo_hiveORCprofile(location text, month
text, number_of_orders int, total_sales float8)
             LOCATION ('pxf://default.sales_info_ORC?PROFILE=hive:orc')
             FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

6. Query the external table:

```
postgres=> SELECT * FROM salesinfo_hiveORCprofile;
```

```
   location     | month | number_of_orders | total_sales
----------------+-------+------------------+-------------
 Prague         | Jan   |              101 |     4875.33
 Rome           | Mar   |               87 |     1557.39
 Bangalore      | May   |              317 |     8936.99
 ...

Time: 425.416 ms
```

Example: Using the Vectorized hive:orc Profile

In the following example, you will use the vectorized `hive:orc` profile to query the `sales_info_ORC` Hive table that you created in the previous example.

1. Start the `psql` subsystem:

```
$ psql -d postgres
```

2. Use the PXF `hive:orc` profile to create a readable Greenplum Database external table that references the Hive table named `sales_info_ORC` that you created in Step 1 of the previous example. The `FORMAT` clause must specify `'CUSTOM'`. The `hive:orc CUSTOM` format supports only the built-in `'pxfwritable_import' formatter`.

```
postgres=> CREATE EXTERNAL TABLE salesinfo_hiveVectORC(location text, month tex
t, number_of_orders int, total_sales float8)
             LOCATION ('pxf://default.sales_info_ORC?PROFILE=hive:orc&VECTORIZE
=true')
             FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

3. Query the external table:

```
postgres=> SELECT * FROM salesinfo_hiveVectORC;
```

```
   location    | month | number_of_orders | total_sales
---------------+-------+------------------+-------------
 Prague        | Jan   |              101 |     4875.33
 Rome          | Mar   |               87 |     1557.39
 Bangalore     | May   |              317 |     8936.99
 ...

Time: 425.416 ms
```

Accessing Parquet-Format Hive Tables

The PXF `hive` profile supports both non-partitioned and partitioned Hive tables that use the Parquet storage format. Map the table columns using equivalent Greenplum Database data types. For example, if a Hive table is created in the `default` schema using:

```
hive> CREATE TABLE hive_parquet_table (location string, month string,
            number_of_orders int, total_sales double)
        STORED AS parquet;
```

Define the Greenplum Database external table:

```
postgres=# CREATE EXTERNAL TABLE pxf_parquet_table (location text, month text, number_
of_orders int, total_sales double precision)
    LOCATION ('pxf://default.hive_parquet_table?profile=hive')
    FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

And query the table:

```
postgres=# SELECT month, number_of_orders FROM pxf_parquet_table;
```

Accessing Avro-Format Hive Tables

The PXF `hive` profile supports accessing Hive tables that use the Avro storage format. Map the table columns using equivalent Greenplum Database data types. For example, if a Hive table is created in the `default` schema using:

```
hive> CREATE TABLE hive_avro_data_table (id int, name string, user_id string)
    ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.avro.AvroSerDe'
    STORED AS INPUTFORMAT 'org.apache.hadoop.hive.ql.io.avro.AvroContainerInputFormat'
    OUTPUTFORMAT 'org.apache.hadoop.hive.ql.io.avro.AvroContainerOutputFormat';
```

Define the Greenplum Database external table:

```
postgres=# CREATE EXTERNAL TABLE userinfo_hiveavro(id int, name text, user_id text)
    LOCATION ('pxf://default.hive_avro_data_table?profile=hive')
    FORMAT 'custom' (FORMATTER='pxfwritable_import');
```

And query the table:

```
postgres=# SELECT * FROM userinfo_hiveavro;
```

Working with Complex Data Types

Example: Using the hive Profile with Complex Data Types

This example employs the `hive` profile and the array and map complex types, specifically an array of integers and a string key/value pair map.

The data schema for this example includes fields with the following names and data types:

Column Name	Data Type
index	int
name	string
intarray	array of integers
propmap	map of string key and value pairs

When you specify an array field in a Hive table, you must identify the terminator for each item in the collection. Similarly, you must also specify the map key termination character.

1. Create a text file from which you will load the data set:

   ```
   $ vi /tmp/pxf_hive_complex.txt
   ```

2. Add the following text to `pxf_hive_complex.txt`. This data uses a comma `,` to separate field values, the percent symbol `%` to separate collection items, and a `:` to terminate map key values:

   ```
   3,Prague,1%2%3,zone:euro%status:up
   89,Rome,4%5%6,zone:euro
   400,Bangalore,7%8%9,zone:apac%status:pending
   183,Beijing,0%1%2,zone:apac
   94,Sacramento,3%4%5,zone:noam%status:down
   101,Paris,6%7%8,zone:euro%status:up
   56,Frankfurt,9%0%1,zone:euro
   202,Jakarta,2%3%4,zone:apac%status:up
   313,Sydney,5%6%7,zone:apac%status:pending
   76,Atlanta,8%9%0,zone:noam%status:down
   ```

3. Create a Hive table to represent this data:

```
$ HADOOP_USER_NAME=hdfs hive
```

```
hive> CREATE TABLE table_complextypes( index int, name string, intarray ARRAY<i
nt>, propmap MAP<string, string>)
         ROW FORMAT DELIMITED FIELDS TERMINATED BY ','
         COLLECTION ITEMS TERMINATED BY '%'
         MAP KEYS TERMINATED BY ':'
         STORED AS TEXTFILE;
```

Notice that:

   - `FIELDS TERMINATED BY` identifies a comma as the field terminator.
   - The `COLLECTION ITEMS TERMINATED BY` subclause specifies the percent sign as the collection items (array item, map key/value pair) terminator.
   - `MAP KEYS TERMINATED BY` identifies a colon as the terminator for map keys.

4. Load the `pxf_hive_complex.txt` sample data file into the `table_complextypes` table that you just created:

```
hive> LOAD DATA LOCAL INPATH '/tmp/pxf_hive_complex.txt' INTO TABLE table_compl
extypes;
```

5. Perform a query on Hive table `table_complextypes` to verify that the data was loaded successfully:

```
hive> SELECT * FROM table_complextypes;
```

```
3   Prague   [1,2,3] {"zone":"euro","status":"up"}
89  Rome     [4,5,6] {"zone":"euro"}
400 Bangalore   [7,8,9] {"zone":"apac","status":"pending"}
...
```

6. Use the PXF `hive` profile to create a readable Greenplum Database external table that references the Hive table named `table_complextypes`:

```
postgres=# CREATE EXTERNAL TABLE complextypes_hiveprofile(index int, name text,
 intarray text, propmap text)
             LOCATION ('pxf://table_complextypes?PROFILE=hive')
           FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Notice that the integer array and map complex types are mapped to Greenplum Database data type text.

7. Query the external table:

```
postgres=# SELECT * FROM complextypes_hiveprofile;
```

```
 index |    name    | intarray |              propmap
-------+------------+----------+------------------------------------
     3 | Prague     | [1,2,3]  | {"zone":"euro","status":"up"}
    89 | Rome       | [4,5,6]  | {"zone":"euro"}
```

```
  400 | Bangalore  | [7,8,9]  | {"zone":"apac","status":"pending"}
  183 | Beijing    | [0,1,2]  | {"zone":"apac"}
   94 | Sacramento | [3,4,5]  | {"zone":"noam","status":"down"}
  101 | Paris      | [6,7,8]  | {"zone":"euro","status":"up"}
   56 | Frankfurt  | [9,0,1]  | {"zone":"euro"}
  202 | Jakarta    | [2,3,4]  | {"zone":"apac","status":"up"}
  313 | Sydney     | [5,6,7]  | {"zone":"apac","status":"pending"}
   76 | Atlanta    | [8,9,0]  | {"zone":"noam","status":"down"}
(10 rows)
```

`intarray` and `propmap` are each serialized as text strings.

Example: Using the hive:orc Profile with Complex Data Types

In the following example, you will create and populate a Hive table stored in ORC format. You will use the `hive:orc` profile to query the complex types in this Hive table.

1. Create a Hive table with ORC storage format:

```
$ HADOOP_USER_NAME=hdfs hive
```

```
hive> CREATE TABLE table_complextypes_ORC( index int, name string, intarray ARR
AY<int>, propmap MAP<string, string>)
        ROW FORMAT DELIMITED FIELDS TERMINATED BY ','
        COLLECTION ITEMS TERMINATED BY '%'
        MAP KEYS TERMINATED BY ':'
      STORED AS ORC;
```

2. Insert the data from the `table_complextypes` table that you created in the previous example into `table_complextypes_ORC`:

```
hive> INSERT INTO TABLE table_complextypes_ORC SELECT * FROM table_complextypes
;
```

   A copy of the sample data set is now stored in ORC format in `table_complextypes_ORC`.

3. Perform a Hive query on `table_complextypes_ORC` to verify that the data was loaded successfully:

```
hive> SELECT * FROM table_complextypes_ORC;
```

```
OK
3       Prague       [1,2,3]    {"zone":"euro","status":"up"}
89      Rome         [4,5,6]    {"zone":"euro"}
400     Bangalore    [7,8,9]    {"zone":"apac","status":"pending"}
...
```

4. Start the `psql` subsystem:

```
$ psql -d postgres
```

5. Use the PXF `hive:orc` profile to create a readable Greenplum Database external table from the Hive table named `table_complextypes_ORC` you created in Step 1. The `FORMAT` clause must specify `'CUSTOM'`. The `hive:orc` `CUSTOM` format supports only the built-in

`'pxfwritable_import' formatter`.

```
postgres=> CREATE EXTERNAL TABLE complextypes_hiveorc(index int, name text, int
array text, propmap text)
           LOCATION ('pxf://default.table_complextypes_ORC?PROFILE=hive:orc')
             FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Notice that the integer array and map complex types are again mapped to Greenplum Database data type text.

6. Query the external table:

```
postgres=> SELECT * FROM complextypes_hiveorc;
```

```
 index |    name    | intarray |              propmap
-------+------------+----------+------------------------------------
     3 | Prague     | [1,2,3]  | {"zone":"euro","status":"up"}
    89 | Rome       | [4,5,6]  | {"zone":"euro"}
   400 | Bangalore  | [7,8,9]  | {"zone":"apac","status":"pending"}
 ...
```

`intarray` and `propmap` are again serialized as text strings.

Partition Pruning

The PXF Hive connector supports Hive partition pruning and the Hive partition directory structure. This enables partition exclusion on selected HDFS files comprising a Hive table. To use the partition filtering feature to reduce network traffic and I/O, run a query on a PXF external table using a `WHERE` clause that refers to a specific partition column in a partitioned Hive table.

The PXF Hive Connector partition filtering support for Hive string and integral types is described below:

- The relational operators `=`, `<`, `<=`, `>`, `>=`, and `<>` are supported on string types.
- The relational operators `=` and `<>` are supported on integral types (To use partition filtering with Hive integral types, you must update the Hive configuration as described in the Prerequisites).
- The logical operators `AND` and `OR` are supported when used with the relational operators mentioned above.
- The `LIKE` string operator is not supported.

To take advantage of PXF partition filtering pushdown, the Hive and PXF partition field names must be the same. Otherwise, PXF ignores partition filtering and the filtering is performed on the Greenplum Database side, impacting performance.

The PXF Hive connector filters only on partition columns, not on other table attributes. Additionally, filter pushdown is supported only for those data types and operators identified above.

PXF filter pushdown is enabled by default. You configure PXF filter pushdown as described in About Filter Pushdown.

Example: Using the hive Profile to Access Partitioned Homogenous

Data

In this example, you use the `hive` profile to query a Hive table named `sales_part` that you partition on the `delivery_state` and `delivery_city` fields. You then create a Greenplum Database external table to query `sales_part`. The procedure includes specific examples that illustrate filter pushdown.

1. Create a Hive table named `sales_part` with two partition columns, `delivery_state` and `delivery_city`:

    ```
    hive> CREATE TABLE sales_part (cname string, itype string, supplier_key int, pr
    ice double)
            PARTITIONED BY (delivery_state string, delivery_city string)
            ROW FORMAT DELIMITED FIELDS TERMINATED BY ',';
    ```

2. Load data into this Hive table and add some partitions:

    ```
    hive> INSERT INTO TABLE sales_part
            PARTITION(delivery_state = 'CALIFORNIA', delivery_city = 'Fresno')
            VALUES ('block', 'widget', 33, 15.17);
    hive> INSERT INTO TABLE sales_part
            PARTITION(delivery_state = 'CALIFORNIA', delivery_city = 'Sacramento')
            VALUES ('cube', 'widget', 11, 1.17);
    hive> INSERT INTO TABLE sales_part
            PARTITION(delivery_state = 'NEVADA', delivery_city = 'Reno')
            VALUES ('dowel', 'widget', 51, 31.82);
    hive> INSERT INTO TABLE sales_part
            PARTITION(delivery_state = 'NEVADA', delivery_city = 'Las Vegas')
            VALUES ('px49', 'pipe', 52, 99.82);
    ```

3. Query the `sales_part` table:

    ```
    hive> SELECT * FROM sales_part;
    ```

    A `SELECT *` statement on a Hive partitioned table shows the partition fields at the end of the record.

4. Examine the Hive/HDFS directory structure for the `sales_part` table:

    ```
    $ sudo -u hdfs hdfs dfs -ls -R /apps/hive/warehouse/sales_part
    /apps/hive/warehouse/sales_part/delivery_state=CALIFORNIA/delivery_city=Fresno/
    /apps/hive/warehouse/sales_part/delivery_state=CALIFORNIA/delivery_city=Sacrame
    nto/
    /apps/hive/warehouse/sales_part/delivery_state=NEVADA/delivery_city=Reno/
    /apps/hive/warehouse/sales_part/delivery_state=NEVADA/delivery_city=Las Vegas/
    ```

5. Create a PXF external table to read the partitioned `sales_part` Hive table. To take advantage of partition filter push-down, define fields corresponding to the Hive partition fields at the end of the `CREATE EXTERNAL TABLE` attribute list.

    ```
    $ psql -d postgres
    ```

    ```
    postgres=# CREATE EXTERNAL TABLE pxf_sales_part(
                 cname TEXT, itype TEXT,
                 supplier_key INTEGER, price DOUBLE PRECISION,
                 delivery_state TEXT, delivery_city TEXT)
    ```

```
            LOCATION ('pxf://sales_part?PROFILE=hive')
            FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

6. Query the table:

```
postgres=# SELECT * FROM pxf_sales_part;
```

7. Perform another query (no pushdown) on `pxf_sales_part` to return records where the `delivery_city` is `Sacramento` and `cname` is `cube`:

```
postgres=# SELECT * FROM pxf_sales_part WHERE delivery_city = 'Sacramento' AND
cname = 'cube';
```

   The query filters the `delivery_city` partition `Sacramento`. The filter on `cname` is not pushed down, since it is not a partition column. It is performed on the Greenplum Database side after all the data in the `Sacramento` partition is transferred for processing.

8. Query (with pushdown) for all records where `delivery_state` is `CALIFORNIA`:

```
postgres=# SET gp_external_enable_filter_pushdown=on;
postgres=# SELECT * FROM pxf_sales_part WHERE delivery_state = 'CALIFORNIA';
```

   This query reads all of the data in the `CALIFORNIA` `delivery_state` partition, regardless of the city.

Example: Using the hive Profile to Access Partitioned Heterogeneous Data

You can use the PXF `hive` profile with any Hive file storage types. With the `hive` profile, you can access heterogeneous format data in a single Hive table where the partitions may be stored in different file formats.

In this example, you create a partitioned Hive external table. The table is composed of the HDFS data files associated with the `sales_info` (text format) and `sales_info_rcfile` (RC format) Hive tables that you created in previous exercises. You will partition the data by year, assigning the data from `sales_info` to the year 2013, and the data from `sales_info_rcfile` to the year 2016. (Ignore at the moment the fact that the tables contain the same data.) You will then use the PXF `hive` profile to query this partitioned Hive external table.

1. Create a Hive external table named `hive_multiformpart` that is partitioned by a string field named `year`:

```
$ HADOOP_USER_NAME=hdfs hive
```

```
hive> CREATE EXTERNAL TABLE hive_multiformpart( location string, month string,
number_of_orders int, total_sales double)
        PARTITIONED BY( year string )
        ROW FORMAT DELIMITED FIELDS TERMINATED BY ',';
```

2. Describe the `sales_info` and `sales_info_rcfile` tables, noting the HDFS file `location` for each table:

```
hive> DESCRIBE EXTENDED sales_info;
```

```
hive> DESCRIBE EXTENDED sales_info_rcfile;
```

3. Create partitions in the `hive_multiformpart` table for the HDFS file locations associated with each of the `sales_info` and `sales_info_rcfile` tables:

```
hive> ALTER TABLE hive_multiformpart ADD PARTITION (year = '2013') LOCATION 'hd
fs://namenode:8020/apps/hive/warehouse/sales_info';
hive> ALTER TABLE hive_multiformpart ADD PARTITION (year = '2016') LOCATION 'hd
fs://namenode:8020/apps/hive/warehouse/sales_info_rcfile';
```

4. Explicitly identify the file format of the partition associated with the `sales_info_rcfile` table:

```
hive> ALTER TABLE hive_multiformpart PARTITION (year='2016') SET FILEFORMAT RCF
ILE;
```

   You need not specify the file format of the partition associated with the `sales_info` table, as `TEXTFILE` format is the default.

5. Query the `hive_multiformpart` table:

```
hive> SELECT * from hive_multiformpart;
...
Bangalore   Jul 271 8320.55 2016
Beijing Dec 100 4248.41 2016
Prague  Jan 101 4875.33 2013
Rome    Mar 87  1557.39 2013
...
hive> SELECT * from hive_multiformpart WHERE year='2013';
hive> SELECT * from hive_multiformpart WHERE year='2016';
```

6. Show the partitions defined for the `hive_multiformpart` table and exit `hive`:

```
hive> SHOW PARTITIONS hive_multiformpart;
year=2013
year=2016
hive> quit;
```

7. Start the `psql` subsystem:

```
$ psql -d postgres
```

8. Use the PXF `hive` profile to create a readable Greenplum Database external table that references the Hive `hive_multiformpart` external table that you created in the previous steps:

```
postgres=# CREATE EXTERNAL TABLE pxf_multiformpart(location text, month text, n
umber_of_orders int, total_sales float8, year text)
             LOCATION ('pxf://default.hive_multiformpart?PROFILE=hive')
           FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

9. Query the PXF external table:

```
postgres=# SELECT * FROM pxf_multiformpart;
```

```
   location    | month | number_of_orders | total_sales | year
---------------+-------+------------------+-------------+--------
 ....
 Prague        | Dec   |              333 |     9894.77 | 2013
 Bangalore     | Jul   |              271 |     8320.55 | 2013
 Beijing       | Dec   |              100 |     4248.41 | 2013
 Prague        | Jan   |              101 |     4875.33 | 2016
 Rome          | Mar   |               87 |     1557.39 | 2016
 Bangalore     | May   |              317 |     8936.99 | 2016
 ....
```

10. Perform a second query to calculate the total number of orders for the year 2013:

```
postgres=# SELECT sum(number_of_orders) FROM pxf_multiformpart WHERE month='Dec
' AND year='2013';
 sum
-----
 433
```

Using PXF with Hive Default Partitions

This topic describes a difference in query results between Hive and PXF queries when Hive tables use a default partition. When dynamic partitioning is enabled in Hive, a partitioned table may store data in a default partition. Hive creates a default partition when the value of a partitioning column does not match the defined type of the column (for example, when a NULL value is used for any partitioning column). In Hive, any query that includes a filter on a partition column *excludes* any data that is stored in the table's default partition.

Similar to Hive, PXF represents a table's partitioning columns as columns that are appended to the end of the table. However, PXF translates any column value in a default partition to a NULL value. This means that a Greenplum Database query that includes an `IS NULL` filter on a partitioning column can return different results than the same Hive query.

Consider a Hive partitioned table that is created with the statement:

```
hive> CREATE TABLE sales (order_id bigint, order_amount float) PARTITIONED BY (xdate d
ate);
```

The table is loaded with five rows that contain the following data:

```
1.0    1900-01-01
2.2    1994-04-14
3.3    2011-03-31
4.5    NULL
5.0    2013-12-06
```

Inserting row 4 creates a Hive default partition, because the partition column `xdate` contains a null value.

In Hive, any query that filters on the partition column omits data in the default partition. For example, the following query returns no rows:

```
hive> SELECT * FROM sales WHERE xdate IS null;
```

However, if you map this Hive table to a PXF external table in Greenplum Database, all default

partition values are translated into actual NULL values. In Greenplum Database, executing the same query against the PXF external table returns row 4 as the result, because the filter matches the NULL value.

Keep this behavior in mind when you execute `IS NULL` queries on Hive partitioned tables.

Reading HBase Table Data

Apache HBase is a distributed, versioned, non-relational database on Hadoop.

The PXF HBase connector reads data stored in an HBase table. The HBase connector supports filter pushdown.

This section describes how to use the PXF HBase connector.

Prerequisites

Before working with HBase table data, ensure that you have:

- Copied `$PXF_HOME/share/pxf-hbase-*.jar` to each node in your HBase cluster, and that the location of this PXF JAR file is in the `$HBASE_CLASSPATH`. This configuration is required for the PXF HBase connector to support filter pushdown.
- Met the PXF Hadoop Prerequisites.

HBase Primer

This topic assumes that you have a basic understanding of the following HBase concepts:

- An HBase column includes two components: a column family and a column qualifier. These components are delimited by a colon `:` character, <column-family>:<column-qualifier>.
- An HBase row consists of a row key and one or more column values. A row key is a unique identifier for the table row.
- An HBase table is a multi-dimensional map comprised of one or more columns and rows of data. You specify the complete set of column families when you create an HBase table.
- An HBase cell is comprised of a row (column family, column qualifier, column value) and a timestamp. The column value and timestamp in a given cell represent a version of the value.

For detailed information about HBase, refer to the Apache HBase Reference Guide.

HBase Shell

The HBase shell is a subsystem similar to that of `psql`. To start the HBase shell:

```
$ hbase shell
<hbase output>
hbase(main):001:0>
```

The default HBase namespace is named `default`.

Example: Creating an HBase Table

Create a sample HBase table.

1. Create an HBase table named `order_info` in the `default` namespace. `order_info` has two column families: `product` and `shipping_info`:

```
hbase(main):> create 'order_info', 'product', 'shipping_info'
```

2. The `order_info` `product` column family has qualifiers named `name` and `location`. The `shipping_info` column family has qualifiers named `state` and `zipcode`. Add some data to the `order_info` table:

```
put 'order_info', '1', 'product:name', 'tennis racquet'
put 'order_info', '1', 'product:location', 'out of stock'
put 'order_info', '1', 'shipping_info:state', 'CA'
put 'order_info', '1', 'shipping_info:zipcode', '12345'
put 'order_info', '2', 'product:name', 'soccer ball'
put 'order_info', '2', 'product:location', 'on floor'
put 'order_info', '2', 'shipping_info:state', 'CO'
put 'order_info', '2', 'shipping_info:zipcode', '56789'
put 'order_info', '3', 'product:name', 'snorkel set'
put 'order_info', '3', 'product:location', 'warehouse'
put 'order_info', '3', 'shipping_info:state', 'OH'
put 'order_info', '3', 'shipping_info:zipcode', '34567'
```

   You will access the `orders_info` HBase table directly via PXF in examples later in this topic.

3. Display the contents of the `order_info` table:

```
hbase(main):> scan 'order_info'
ROW     COLUMN+CELL
 1      column=product:location, timestamp=1499074825516, value=out of stock

 1      column=product:name, timestamp=1499074825491, value=tennis racquet

 1      column=shipping_info:state, timestamp=1499074825531, value=CA

 1      column=shipping_info:zipcode, timestamp=1499074825548, value=12345

 2      column=product:location, timestamp=1499074825573, value=on floor
 ...
3 row(s) in 0.0400 seconds
```

Querying External HBase Data

The PXF HBase connector supports a single profile named `hbase`.

Use the following syntax to create a Greenplum Database external table that references an HBase table:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<hbase-table-name>?PROFILE=hbase[&SERVER=<server_name>]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

HBase connector-specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE call are described below.

Keyword	Value
<hbase-table-name>	The name of the HBase table.
PROFILE	The `PROFILE` keyword must specify `hbase`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.
FORMAT	The `FORMAT` clause must specify `'CUSTOM' (FORMATTER='pxfwritable_import')`.

Data Type Mapping

HBase is byte-based; it stores all data types as an array of bytes. To represent HBase data in Greenplum Database, select a data type for your Greenplum Database column that matches the underlying content of the HBase column qualifier values.

Note: PXF does not support complex HBase objects.

Column Mapping

You can create a Greenplum Database external table that references all, or a subset of, the column qualifiers defined in an HBase table. PXF supports direct or indirect mapping between a Greenplum Database table column and an HBase table column qualifier.

Direct Mapping

When you use direct mapping to map Greenplum Database external table column names to HBase qualifiers, you specify column-family-qualified HBase qualifier names as quoted values. The PXF HBase connector passes these column names as-is to HBase as it reads the table data.

For example, to create a Greenplum Database external table accessing the following data:

- qualifier `name` in the column family named `product`
- qualifier `zipcode` in the column family named `shipping_info`

from the `order_info` HBase table that you created in Example: Creating an HBase Table, use this `CREATE EXTERNAL TABLE` syntax:

```
CREATE EXTERNAL TABLE orderinfo_hbase ("product:name" varchar, "shipping_info:zipcode"
 int)
    LOCATION ('pxf://order_info?PROFILE=hbase')
    FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Indirect Mapping via Lookup Table

When you use indirect mapping to map Greenplum Database external table column names to HBase qualifiers, you specify the mapping in a lookup table that you create in HBase. The lookup table maps a <column-family>:<column-qualifier> to a column name alias that you specify when you create the Greenplum Database external table.

You must name the HBase PXF lookup table `pxflookup`. And you must define this table with a single column family named `mapping`. For example:

```
hbase(main):> create 'pxflookup', 'mapping'
```

While the direct mapping method is fast and intuitive, using indirect mapping allows you to create a shorter, character-based alias for the HBase <column-family>:<column-qualifier> name. This better reconciles HBase column qualifier names with Greenplum Database due to the following:

- HBase qualifier names can be very long. Greenplum Database has a 63 character limit on the size of the column name.
- HBase qualifier names can include binary or non-printable characters. Greenplum Database column names are character-based.

When populating the `pxflookup` HBase table, add rows to the table such that the:

- row key specifies the HBase table name
- `mapping` column family qualifier identifies the Greenplum Database column name, and the value identifies the HBase `<column-family>:<column-qualifier>` for which you are creating the alias.

For example, to use indirect mapping with the `order_info` table, add these entries to the `pxflookup` table:

```
hbase(main):> put 'pxflookup', 'order_info', 'mapping:pname', 'product:name'
hbase(main):> put 'pxflookup', 'order_info', 'mapping:zip', 'shipping_info:zipcode'
```

Then create a Greenplum Database external table using the following `CREATE EXTERNAL TABLE` syntax:

```
CREATE EXTERNAL TABLE orderinfo_map (pname varchar, zip int)
    LOCATION ('pxf://order_info?PROFILE=hbase')
    FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Row Key

The HBase table row key is a unique identifier for the table row. PXF handles the row key in a special way.

To use the row key in the Greenplum Database external table query, define the external table using the PXF reserved column named `recordkey.` The `recordkey` column name instructs PXF to return the HBase table record key for each row.

Define the `recordkey` using the Greenplum Database data type `bytea`.

For example:

```
CREATE EXTERNAL TABLE <table_name> (recordkey bytea, ... )
    LOCATION ('pxf://<hbase_table_name>?PROFILE=hbase')
    FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

After you have created the external table, you can use the `recordkey` in a `WHERE` clause to filter the HBase table on a range of row key values.

Note: To enable filter pushdown on the `recordkey`, define the field as `text`.

Accessing Azure, Google Cloud Storage, Minio, and S3 Object Stores with PXF

PXF is installed with connectors to Azure Blob Storage, Azure Data Lake, Google Cloud Storage, Minio, and S3 object stores.

Prerequisites

Before working with object store data using PXF, ensure that:

- You have configured PXF, and PXF is running on each Greenplum Database host. See Configuring PXF for additional information.
- You have configured the PXF Object Store Connectors that you plan to use. Refer to Configuring Connectors to Azure and Google Cloud Storage Object Stores and Configuring Connectors to Minio and S3 Object Stores for instructions.
- Time is synchronized between the Greenplum Database hosts and the external object store systems.

Connectors, Data Formats, and Profiles

The PXF object store connectors provide built-in profiles to support the following data formats:

- Text
- CSV
- Avro
- JSON
- ORC
- Parquet
- AvroSequenceFile
- SequenceFile

The PXF connectors to Azure expose the following profiles to read, and in many cases write, these supported data formats:

Data Format	Azure Blob Storage	Azure Data Lake	Supported Operations
delimited single line plain text	wasbs:text	adl:text	Read, Write
delimited single line comma-separated values of plain text	wasbs:csv	adl:csv	Read, Write

Data Format	Azure Blob Storage	Azure Data Lake	Supported Operations
delimited text with quoted linefeeds	wasbs:text:multi	adl:text:multi	Read
Avro	wasbs:avro	adl:avro	Read, Write
JSON	wasbs:json	adl:json	Read
ORC	wasbs:orc	adl:orc	Read
Parquet	wasbs:parquet	adl:parquet	Read, Write
AvroSequenceFile	wasbs:AvroSequenceFile	adl:AvroSequenceFile	Read, Write
SequenceFile	wasbs:SequenceFile	adl:SequenceFile	Read, Write

Similarly, the PXF connectors to Google Cloud Storage, Minio, and S3 expose these profiles:

Data Format	Google Cloud Storage	S3 or Minio	Supported Operations
delimited single line plain text	gs:text	s3:text	Read, Write
delimited single line comma-separated values of plain text	gs:csv	s3:csv	Read, Write
delimited text with quoted linefeeds	gs:text:multi	s3:text:multi	Read
Avro	gs:avro	s3:avro	Read, Write
JSON	gs:json	s3:json	Read
ORC	gs:orc	s3:orc	Read
Parquet	gs:parquet	s3:parquet	Read, Write
AvroSequenceFile	gs:AvroSequenceFile	s3:AvroSequenceFile	Read, Write
SequenceFile	gs:SequenceFile	s3:SequenceFile	Read, Write

You provide the profile name when you specify the `pxf` protocol on a `CREATE EXTERNAL TABLE` command to create a Greenplum Database external table that references a file or directory in the specific object store.

Sample CREATE EXTERNAL TABLE Commands

When you create an external table that references a file or directory in an object store, you must specify a `SERVER` in the `LOCATION` URI.

The following command creates an external table that references a text file on S3. It specifies the profile named `s3:text` and the server configuration named `s3srvcfg`:

```
CREATE EXTERNAL TABLE pxf_s3_text(location text, month text, num_orders int, total_sales float8)
  LOCATION ('pxf://S3_BUCKET/pxf_examples/pxf_s3_simple.txt?PROFILE=s3:text&SERVER=s3srvcfg')
FORMAT 'TEXT' (delimiter=E',');
```

The following command creates an external table that references a text file on Azure Blob Storage. It

specifies the profile named `wasbs:text` and the server configuration named `wasbssrvcfg`. You would provide the Azure Blob Storage container identifier and your Azure Blob Storage account name.

```
CREATE EXTERNAL TABLE pxf_wasbs_text(location text, month text, num_orders int, total_sales float8)
  LOCATION ('pxf://AZURE_CONTAINER@YOUR_AZURE_BLOB_STORAGE_ACCOUNT_NAME.blob.core.windows.net/path/to/blob/file?PROFILE=wasbs:text&SERVER=wasbssrvcfg')
FORMAT 'TEXT';
```

The following command creates an external table that references a text file on Azure Data Lake. It specifies the profile named `adl:text` and the server configuration named `adlsrvcfg`. You would provide your Azure Data Lake account name.

```
CREATE EXTERNAL TABLE pxf_adl_text(location text, month text, num_orders int, total_sales float8)
  LOCATION ('pxf://YOUR_ADL_ACCOUNT_NAME.azuredatalakestore.net/path/to/file?PROFILE=adl:text&SERVER=adlsrvcfg')
FORMAT 'TEXT';
```

The following command creates an external table that references a JSON file on Google Cloud Storage. It specifies the profile named `gs:json` and the server configuration named `gcssrvcfg`:

```
CREATE EXTERNAL TABLE pxf_gsc_json(location text, month text, num_orders int, total_sales float8)
  LOCATION ('pxf://dir/subdir/file.json?PROFILE=gs:json&SERVER=gcssrvcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

About Accessing the S3 Object Store

PXF is installed with a connector to the S3 object store. PXF supports the following additional runtime features with this connector:

- Overriding the S3 credentials specified in the server configuration by providing them in the `CREATE EXTERNAL TABLE` command DDL.
- Using the Amazon S3 Select service to read certain CSV and Parquet data from S3.

Overriding the S3 Server Configuration with DDL

If you are accessing an S3-compatible object store, you can override the credentials in an S3 server configuration by directly specifying the S3 access ID and secret key via these custom options in the `CREATE EXTERNAL TABLE LOCATION` clause:

Custom Option	Value Description
accesskey	The AWS account access key ID.
secretkey	The secret key associated with the AWS access key ID.

For example:

```
CREATE EXTERNAL TABLE pxf_ext_tbl(name text, orders int)
  LOCATION ('pxf://S3_BUCKET/dir/file.txt?PROFILE=s3:text&SERVER=s3srvcfg&accesskey=YOURKEY&secretkey=YOURSECRET')
```

```
FORMAT 'TEXT' (delimiter=E',');
```

Credentials that you provide in this manner are visible as part of the external table definition. Do not use this method of passing credentials in a production environment.

PXF does not support overriding Azure, Google Cloud Storage, and Minio server credentials in this manner at this time.

Refer to Configuration Property Precedence for detailed information about the precedence rules that PXF uses to obtain configuration property settings for a Greenplum Database user.

Using the Amazon S3 Select Service

Refer to Reading CSV and Parquet Data from S3 Using S3 Select for specific information on how PXF can use the Amazon S3 Select service to read CSV and Parquet files stored on S3.

Reading and Writing Text Data in an Object Store

The PXF object store connectors support plain delimited and comma-separated value format text data. This section describes how to use PXF to access text data in an object store, including how to create, query, and insert data into an external table that references files in the object store.

Note: Accessing text data from an object store is very similar to accessing text data in HDFS.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from or write data to an object store.

Reading Text Data

Use the `<objstore>:text` profile when you read plain text delimited and `<objstore>:csv` when reading .csv data from an object store where each row is a single record. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

The following syntax creates a Greenplum Database readable external table that references a simple text file in an object store:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>?PROFILE=<objstore>:text|csv&SERVER=<server_name>[&IGNO
RE_MISSING_PATH=<boolean>][&SKIP_HEADER_COUNT=<numlines>][&<custom-option>=<value>[...
]]')
```

```
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-file>	The path to the directory or file in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:text PROFILE=<objstore>:csv	The `PROFILE` keyword must identify the specific object store. For example, `s3:text`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
SKIP_HEADER_COUNT=<numlines>	Specify the number of header lines that PXF should skip in the first split of each <file> before reading the data. The default value is 0, do not skip any lines.
FORMAT	Use `FORMAT 'TEXT'` when <path-to-file> references plain text delimited data. Use `FORMAT 'CSV'` when <path-to-file> references comma-separated value data.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.

Note: PXF does not support the `(HEADER)` formatter option in the `CREATE EXTERNAL TABLE` command. If your text file includes header line(s), use `SKIP_HEADER_COUNT` to specify the number of lines that PXF should skip at the beginning of the first split of each file.

If you are accessing an S3 object store:

- You can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.
- If you are reading CSV-format data from S3, you can direct PXF to use the S3 Select Amazon service to retrieve the data. Refer to Using the Amazon S3 Select Service for more information about the PXF custom option used for this purpose.

Example: Reading Text Data from S3

Perform the following procedure to create a sample text file, copy the file to S3, and use the `s3:text` and `s3:csv` profiles to create two PXF external tables to query the data.

To run this example, you must:

- Have the AWS CLI tools installed on your system
- Know your AWS access ID and secret key
- Have write permission to an S3 bucket

1. Create a directory in S3 for PXF example data files. For example, if you have write access to an S3 bucket named `BUCKET`:

```
$ aws s3 mb s3://BUCKET/pxf_examples
```

2. Locally create a delimited plain text data file named `pxf_s3_simple.txt`:

```
$ echo 'Prague,Jan,101,4875.33
Rome,Mar,87,1557.39
Bangalore,May,317,8936.99
Beijing,Jul,411,11600.67' > /tmp/pxf_s3_simple.txt
```

   Note the use of the comma `,` to separate the four data fields.

3. Copy the data file to the S3 directory you created in Step 1:

```
$ aws s3 cp /tmp/pxf_s3_simple.txt s3://BUCKET/pxf_examples/
```

4. Verify that the file now resides in S3:

```
$ aws s3 ls s3://BUCKET/pxf_examples/pxf_s3_simple.txt
```

5. Start the `psql` subsystem:

```
$ psql -d postgres
```

6. Use the PXF `s3:text` profile to create a Greenplum Database external table that references the `pxf_s3_simple.txt` file that you just created and added to S3. For example, if your server name is `s3srvcfg`:

```
postgres=# CREATE EXTERNAL TABLE pxf_s3_textsimple(location text, month text, n
um_orders int, total_sales float8)
            LOCATION ('pxf://BUCKET/pxf_examples/pxf_s3_simple.txt?PROFILE=s3:t
ext&SERVER=s3srvcfg')
          FORMAT 'TEXT' (delimiter=E',');
```

7. Query the external table:

```
postgres=# SELECT * FROM pxf_s3_textsimple;
```

```
   location    | month | num_orders | total_sales
---------------+-------+------------+-------------
 Prague        | Jan   |        101 |     4875.33
 Rome          | Mar   |         87 |     1557.39
 Bangalore     | May   |        317 |     8936.99
 Beijing       | Jul   |        411 |    11600.67
(4 rows)
```

8. Create a second external table that references `pxf_s3_simple.txt`, this time specifying the `s3:csv` `PROFILE` and the `CSV` `FORMAT`:

```
postgres=# CREATE EXTERNAL TABLE pxf_s3_textsimple_csv(location text, month tex
t, num_orders int, total_sales float8)
            LOCATION ('pxf://BUCKET/pxf_examples/pxf_s3_simple.txt?PROFILE=s3:c
sv&SERVER=s3srvcfg')
          FORMAT 'CSV';
postgres=# SELECT * FROM pxf_s3_textsimple_csv;
```

When you specify `FORMAT 'CSV'` for comma-separated value data, no `delimiter` formatter option is required because comma is the default delimiter value.

Reading Text Data with Quoted Linefeeds

Use the `<objstore>:text:multi` profile to read plain text data with delimited single- or multi- line records that include embedded (quoted) linefeed characters. The following syntax creates a Greenplum Database readable external table that references such a text file in an object store:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>?PROFILE=<objstore>:text:multi&SERVER=<server_name>[&IG
NORE_MISSING_PATH=<boolean>][&SKIP_HEADER_COUNT=<numlines>][&<custom-option>=<value>[.
..]]')
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
```

The specific keywords and values used in the CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-file>	The path to the directory or file in the data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:text:multi	The `PROFILE` keyword must identify the specific object store. For example, `s3:text:multi`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-file> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
SKIP_HEADER_COUNT=<numlines>	Specify the number of header lines that PXF should skip in the first split of each <file> before reading the data. The default value is 0, do not skip any lines.
FORMAT	Use `FORMAT 'TEXT'` when <path-to-file> references plain text delimited data. Use `FORMAT 'CSV'` when <path-to-file> references comma-separated value data.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.

Note: PXF does not support the `(HEADER)` formatter option in the `CREATE EXTERNAL TABLE` command. If your text file includes header line(s), use `SKIP_HEADER_COUNT` to specify the number of lines that PXF should skip at the beginning of the first split of each file.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example: Reading Multi-Line Text Data from S3

Perform the following steps to create a sample text file, copy the file to S3, and use the PXF `s3:text:multi` profile to create a Greenplum Database readable external table to query the data.

To run this example, you must:

- Have the AWS CLI tools installed on your system
- Know your AWS access ID and secret key
- Have write permission to an S3 bucket

1. Create a second delimited plain text file:

```
$ vi /tmp/pxf_s3_multi.txt
```

2. Copy/paste the following data into `pxf_s3_multi.txt`:

```
"4627 Star Rd.
San Francisco, CA  94107":Sept:2017
"113 Moon St.
San Diego, CA  92093":Jan:2018
"51 Belt Ct.
Denver, CO  90123":Dec:2016
"93114 Radial Rd.
Chicago, IL  60605":Jul:2017
"7301 Brookview Ave.
Columbus, OH  43213":Dec:2018
```

Notice the use of the colon `:` to separate the three fields. Also notice the quotes around the first (address) field. This field includes an embedded line feed separating the street address from the city and state.

3. Copy the text file to S3:

```
$ aws s3 cp /tmp/pxf_s3_multi.txt s3://BUCKET/pxf_examples/
```

4. Use the `s3:text:multi` profile to create an external table that references the `pxf_s3_multi.txt` S3 file, making sure to identify the `:` (colon) as the field separator. For example, if your server name is `s3srvcfg`:

```
postgres=# CREATE EXTERNAL TABLE pxf_s3_textmulti(address text, month text, yea
r int)
            LOCATION ('pxf://BUCKET/pxf_examples/pxf_s3_multi.txt?PROFILE=s3:te
xt:multi&SERVER=s3srvcfg')
          FORMAT 'CSV' (delimiter ':');
```

Notice the alternate syntax for specifying the `delimiter`.

5. Query the `pxf_s3_textmulti` table:

```
postgres=# SELECT * FROM pxf_s3_textmulti;
```

```
         address          | month | year
--------------------------+-------+------
 4627 Star Rd.            | Sept  | 2017
 San Francisco, CA  94107
 113 Moon St.             | Jan   | 2018
 San Diego, CA  92093
 51 Belt Ct.              | Dec   | 2016
```

```
 Denver, CO  90123
 93114 Radial Rd.         | Jul    | 2017
 Chicago, IL  60605
 7301 Brookview Ave.      | Dec    | 2018
 Columbus, OH  43213
(5 rows)
```

Writing Text Data

The `<objstore>:text|csv` profiles support writing single line plain text data to an object store. When you create a writable external table with PXF, you specify the name of a directory. When you insert records into a writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specified.

Note: External tables that you create with a writable profile can only be used for `INSERT` operations. If you want to query the data that you inserted, you must create a separate readable external table that references the directory.

Use the following syntax to create a Greenplum Database writable external table that references an object store directory:

```
CREATE WRITABLE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-dir>
    ?PROFILE=<objstore>:text|csv&SERVER=<server_name>[&<custom-option>=<value>[...]]')
FORMAT '[TEXT|CSV]' (delimiter[=|<space>][E]'<delim_value>');
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-dir>	The path to the directory in the data store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-dir> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-dir> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:text PROFILE=<objstore>:csv	The `PROFILE` keyword must identify the specific object store. For example, `s3:text`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	<custom-option>s are described below.
FORMAT	Use `FORMAT 'TEXT'` to write plain, delimited text to <path-to-dir>. Use `FORMAT 'CSV'` to write comma-separated value text to <path-to-dir>.
delimiter	The delimiter character in the data. For `FORMAT 'CSV'`, the default <delim_value> is a comma `,`. Preface the <delim_value> with an `E` when the value is an escape sequence. Examples: `(delimiter=E'\t')`, `(delimiter ':')`.

Keyword	Value
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

Writable external tables that you create using an `<objstore>:text|csv` profile can optionally use record or block compression. You specify the compression type and codec via custom options in the `CREATE EXTERNAL TABLE LOCATION` clause. The `<objstore>:text|csv` profiles support the following custom write options:

Option	Value Description
COMPRESSION_CODEC	The compression codec alias. Supported compression codecs for writing text data include: `default`, `bzip2`, `gzip`, and `uncompressed`. If this option is not provided, Greenplum Database performs no data compression.
COMPRESSION_TYPE	The compression type to employ; supported values are `RECORD` (the default) or `BLOCK`.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example: Writing Text Data to S3

This example utilizes the data schema introduced in Example: Reading Text Data from S3.

Column Name	Data Type
location	text
month	text
number_of_orders	int
total_sales	float8

This example also optionally uses the Greenplum Database external table named `pxf_s3_textsimple` that you created in that exercise.

Procedure

Perform the following procedure to create Greenplum Database writable external tables utilizing the same data schema as described above, one of which will employ compression. You will use the PXF `s3:text` profile to write data to S3. You will also create a separate, readable external table to read the data that you wrote to S3.

1. Create a Greenplum Database writable external table utilizing the data schema described above. Write to the S3 directory `BUCKET/pxf_examples/pxfwrite_s3_textsimple1`. Create the table specifying a comma `,` as the delimiter. For example, if your server name is `s3srvcfg`:

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_s3_writetbl_1(location text, mont
h text, num_orders int, total_sales float8)
            LOCATION ('pxf://BUCKET/pxf_examples/pxfwrite_s3_textsimple1?PROFIL
E=s3:text|csv&SERVER=s3srvcfg')
          FORMAT 'TEXT' (delimiter=',');
```

You specify the `FORMAT` subclause `delimiter` value as the single ascii comma character `,`.

2. Write a few individual records to the `pxfwrite_s3_textsimple1` S3 directory by invoking the SQL `INSERT` command on `pxf_s3_writetbl_1`:

```
postgres=# INSERT INTO pxf_s3_writetbl_1 VALUES ( 'Frankfurt', 'Mar', 777, 3956
.98 );
postgres=# INSERT INTO pxf_s3_writetbl_1 VALUES ( 'Cleveland', 'Oct', 3812, 966
45.37 );
```

3. (Optional) Insert the data from the `pxf_s3_textsimple` table that you created in Example: Reading Text Data from S3 into `pxf_s3_writetbl_1`:

```
postgres=# INSERT INTO pxf_s3_writetbl_1 SELECT * FROM pxf_s3_textsimple;
```

4. Greenplum Database does not support directly querying a writable external table. To query the data that you just added to S3, you must create a readable external Greenplum Database table that references the S3 directory:

```
postgres=# CREATE EXTERNAL TABLE pxf_s3_textsimple_r1(location text, month text
, num_orders int, total_sales float8)
            LOCATION ('pxf://BUCKET/pxf_examples/pxfwrite_s3_textsimple1?PROFIL
E=s3:text&SERVER=s3srvcfg')
            FORMAT 'CSV';
```

You specify the `'CSV'` `FORMAT` when you create the readable external table because you created the writable table with a comma `,` as the delimiter character, the default delimiter for `'CSV'` `FORMAT`.

5. Query the readable external table:

```
postgres=# SELECT * FROM pxf_s3_textsimple_r1 ORDER BY total_sales;
```

```
 location  | month | num_orders | total_sales
-----------+-------+------------+-------------
 Rome      | Mar   |         87 |     1557.39
 Frankfurt | Mar   |        777 |     3956.98
 Prague    | Jan   |        101 |     4875.33
 Bangalore | May   |        317 |     8936.99
 Beijing   | Jul   |        411 |    11600.67
 Cleveland | Oct   |       3812 |    96645.37
(6 rows)
```

The `pxf_s3_textsimple_r1` table includes the records you individually inserted, as well as the full contents of the `pxf_s3_textsimple` table if you performed the optional step.

6. Create a second Greenplum Database writable external table, this time using Gzip compression and employing a colon `:` as the delimiter:

```
postgres=# CREATE WRITABLE EXTERNAL TABLE pxf_s3_writetbl_2 (location text, mon
th text, num_orders int, total_sales float8)
            LOCATION ('pxf://BUCKET/pxf_examples/pxfwrite_s3_textsimple2?PROFIL
E=s3:text&SERVER=s3srvcfg&COMPRESSION_CODEC=gzip')
          FORMAT 'TEXT' (delimiter=':');
```

7. Write a few records to the `pxfwrite_s3_textsimple2` S3 directory by inserting directly into the `pxf_s3_writetbl_2` table:

```
gpadmin=# INSERT INTO pxf_s3_writetbl_2 VALUES ( 'Frankfurt', 'Mar', 777, 3956.
98 );
gpadmin=# INSERT INTO pxf_s3_writetbl_2 VALUES ( 'Cleveland', 'Oct', 3812, 9664
5.37 );
```

8. To query data from the newly-created S3 directory named `pxfwrite_s3_textsimple2`, you can create a readable external Greenplum Database table as described above that references this S3 directory and specifies `FORMAT 'CSV' (delimiter=':')`.

Reading and Writing Avro Data in an Object Store

The PXF object store connectors support reading Avro-format data. This section describes how to use PXF to read and write Avro data in an object store, including how to create, query, and insert into an external table that references an Avro file in the store.

Note: Accessing Avro-format data from an object store is very similar to accessing Avro-format data in HDFS. This topic identifies object store-specific information required to read Avro data, and links to the PXF HDFS Avro documentation where appropriate for common information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from an object store.

Working with Avro Data

Refer to Working with Avro Data in the PXF HDFS Avro documentation for a description of the Apache Avro data serialization framework.

When you read or write Avro data in an object store:

- If the Avro schema file resides in the object store:
  - You must include the bucket in the schema file path. This bucket need not specify the same bucket as the Avro data file.
  - The secrets that you specify in the `SERVER` configuration must provide access to both the data file and schema file buckets.
- The schema file path must not include spaces.

Creating the External Table

Use the `<objstore>:avro` profiles to read and write Avro-format files in an object store. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl

Object Store	Profile Prefix
Google Cloud Storage	gs
Minio	s3
S3	s3

The following syntax creates a Greenplum Database external table that references an Avro-format file:

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>?PROFILE=<objstore>:avro&SERVER=<server_name>[&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-file>	The path to the directory or file in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:avro	The `PROFILE` keyword must identify the specific object store. For example, `s3:avro`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	Avro-specific custom options are described in the PXF HDFS Avro documentation.
FORMAT 'CUSTOM'	Use `FORMAT 'CUSTOM'` with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example

Refer to Example: Reading Avro Data in the PXF HDFS Avro documentation for an Avro example. Modifications that you must make to run the example with an object store include:

- Copying the file to the object store instead of HDFS. For example, to copy the file to S3:

```
$ aws s3 cp /tmp/pxf_avro.avro s3://BUCKET/pxf_examples/
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE pxf_s3_avro(id bigint, username text, followers text[], fmap text, relationship text, address text)
  LOCATION ('pxf://BUCKET/pxf_examples/pxf_avro.avro?PROFILE=s3:avro&SERVER=s3s
```

```
rvcfg&COLLECTION_DELIM=,&MAPKEY_DELIM=:&RECORDKEY_DELIM=:')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

You make similar modifications to follow the steps in Example: Writing Avro Data.

Reading JSON Data from an Object Store

The PXF object store connectors support reading JSON-format data. This section describes how to use PXF to access JSON data in an object store, including how to create and query an external table that references a JSON file in the store.

Note: Accessing JSON-format data from an object store is very similar to accessing JSON-format data in HDFS. This topic identifies object store-specific information required to read JSON data, and links to the PXF HDFS JSON documentation where appropriate for common information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from an object store.

Working with JSON Data

Refer to Working with JSON Data in the PXF HDFS JSON documentation for a description of the JSON text-based data-interchange format.

Creating the External Table

Use the `<objstore>:json` profile to read JSON-format files from an object store. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

The following syntax creates a Greenplum Database readable external table that references a JSON-format file:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>?PROFILE=<objstore>:json&SERVER=<server_name>[&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-file>	The path to the directory or file in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:json	The `PROFILE` keyword must identify the specific object store. For example, `s3:json`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	JSON supports the custom option named `IDENTIFIER` as described in the PXF HDFS JSON documentation.
FORMAT 'CUSTOM'	Use `FORMAT 'CUSTOM'` with the `<objstore>:json` profile. The `CUSTOM FORMAT` requires that you specify `(FORMATTER='pxfwritable_import')`.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example

Refer to Loading the Sample JSON Data to HDFS and Example: Reading a JSON File with Single Line Records in the PXF HDFS JSON documentation for a JSON example. Modifications that you must make to run the example with an object store include:

- Copying the file to the object store instead of HDFS. For example, to copy the file to S3:

```
$ aws s3 cp /tmp/singleline.json s3://BUCKET/pxf_examples/
$ aws s3 cp /tmp/multiline.json s3://BUCKET/pxf_examples/
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE singleline_json_s3(
  created_at TEXT,
  id_str TEXT,
  "user.id" INTEGER,
  "user.location" TEXT,
  "coordinates.values" TEXT[]
)
  LOCATION('pxf://BUCKET/pxf_examples/singleline.json?PROFILE=s3:json&SERVER=s3
srvcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

- If you want to access specific elements of the `coordinates.values` array, you can specify the array subscript number in square brackets:

```
SELECT "coordinates.values"[1], "coordinates.values"[2] FROM singleline_json_s3
;
```

Reading ORC Data from an Object Store

The PXF object store connectors support reading ORC-format data. This section describes how to

use PXF to access ORC data in an object store, including how to create and query an external table that references a file in the store.

Note: Accessing ORC-format data from an object store is very similar to accessing ORC-format data in HDFS. This topic identifies object store-specific information required to read ORC data, and links to the PXF Hadoop ORC documentation where appropriate for common information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from an object store.

Data Type Mapping

Refer to Data Type Mapping in the PXF Hadoop ORC documentation for a description of the mapping between Greenplum Database and ORC data types.

Creating the External Table

Use the `<objstore>:orc` profile to read ORC-format files from an object store. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

The following syntax creates a Greenplum Database readable external table that references an ORC-format file:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>?PROFILE=<objstore>:orc&SERVER=<server_name>[&<custom-o
ption>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-file>	The path to the directory or file in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-file> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-file> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:orc	The `PROFILE` keyword must identify the specific object store. For example, `s3:orc`.

Keyword	Value
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	ORC supports customs options as described in the PXF Hadoop ORC documentation.
FORMAT 'CUSTOM'	Use `FORMAT 'CUSTOM'` with the `<objstore>:orc` profile. The `CUSTOM FORMAT` requires that you specify `(FORMATTER='pxfwritable_import')`.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example

Refer to Example: Reading an ORC File on HDFS in the PXF Hadoop ORC documentation for an example. Modifications that you must make to run the example with an object store include:

- Copying the file to the object store instead of HDFS. For example, to copy the file to S3:

```
$ aws s3 cp /tmp/sampledata.orc s3://BUCKET/pxf_examples/
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE sample_orc( location TEXT, month TEXT, num_orders INTEGER
, total_sales NUMERIC(10,2), items_sold TEXT[] )
  LOCATION('pxf://BUCKET/pxf_examples/sampledata.orc?PROFILE=s3:orc&SERVER=s3sr
vcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Reading and Writing Parquet Data in an Object Store

The PXF object store connectors support reading and writing Parquet-format data. This section describes how to use PXF to access Parquet-format data in an object store, including how to create and query external tables that references a Parquet file in the store.

Note: Accessing Parquet-format data from an object store is very similar to accessing Parquet-format data in HDFS. This topic identifies object store-specific information required to read and write Parquet data, and links to the PXF HDFS Parquet documentation where appropriate for common information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from or write data to an object store.

Data Type Mapping

Refer to Data Type Mapping in the PXF HDFS Parquet documentation for a description of the mapping between Greenplum Database and Parquet data types.

Creating the External Table

The PXF `<objstore>:parquet` profiles support reading and writing data in Parquet-format. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

Use the following syntax to create a Greenplum Database external table that references an HDFS directory. When you insert records into a writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specified.

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-dir>
    ?PROFILE=<objstore>:parquet&SERVER=<server_name>[&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export');
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-dir>	The path to the directory in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-dir> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-dir> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:parquet	The `PROFILE` keyword must identify the specific object store. For example, `s3:parquet`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	Parquet-specific custom options are described in the PXF HDFS Parquet documentation.
FORMAT 'CUSTOM'	Use `FORMAT` '`CUSTOM`' with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

If you are accessing an S3 object store:

- You can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.
- If you are reading Parquet data from S3, you can direct PXF to use the S3 Select Amazon

service to retrieve the data. Refer to Using the Amazon S3 Select Service for more information about the PXF custom option used for this purpose.

Example

Refer to the Example in the PXF HDFS Parquet documentation for a Parquet write/read example. Modifications that you must make to run the example with an object store include:

- Using the `CREATE WRITABLE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above for the writable external table. For example, if your server name is `s3srvcfg`:

```
CREATE WRITABLE EXTERNAL TABLE pxf_tbl_parquet_s3 (location text, month text, n
umber_of_orders int, total_sales double precision)
  LOCATION ('pxf://BUCKET/pxf_examples/pxf_parquet?PROFILE=s3:parquet&SERVER=s3
srvcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above for the readable external table. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE read_pxf_parquet_s3(location text, month text, number_of_
orders int, total_sales double precision)
  LOCATION ('pxf://BUCKET/pxf_examples/pxf_parquet?PROFILE=s3:parquet&SERVER=s3
srvcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Reading and Writing SequenceFile Data in an Object Store

The PXF object store connectors support SequenceFile format binary data. This section describes how to use PXF to read and write SequenceFile data, including how to create, insert, and query data in external tables that reference files in an object store.

Note: Accessing SequenceFile-format data from an object store is very similar to accessing SequenceFile-format data in HDFS. This topic identifies object store-specific information required to read and write SequenceFile data, and links to the PXF HDFS SequenceFile documentation where appropriate for common information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from or write data to an object store.

Creating the External Table

The PXF `<objstore>:SequenceFile` profiles support reading and writing binary data in SequenceFile-format. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs

Object Store	Profile Prefix
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

Use the following syntax to create a Greenplum Database external table that references an HDFS directory. When you insert records into a writable external table, the block(s) of data that you insert are written to one or more files in the directory that you specified.

```
CREATE [WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-dir>
    ?PROFILE=<objstore>:SequenceFile&SERVER=<server_name>[&<custom-option>=<value>[...
]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export')
[DISTRIBUTED BY (<column_name> [, ... ] ) | DISTRIBUTED RANDOMLY];
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-dir>	The path to the directory in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-dir> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-dir> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:SequenceFile	The `PROFILE` keyword must identify the specific object store. For example, `s3:SequenceFile`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
<custom-option>=<value>	SequenceFile-specific custom options are described in the PXF HDFS SequenceFile documentation.
FORMAT 'CUSTOM'	Use `FORMAT` '`CUSTOM`' with `(FORMATTER='pxfwritable_export')` (write) or `(FORMATTER='pxfwritable_import')` (read).
DISTRIBUTED BY	If you want to load data from an existing Greenplum Database table into the writable external table, consider specifying the same distribution policy or <column_name> on both tables. Doing so will avoid extra motion of data between segments on the load operation.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example

Refer to Example: Writing Binary Data to HDFS in the PXF HDFS SequenceFile documentation for a write/read example. Modifications that you must make to run the example with an object store include:

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described

above for the writable external table. For example, if your server name is `s3srvcfg`:

```
CREATE WRITABLE EXTERNAL TABLE pxf_tbl_seqfile_s3(location text, month text, nu
mber_of_orders integer, total_sales real)
  LOCATION ('pxf://BUCKET/pxf_examples/pxf_seqfile?PROFILE=s3:SequenceFile&DATA
-SCHEMA=com.example.pxf.hdfs.writable.dataschema.PxfExample_CustomWritable&COMP
RESSION_TYPE=BLOCK&COMPRESSION_CODEC=bzip2&SERVER=s3srvcfg')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above for the readable external table. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE read_pxf_tbl_seqfile_s3(location text, month text, number
_of_orders integer, total_sales real)
  LOCATION ('pxf://BUCKET/pxf_examples/pxf_seqfile?PROFILE=s3:SequenceFile&DATA
-SCHEMA=com.example.pxf.hdfs.writable.dataschema.PxfExample_CustomWritable&SERV
ER=s3srvcfg')
 FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

Reading a Multi-Line Text File into a Single Table Row

The PXF object store connectors support reading a multi-line text file as a single table row. This section describes how to use PXF to read multi-line text and JSON data files in an object store, including how to create an external table that references multiple files in the store.

PXF supports reading only text and JSON files in this manner.

Note: Accessing multi-line files from an object store is very similar to accessing multi-line files in HDFS. This topic identifies the object store-specific information required to read these files. Refer to the PXF HDFS documentation for more information.

Prerequisites

Ensure that you have met the PXF Object Store Prerequisites before you attempt to read data from multiple files residing in an object store.

Creating the External Table

Use the `<objstore>:text:multi` profile to read multiple files in an object store each into a single table row. PXF supports the following `<objstore>` profile prefixes:

Object Store	Profile Prefix
Azure Blob Storage	wasbs
Azure Data Lake	adl
Google Cloud Storage	gs
Minio	s3
S3	s3

The following syntax creates a Greenplum Database readable external table that references one or more text files in an object store:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> text|json | LIKE <other_table> )
  LOCATION ('pxf://<path-to-files>?PROFILE=<objstore>:text:multi&SERVER=<server_name>[
&IGNORE_MISSING_PATH=<boolean>]&FILE_AS_ROW=true')
  FORMAT 'CSV');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<path-to-files>	The path to the directory or files in the object store. When the `<server_name>` configuration includes a `pxf.fs.basePath` property setting, PXF considers <path-to-files> to be relative to the base path specified. Otherwise, PXF considers it to be an absolute path. <path-to-files> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE=<objstore>:text:multi	The `PROFILE` keyword must identify the specific object store. For example, `s3:text:multi`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data.
IGNORE_MISSING_PATH=<boolean>	Specify the action to take when <path-to-files> is missing or invalid. The default value is `false`, PXF returns an error in this situation. When the value is `true`, PXF ignores missing path errors and returns an empty fragment.
FILE_AS_ROW=true	The required option that instructs PXF to read each file into a single table row.
FORMAT	The `FORMAT` must specify `'CSV'`.

If you are accessing an S3 object store, you can provide S3 credentials via custom options in the `CREATE EXTERNAL TABLE` command as described in Overriding the S3 Server Configuration with DDL.

Example

Refer to Example: Reading an HDFS Text File into a Single Table Row in the PXF HDFS documentation for an example. Modifications that you must make to run the example with an object store include:

- Copying the file to the object store instead of HDFS. For example, to copy the file to S3:

```
$ aws s3 cp /tmp/file1.txt s3://BUCKET/pxf_examples/tdir
$ aws s3 cp /tmp/file2.txt s3://BUCKET/pxf_examples/tdir
$ aws s3 cp /tmp/file3.txt s3://BUCKET/pxf_examples/tdir
```

- Using the `CREATE EXTERNAL TABLE` syntax and `LOCATION` keywords and settings described above. For example, if your server name is `s3srvcfg`:

```
CREATE EXTERNAL TABLE pxf_readfileasrow_s3( c1 text )
  LOCATION('pxf://BUCKET/pxf_examples/tdir?PROFILE=s3:text:multi&SERVER=s3srvcf
g&FILE_AS_ROW=true')
FORMAT 'CSV'
```

Reading CSV and Parquet Data from S3 Using S3 Select

The PXF S3 connector supports reading certain CSV- and Parquet-format data from S3 using the Amazon S3 Select service. S3 Select provides direct query-in-place features on data stored in Amazon S3.

When you enable it, PXF uses S3 Select to filter the contents of S3 objects to retrieve the subset of data that you request. This typically reduces both the amount of data transferred to Greenplum Database and the query time.

You can use the PXF S3 Connector with S3 Select to read:

- `gzip`- or `bzip2`-compressed `CSV` files
- `Parquet` files with `gzip`- or `snappy`-compressed columns

The data must be `UTF-8`-encoded, and may be server-side encrypted.

PXF supports column projection as well as predicate pushdown for `AND`, `OR`, and `NOT` operators when using S3 Select.

Using the Amazon S3 Select service may increase the cost of data access and retrieval. Be sure to consider the associated costs before you enable PXF to use the S3 Select service.

Enabling PXF to Use S3 Select

The `S3_SELECT` external table custom option governs PXF's use of S3 Select when accessing the S3 object store. You can provide the following values when you set the `S3_SELECT` option:

S3-SELECT Value	Description
OFF	PXF does not use S3 Select; the default.
ON	PXF always uses S3 Select.
AUTO	PXF uses S3 Select when it will benefit access or performance.

By default, PXF does not use S3 Select (`S3_SELECT=OFF`). You can enable PXF to always use S3 Select, or to use S3 Select only when PXF determines that it could be beneficial for performance. For example, when `S3_SELECT=AUTO`, PXF automatically uses S3 Select when a query on the external table utilizes column projection or predicate pushdown, or when the referenced CSV file has a header row.

Note: The `IGNORE_MISSING_PATH` custom option is not available when you use a PXF external table to read CSV text and Parquet data from S3 using S3 Select.

Reading Parquet Data with S3 Select

PXF supports reading Parquet data from S3 as described in Reading and Writing Parquet Data in an Object Store. If you want PXF to use S3 Select when reading the Parquet data, you add the `S3_SELECT` custom option and value to the `CREATE EXTERNAL TABLE LOCATION` URI.

Specifying the Parquet Column Compression Type

If columns in the Parquet file are `gzip`- or `snappy`-compressed, use the `COMPRESSION_CODEC` custom option in the `LOCATION` URI to identify the compression codec alias. For example:

```
&COMPRESSION_CODEC=gzip
```

Or,

```
&COMPRESSION_CODEC=snappy
```

Creating the External Table

Use the following syntax to create a Greenplum Database external table that references a Parquet file on S3 that you want PXF to access with the S3 Select service:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
  LOCATION ('pxf://<path-to-file>?PROFILE=s3:parquet&SERVER=<server_name>&S3_SELECT=ON
|AUTO[&<other-custom-option>=<value>[...]]')
FORMAT 'CSV';
```

You *must* specify `FORMAT 'CSV'` when you enable PXF to use S3 Select on an external table that accesses a Parquet file on S3.

For example, use the following command to have PXF use S3 Select to access a Parquet file on S3 when optimal:

```
CREATE EXTERNAL TABLE parquet_on_s3 ( LIKE table1 )
  LOCATION ('pxf://bucket/file.parquet?PROFILE=s3:parquet&SERVER=s3srvcfg&S3_SELECT=AU
TO')
FORMAT 'CSV';
```

Reading CSV files with S3 Select

PXF supports reading CSV data from S3 as described in Reading and Writing Text Data in an Object Store. If you want PXF to use S3 Select when reading the CSV data, you add the `S3_SELECT` custom option and value to the `CREATE EXTERNAL TABLE LOCATION` URI. You may also specify the delimiter formatter option and the file header and compression custom options.

Handling the CSV File Header

PXF can read a CSV file with a header row *only* when the S3 Connector uses the Amazon S3 Select service to access the file on S3. PXF does not support reading a CSV file that includes a header row from any other external data store.

CSV files may include a header line. When you enable PXF to use S3 Select to access a CSV-format file, you use the `FILE_HEADER` custom option in the `LOCATION` URI to identify whether or not the CSV file has a header row and, if so, how you want PXF to handle the header. PXF never returns the header row.

Note: You *must* specify `S3_SELECT=ON` or `S3_SELECT=AUTO` when the CSV file has a header row. Do not specify `S3_SELECT=OFF` in this case.

The `FILE_HEADER` option takes the following values:

FILE_HEADER Value	Description

NONE	The file has no header row; the default.
IGNORE	The file has a header row; ignore the header. Use when the order of the columns in the external table and the CSV file are the same. (When the column order is the same, the column names and the CSV header names may be different.)
USE	The file has a header row; read the header. Use when the external table column names and the CSV header names are the same, but are in a different order.

If both the order and the names of the external table columns and the CSV header are the same, you can specify either `FILE_HEADER=IGNORE` or `FILE_HEADER=USE`.

PXF cannot match the CSV data with the external table definition when both the order and the names of the external table columns are different from the CSV header columns. Any query on an external table with these conditions fails with the error `Some headers in the query are missing from the file`.

For example, if the order of the columns in the CSV file header and the external table are the same, add the following to the `CREATE EXTERNAL TABLE LOCATION` URI to have PXF ignore the CSV header:

```
&FILE_HEADER=IGNORE
```

Specifying the CSV File Compression Type

If the CSV file is `gzip`- or `bzip2`-compressed, use the `COMPRESSION_CODEC` custom option in the `LOCATION` URI to identify the compression codec alias. For example:

```
&COMPRESSION_CODEC=gzip
```

Or,

```
&COMPRESSION_CODEC=bzip2
```

Creating the External Table

Use the following syntax to create a Greenplum Database external table that references a CSV file on S3 that you want PXF to access with the S3 Select service:

```
CREATE EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<path-to-file>
    ?PROFILE=s3:text&SERVER=<server_name>&S3_SELECT=ON|AUTO[&FILE_HEADER=IGNORE|USE][&
COMPRESSION_CODEC=gzip|bzip2][&<other-custom-option>=<value>[...]]')
FORMAT 'CSV' [(delimiter '<delim_char>')];
```

Note: Do not use the `(HEADER)` formatter option in the `CREATE EXTERNAL TABLE` command.

Note: PXF does not support the `SKIP_HEADER_COUNT` custom option when you read a CSV file on S3 using the S3 Select service.

For example, use the following command to have PXF always use S3 Select to access a `gzip`-compressed file on S3, where the field delimiter is a pipe ('|') character and the external table and CSV header columns are in the same order.

```
CREATE EXTERNAL TABLE gzippedcsv_on_s3 ( LIKE table2 )
  LOCATION ('pxf://bucket/file.csv.gz?PROFILE=s3:text&SERVER=s3srvcfg&S3_SELECT=ON&FIL
E_HEADER=USE')
FORMAT 'CSV' (delimiter '|');
```

Accessing an SQL Database with PXF (JDBC)

Some of your data may already reside in an external SQL database. PXF provides access to this data via the PXF JDBC connector. The JDBC connector is a JDBC client. It can read data from and write data to SQL databases including MySQL, ORACLE, Microsoft SQL Server, DB2, PostgreSQL, Hive, and Apache Ignite.

This section describes how to use the PXF JDBC connector to access data in an external SQL database, including how to create and query or insert data into a PXF external table that references a table in an external database.

The JDBC connector does not guarantee consistency when writing to an external SQL database. Be aware that if an `INSERT` operation fails, some data may be written to the external database table. If you require consistency for writes, consider writing to a staging table in the external database, and loading to the target table only after verifying the write operation.

Prerequisites

Before you access an external SQL database using the PXF JDBC connector, ensure that:

- You can identify the PXF runtime configuration directory (`$PXF_BASE`).
- You have configured PXF, and PXF is running on each Greenplum Database host. See Configuring PXF for additional information.
- Connectivity exists between all Greenplum Database hosts and the external SQL database.
- You have configured your external SQL database for user access from all Greenplum Database hosts.
- You have registered any JDBC driver JAR dependencies.
- (Recommended) You have created one or more named PXF JDBC connector server configurations as described in Configuring the PXF JDBC Connector.

Data Types Supported

The PXF JDBC connector supports the following data types:

- INTEGER, BIGINT, SMALLINT
- REAL, FLOAT8
- NUMERIC
- BOOLEAN
- VARCHAR, BPCHAR, TEXT

- DATE
- TIMESTAMP
- BYTEA

Any data type not listed above is not supported by the PXF JDBC connector.

About Accessing Hive via JDBC

PXF includes version 1.1.0 of the Hive JDBC driver. This version does **not** support the following data types when you use the PXF JDBC connector to operate on a Hive table:

Data Type	Fixed in Hive JDBC Driver	Upstream Issue	Operations Not Supported
NUMERIC	2.3.0	HIVE-13614	Write
TIMESTAMP	2.0.0	HIVE-11748	Write
DATE	1.3.0, 2.0.0	HIVE-11024	Write
BYTEA	N/A	N/A	Read, Write

Accessing an External SQL Database

The PXF JDBC connector supports a single profile named `jdbc`. You can both read data from and write data to an external SQL database table with this profile. You can also use the connector to run a static, named query in external SQL database and read the results.

To access data in a remote SQL database, you create a readable or writable Greenplum Database external table that references the remote database table. The Greenplum Database external table and the remote database table or query result tuple must have the same definition; the column names and types must match.

Use the following syntax to create a Greenplum Database external table that references a remote SQL database table or a query result from the remote database:

```
CREATE [READABLE | WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<external-table-name>|query:<query_name>?PROFILE=jdbc[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import'|'pxfwritable_export');
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<external-table-name>	The full name of the external table. Depends on the external SQL database, may include a schema name and a table name.
query:<query_name>	The name of the query to execute in the remote SQL database.
PROFILE	The `PROFILE` keyword value must specify `jdbc`.
SERVER=<server_name>	The named server configuration that PXF uses to access the data. PXF uses the `default` server if not specified.

Keyword	Value
<custom-option>=<value>	<custom-option> is profile-specific. `jdbc` profile-specific options are discussed in the next section.
FORMAT 'CUSTOM'	The JDBC `CUSTOM FORMAT` supports the built-in `'pxfwritable_import' FORMATTER` function for read operations and the built-in `'pxfwritable_export'` function for write operations.

Note: You cannot use the `HEADER` option in your `FORMAT` specification when you create a PXF external table.

JDBC Custom Options

You include JDBC connector custom options in the `LOCATION` URI, prefacing each option with an ampersand `&`. `CREATE EXTERNAL TABLE` <custom-option>s supported by the `jdbc` profile include:

Option Name	Operation	Description
BATCH_SIZE	Write	Integer that identifies the number of `INSERT` operations to batch to the external SQL database. Write batching is enabled by default; the default value is 100.
FETCH_SIZE	Read	Integer that identifies the number of rows to buffer when reading from an external SQL database. Read row batching is enabled by default. The default read fetch size for MySQL is `-2147483648` (`Integer.MIN_VALUE`). The default read fetch size for all other databases is 1000.
QUERY_TIMEOUT	Read/Write	Integer that identifies the amount of time (in seconds) that the JDBC driver waits for a statement to execute. The default wait time is infinite.
POOL_SIZE	Write	Enable thread pooling on `INSERT` operations and identify the number of threads in the pool. Thread pooling is disabled by default.
PARTITION_BY	Read	Enables read partitioning. The partition column, <column-name>:<column-type>. You may specify only one partition column. The JDBC connector supports `date`, `int`, and `enum` <column-type> values, where `int` represents any JDBC integral type. If you do not identify a `PARTITION_BY` column, a single PXF instance services the read request.
RANGE	Read	Required when `PARTITION_BY` is specified. The query range; used as a hint to aid the creation of partitions. The `RANGE` format is dependent upon the data type of the partition column. When the partition column is an `enum` type, `RANGE` must specify a list of values, <value>:<value>[:<value>[…]], each of which forms its own fragment. If the partition column is an `int` or `date` type, `RANGE` must specify <start-value>:<end-value> and represents the interval from <start-value> through <end-value>, inclusive. The `RANGE` for an `int` partition column may span any 64-bit signed integer values. If the partition column is a `date` type, use the `yyyy-MM-dd` date format.
INTERVAL	Read	Required when `PARTITION_BY` is specified and of the `int`, `bigint`, or `date` type. The interval, <interval-value>[:<interval-unit>], of one fragment. Used with `RANGE` as a hint to aid the creation of partitions. Specify the size of the fragment in <interval-value>. If the partition column is a `date` type, use the <interval-unit> to specify `year`, `month`, or `day`. PXF ignores `INTERVAL` when the `PARTITION_BY` column is of the `enum` type.

Option Name	Operation	Description
QUOTE_COLUMNS	Read	Controls whether PXF should quote column names when constructing an SQL query to the external database. Specify `true` to force PXF to quote all column names; PXF does not quote column names if any other value is provided. If `QUOTE_COLUMNS` is not specified (the default), PXF automatically quotes *all* column names in the query when *any* column name: - includes special characters, or - is mixed case and the external database does not support unquoted mixed case identifiers.

Batching Insert Operations (Write)

When the JDBC driver of the external SQL database supports it, batching of `INSERT` operations may significantly increase performance.

Write batching is enabled by default, and the default batch size is 100. To disable batching or to modify the default batch size value, create the PXF external table with a `BATCH_SIZE` setting:

- `BATCH_SIZE=0` or `BATCH_SIZE=1` - disables batching
- `BATCH_SIZE=(n>1)` - sets the `BATCH_SIZE` to `n`

When the external database JDBC driver does not support batching, the behaviour of the PXF JDBC connector depends on the `BATCH_SIZE` setting as follows:

- `BATCH_SIZE` omitted - The JDBC connector inserts without batching.
- `BATCH_SIZE=(n>1)` - The `INSERT` operation fails and the connector returns an error.

Batching on Read Operations

By default, the PXF JDBC connector automatically batches the rows it fetches from an external database table. The default row fetch size is 1000. To modify the default fetch size value, specify a `FETCH_SIZE` when you create the PXF external table. For example:

```
FETCH_SIZE=5000
```

If the external database JDBC driver does not support batching on read, you must explicitly disable read row batching by setting `FETCH_SIZE=0`.

Thread Pooling (Write)

The PXF JDBC connector can further increase write performance by processing `INSERT` operations in multiple threads when threading is supported by the JDBC driver of the external SQL database.

Consider using batching together with a thread pool. When used together, each thread receives and processes one complete batch of data. If you use a thread pool without batching, each thread in the pool receives exactly one tuple.

The JDBC connector returns an error when any thread in the thread pool fails. Be aware that if an `INSERT` operation fails, some data may be written to the external database table.

To disable or enable a thread pool and set the pool size, create the PXF external table with a `POOL_SIZE` setting as follows:

- `POOL_SIZE=(n<1)` - thread pool size is the number of CPUs in the system
- `POOL_SIZE=1` - disable thread pooling
- `POOL_SIZE=(n>1)` - set the `POOL_SIZE` to `n`

Partitioning (Read)

The PXF JDBC connector supports simultaneous read access from PXF instances running on multiple Greenplum Database hosts to an external SQL table. This feature is referred to as partitioning. Read partitioning is not enabled by default. To enable read partitioning, set the `PARTITION_BY`, `RANGE`, and `INTERVAL` custom options when you create the PXF external table.

PXF uses the `RANGE` and `INTERVAL` values and the `PARTITON_BY` column that you specify to assign specific data rows in the external table to PXF instances running on the Greenplum Database segment hosts. This column selection is specific to PXF processing, and has no relationship to a partition column that you may have specified for the table in the external SQL database.

Example JDBC <custom-option> substrings that identify partitioning parameters:

```
&PARTITION_BY=id:int&RANGE=1:100&INTERVAL=5
&PARTITION_BY=year:int&RANGE=2011:2013&INTERVAL=1
&PARTITION_BY=createdate:date&RANGE=2013-01-01:2016-01-01&INTERVAL=1:month
&PARTITION_BY=color:enum&RANGE=red:yellow:blue
```

When you enable partitioning, the PXF JDBC connector splits a `SELECT` query into multiple subqueries that retrieve a subset of the data, each of which is called a fragment. The JDBC connector automatically adds extra query constraints (`WHERE` expressions) to each fragment to guarantee that every tuple of data is retrieved from the external database exactly once.

For example, when a user queries a PXF external table created with a `LOCATION` clause that specifies `&PARTITION_BY=id:int&RANGE=1:5&INTERVAL=2`, PXF generates 5 fragments: two according to the partition settings and up to three implicitly generated fragments. The constraints associated with each fragment are as follows:

- Fragment 1: WHERE (id < 1) - implicitly-generated fragment for RANGE start-bounded interval
- Fragment 2: WHERE (id >= 1) AND (id < 3) - fragment specified by partition settings
- Fragment 3: WHERE (id >= 3) AND (id < 5) - fragment specified by partition settings
- Fragment 4: WHERE (id >= 5) - implicitly-generated fragment for RANGE end-bounded interval
- Fragment 5: WHERE (id IS NULL) - implicitly-generated fragment

PXF distributes the fragments among Greenplum Database segments. A PXF instance running on a segment host spawns a thread for each segment on that host that services a fragment. If the number of fragments is less than or equal to the number of Greenplum segments configured on a segment host, a single PXF instance may service all of the fragments. Each PXF instance sends its results back to Greenplum Database, where they are collected and returned to the user.

When you specify the `PARTITION_BY` option, tune the `INTERVAL` value and unit based upon the optimal number of JDBC connections to the target database and the optimal distribution of external data across Greenplum Database segments. The `INTERVAL` low boundary is driven by the number of

Greenplum Database segments while the high boundary is driven by the acceptable number of JDBC connections to the target database. The `INTERVAL` setting influences the number of fragments, and should ideally not be set too high nor too low. Testing with multiple values may help you select the optimal settings.

Examples

Refer to the following topics for examples on how to use PXF to read data from and write data to specific SQL databases:

- Reading From and Writing to a PostgreSQL Table
- Reading From and Writing to a MySQL Table
- Reading From and Writing to an Oracle Table
- Reading From and Writing to a Trino Table

About Using Named Queries

The PXF JDBC Connector allows you to specify a statically-defined query to run against the remote SQL database. Consider using a *named query* when:

- You need to join several tables that all reside in the same external database.
- You want to perform complex aggregation closer to the data source.
- You would use, but are not allowed to create, a `VIEW` in the external database.
- You would rather consume computational resources in the external system to minimize utilization of Greenplum Database resources.
- You want to run a HIVE query and control resource utilization via YARN.

The Greenplum Database administrator defines a query and provides you with the query name to use when you create the external table. Instead of a table name, you specify `query:<query_name>` in the `CREATE EXTERNAL TABLE LOCATION` clause to instruct the PXF JDBC connector to run the static query named `<query_name>` in the remote SQL database.

PXF supports named queries only with readable external tables. You must create a unique Greenplum Database readable external table for each query that you want to run.

The names and types of the external table columns must exactly match the names, types, and order of the columns return by the query result. If the query returns the results of an aggregation or other function, be sure to use the `AS` qualifier to specify a specific column name.

For example, suppose that you are working with PostgreSQL tables that have the following definitions:

```
CREATE TABLE customers(id int, name text, city text, state text);
CREATE TABLE orders(customer_id int, amount int, month int, year int);
```

And this PostgreSQL query that the administrator named `order_rpt`:

```
SELECT c.name, sum(o.amount) AS total, o.month
  FROM customers c JOIN orders o ON c.id = o.customer_id
  WHERE c.state = 'CO'
```

```
GROUP BY c.name, o.month
```

This query returns tuples of type `(name text, total int, month int)`. If the `order_rpt` query is defined for the PXF JDBC server named `pgserver`, you could create a Greenplum Database external table to read these query results as follows:

```
CREATE EXTERNAL TABLE orderrpt_frompg(name text, total int, month int)
  LOCATION ('pxf://query:order_rpt?PROFILE=jdbc&SERVER=pgserver&PARTITION_BY=month:int
&RANGE=1:13&INTERVAL=3')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

This command references a query named `order_rpt` defined in the `pgserver` server configuration. It also specifies JDBC read partitioning options that provide PXF with the information that it uses to split/partition the query result data across its servers/segments.

For a more detailed example see Example: Using a Named Query with PostgreSQL.

The PXF JDBC connector automatically applies column projection and filter pushdown to external tables that reference named queries.

Overriding the JDBC Server Configuration with DDL

You can override certain properties in a JDBC server configuration for a specific external database table by directly specifying the custom option in the `CREATE EXTERNAL TABLE LOCATION` clause:

Custom Option Name	jdbc-site.xml Property Name
JDBC_DRIVER	jdbc.driver
DB_URL	jdbc.url
USER	jdbc.user
PASS	jdbc.password
BATCH_SIZE	jdbc.statement.batchSize
FETCH_SIZE	jdbc.statement.fetchSize
QUERY_TIMEOUT	jdbc.statement.queryTimeout

Example JDBC connection strings specified via custom options:

```
&JDBC_DRIVER=org.postgresql.Driver&DB_URL=jdbc:postgresql://pgserverhost:5432/pgtestdb
&USER=pguser1&PASS=changeme
&JDBC_DRIVER=com.mysql.jdbc.Driver&DB_URL=jdbc:mysql://mysqlhost:3306/testdb&USER=user
1&PASS=changeme
```

For example:

```
CREATE EXTERNAL TABLE pxf_pgtbl(name text, orders int)
  LOCATION ('pxf://public.forpxf_table1?PROFILE=jdbc&JDBC_DRIVER=org.postgresql.Driver&DB
_URL=jdbc:postgresql://pgserverhost:5432/pgtestdb&USER=pxfuser1&PASS=changeme')
FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

Credentials that you provide in this manner are visible as part of the external table definition. Do not use this method of passing credentials in a production environment.

Refer to Configuration Property Precedence for detailed information about the precedence rules that PXF uses to obtain configuration property settings for a Greenplum Database user.

Example: Reading From and Writing to a PostgreSQL Table

In this example, you:

- Create a PostgreSQL database and table, and insert data into the table
- Create a PostgreSQL user and assign all privileges on the table to the user
- Configure the PXF JDBC connector to access the PostgreSQL database
- Create a PXF readable external table that references the PostgreSQL table
- Read the data in the PostgreSQL table using PXF
- Create a PXF writable external table that references the PostgreSQL table
- Write data to the PostgreSQL table using PXF
- Read the data in the PostgreSQL table again

Create a PostgreSQL Table

Perform the following steps to create a PostgreSQL table named `forpxf_table1` in the `public` schema of a database named `pgtestdb`, and grant a user named `pxfuser1` all privileges on this table:

1. Identify the host name and port of your PostgreSQL server.
2. Connect to the default PostgreSQL database as the `postgres` user. For example, if your PostgreSQL server is running on the default port on the host named `pserver`:

   ```
   $ psql -U postgres -h pserver
   ```

3. Create a PostgreSQL database named `pgtestdb` and connect to this database:

   ```
   =# CREATE DATABASE pgtestdb;
   =# \connect pgtestdb;
   ```

4. Create a table named `forpxf_table1` and insert some data into this table:

   ```
   =# CREATE TABLE forpxf_table1(id int);
   =# INSERT INTO forpxf_table1 VALUES (1);
   =# INSERT INTO forpxf_table1 VALUES (2);
   =# INSERT INTO forpxf_table1 VALUES (3);
   ```

5. Create a PostgreSQL user named `pxfuser1`:

   ```
   =# CREATE USER pxfuser1 WITH PASSWORD 'changeme';
   ```

6. Assign user `pxfuser1` all privileges on table `forpxf_table1`, and exit the `psql` subsystem:

   ```
   =# GRANT ALL ON forpxf_table1 TO pxfuser1;
   =# \q
   ```

   With these privileges, `pxfuser1` can read from and write to the `forpxf_table1` table.

7. Update the PostgreSQL configuration to allow user `pxfuser1` to access `pgtestdb` from each Greenplum Database host. This configuration is specific to your PostgreSQL environment. You will update the `/var/lib/pgsql/pg_hba.conf` file and then restart the PostgreSQL server.

Configure the JDBC Connector

You must create a JDBC server configuration for PostgreSQL and synchronize the PXF configuration. The PostgreSQL JAR file is bundled with PXF, so there is no need to manually download it.

This procedure will typically be performed by the Greenplum Database administrator.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Create a JDBC server configuration for PostgreSQL as described in Example Configuration Procedure, naming the server directory `pgsrvcfg`. The `jdbc-site.xml` file contents should look similar to the following (substitute your PostgreSQL host system for `pgserverhost`):

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
<property>
    <name>jdbc.driver</name>
    <value>org.postgresql.Driver</value>
</property>
<property>
    <name>jdbc.url</name>
    <value>jdbc:postgresql://pgserverhost:5432/pgtestdb</value>
</property>
<property>
    <name>jdbc.user</name>
    <value>pxfuser1</value>
</property>
<property>
    <name>jdbc.password</name>
    <value>changeme</value>
</property>
</configuration>
```

3. Synchronize the PXF server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Read from the PostgreSQL Table

Perform the following procedure to create a PXF external table that references the `forpxf_table1` PostgreSQL table that you created in the previous section, and reads the data in the table:

1. Create the PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE EXTERNAL TABLE pxf_tblfrompg(id int)
            LOCATION ('pxf://public.forpxf_table1?PROFILE=jdbc&SERVER=pgsrvcfg'
```

```
)
            FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
```

2. Display all rows of the `pxf_tblfrompg` table:

```
gpadmin=# SELECT * FROM pxf_tblfrompg;
 id
----
  1
  2
  3
(3 rows)
```

Write to the PostgreSQL Table

Perform the following procedure to insert some data into the `forpxf_table1` Postgres table and then read from the table. You must create a new external table for the write operation.

1. Create a writable PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE WRITABLE EXTERNAL TABLE pxf_writeto_postgres(id int)
            LOCATION ('pxf://public.forpxf_table1?PROFILE=jdbc&SERVER=pgsrvcfg'
)
          FORMAT 'CUSTOM' (FORMATTER='pxfwritable_export');
```

2. Insert some data into the `pxf_writeto_postgres` table. For example:

```
=# INSERT INTO pxf_writeto_postgres VALUES (111);
=# INSERT INTO pxf_writeto_postgres VALUES (222);
=# INSERT INTO pxf_writeto_postgres VALUES (333);
```

3. Use the `pxf_tblfrompg` readable external table that you created in the previous section to view the new data in the `forpxf_table1` PostgreSQL table:

```
gpadmin=# SELECT * FROM pxf_tblfrompg ORDER BY id DESC;
 id
-----
 333
 222
 111
   3
   2
   1
(6 rows)
```

Example: Reading From and Writing to a MySQL Table

In this example, you:

- Create a MySQL database and table, and insert data into the table
- Create a MySQL user and assign all privileges on the table to the user
- Configure the PXF JDBC connector to access the MySQL database
- Create a PXF readable external table that references the MySQL table

- Read the data in the MySQL table using PXF
- Create a PXF writable external table that references the MySQL table
- Write data to the MySQL table using PXF
- Read the data in the MySQL table again

Create a MySQL Table

Perform the following steps to create a MySQL table named `names` in a database named `mysqltestdb`, and grant a user named `mysql-user` all privileges on this table:

1. Identify the host name and port of your MySQL server.
2. Connect to the default MySQL database as the `root` user:

```
$ mysql -u root -p
```

3. Create a MySQL database named `mysqltestdb` and connect to this database:

```
> CREATE DATABASE mysqltestdb;
> USE mysqltestdb;
```

4. Create a table named `names` and insert some data into this table:

```
> CREATE TABLE names (id int, name varchar(64), last varchar(64));
> INSERT INTO names values (1, 'John', 'Smith'), (2, 'Mary', 'Blake');
```

5. Create a MySQL user named `mysql-user` and assign the password `my-secret-pw` to it:

```
> CREATE USER 'mysql-user' IDENTIFIED BY 'my-secret-pw';
```

6. Assign user `mysql-user` all privileges on table `names`, and exit the `mysql` subsystem:

```
> GRANT ALL PRIVILEGES ON mysqltestdb.names TO 'mysql-user';
> exit
```

With these privileges, `mysql-user` can read from and write to the `names` table.

Configure the MySQL Connector

You must create a JDBC server configuration for MySQL, download the MySQL driver JAR file to your system, copy the JAR file to the PXF user configuration directory, synchronize the PXF configuration, and then restart PXF.

This procedure will typically be performed by the Greenplum Database administrator.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Download the MySQL JDBC driver and place it under `$PXF_BASE/lib`. If you relocated $PXF_BASE, make sure you use the updated location. You can download a MySQL JDBC

driver from your preferred download location. The following example downloads the driver from Maven Central and places it under `$PXF_BASE/lib`:

1. If you did not relocate `$PXF_BASE`, run the following from the Greenplum master:

    ```
    gpadmin@gpmaster$ cd /usr/local/pxf-gp<version>/lib
    gpadmin@gpmaster$ wget https://repo1.maven.org/maven2/mysql/mysql-connect
    or-java/8.0.21/mysql-connector-java-8.0.21.jar
    ```

2. If you relocated `$PXF_BASE`, run the following from the Greenplum master:

    ```
    gpadmin@gpmaster$ cd $PXF_BASE/lib
    gpadmin@gpmaster$ wget https://repo1.maven.org/maven2/mysql/mysql-connect
    or-java/8.0.21/mysql-connector-java-8.0.21.jar
    ```

3. Synchronize the PXF configuration, and then restart PXF:

    ```
    gpadmin@gpmaster$ pxf cluster sync
    gpadmin@gpmaster$ pxf cluster restart
    ```

4. Create a JDBC server configuration for MySQL as described in Example Configuration Procedure, naming the server directory `mysql`. The `jdbc-site.xml` file contents should look similar to the following (substitute your MySQL host system for `mysqlserverhost`):

    ```
    <?xml version="1.0" encoding="UTF-8"?>
    <configuration>
        <property>
            <name>jdbc.driver</name>
            <value>com.mysql.jdbc.Driver</value>
            <description>Class name of the JDBC driver</description>
        </property>
        <property>
            <name>jdbc.url</name>
            <value>jdbc:mysql://mysqlserverhost:3306/mysqltestdb</value>
            <description>The URL that the JDBC driver can use to connect to the dat
    abase</description>
        </property>
        <property>
            <name>jdbc.user</name>
            <value>mysql-user</value>
            <description>User name for connecting to the database</description>
        </property>
        <property>
            <name>jdbc.password</name>
            <value>my-secret-pw</value>
            <description>Password for connecting to the database</description>
        </property>
    </configuration>
    ```

5. Synchronize the PXF server configuration to the Greenplum Database cluster:

    ```
    gpadmin@gpmaster$ pxf cluster sync
    ```

Read from the MySQL Table

Perform the following procedure to create a PXF external table that references the `names` MySQL

table that you created in the previous section, and reads the data in the table:

1. Create the PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE EXTERNAL TABLE names_in_mysql (id int, name text, last text)
          LOCATION('pxf://names?PROFILE=jdbc&SERVER=mysql')
          FORMAT 'CUSTOM' (formatter='pxfwritable_import');
```

2. Display all rows of the `names_in_mysql` table:

```
gpadmin=# SELECT * FROM names_in_mysql;
 id |    name     |    last
----+-----------+----------
  1 |    John     |    Smith
  2 |    Mary     |    Blake
(2 rows)
```

Write to the MySQL Table

Perform the following procedure to insert some data into the `names` MySQL table and then read from the table. You must create a new external table for the write operation.

1. Create a writable PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE WRITABLE EXTERNAL TABLE names_in_mysql_w (id int, name text, l
ast text)
          LOCATION('pxf://names?PROFILE=jdbc&SERVER=mysql')
          FORMAT 'CUSTOM' (formatter='pxfwritable_export');
```

2. Insert some data into the `names_in_mysql_w` table. For example:

```
=# INSERT INTO names_in_mysql_w VALUES (3, 'Muhammad', 'Ali');
```

3. Use the `names_in_mysql` readable external table that you created in the previous section to view the new data in the `names` MySQL table:

```
gpadmin=#  SELECT * FROM names_in_mysql;
 id |    name      |    last
----+------------+--------
  1 |    John      |    Smith
  2 |    Mary      |    Blake
  3 |    Muhammad |    Ali
(3 rows)
```

Example: Reading From and Writing to an Oracle Table

In this example, you:

- Create an Oracle user and assign all privileges on the table to the user
- Create an Oracle table, and insert data into the table
- Configure the PXF JDBC connector to access the Oracle database
- Create a PXF readable external table that references the Oracle table

- Read the data in the Oracle table using PXF
- Create a PXF writable external table that references the Oracle table
- Write data to the Oracle table using PXF
- Read the data in the Oracle table again

Create an Oracle Table

Perform the following steps to create an Oracle table named `countries` in the schema `oracleuser`, and grant a user named `oracleuser` all the necessary privileges:

1. Identify the host name and port of your Oracle server.
2. Connect to the Oracle database as the `system` user:

```
$ sqlplus system
```

3. Create a user named `oracleuser` and assign the password `mypassword` to it:

```
> CREATE USER oracleuser IDENTIFIED BY mypassword;
```

4. Assign user `oracleuser` enough privileges to login, create and modify a table:

```
> GRANT CREATE SESSION TO oracleuser;
> GRANT CREATE TABLE TO oracleuser;
> GRANT UNLIMITED TABLESPACE TO oracleuser;
> exit
```

5. Log in as user `oracleuser`:

```
$ sqlplus oracleuser
```

6. Create a table named `countries`, insert some data into this table, and commit the transaction:

```
> CREATE TABLE countries (country_id int, country_name varchar(40), population
float);
> INSERT INTO countries (country_id, country_name, population) values (3, 'Port
ugal', 10.28);
> INSERT INTO countries (country_id, country_name, population) values (24, 'Zam
bia', 17.86);
> COMMIT;
```

Configure the Oracle Connector

You must create a JDBC server configuration for Oracle, download the Oracle driver JAR file to your system, copy the JAR file to the PXF user configuration directory, synchronize the PXF configuration, and then restart PXF.

This procedure will typically be performed by the Greenplum Database administrator.

1. Download the Oracle JDBC driver and place it under `$PXF_BASE/lib` of your Greenplum Database master host. If you relocated $PXF_BASE, make sure you use the updated location. You can download a Oracle JDBC driver from your preferred download location.

The following example places a driver downloaded from Oracle webiste under `$PXF_BASE/lib` of the Greenplum Database master:

1. If you did not relocate `$PXF_BASE`, run the following from the Greenplum master:

   ```
   gpadmin@gpmaster$ scp ojdbc10.jar gpadmin@gpmaster:/usr/local/pxf-gp<vers
   ion>/lib/
   ```

2. If you relocated `$PXF_BASE`, run the following from the Greenplum master:

   ```
   gpadmin@gpmaster$ scp ojdbc10.jar gpadmin@gpmaster:$PXF_BASE/lib/
   ```

2. Synchronize the PXF configuration, and then restart PXF:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   gpadmin@gpmaster$ pxf cluster restart
   ```

3. Create a JDBC server configuration for Oracle as described in Example Configuration Procedure, naming the server directory `oracle`. The `jdbc-site.xml` file contents should look similar to the following (substitute your Oracle host system for `oracleserverhost`, and the value of your Oracle service name for `orcl`):

   ```
   <?xml version="1.0" encoding="UTF-8"?>
   <configuration>
       <property>
           <name>jdbc.driver</name>
           <value>oracle.jdbc.driver.OracleDriver</value>
           <description>Class name of the JDBC driver</description>
       </property>
       <property>
           <name>jdbc.url</name>
           <value>jdbc:oracle:thin:@oracleserverhost:1521/orcl</value>
           <description>The URL that the JDBC driver can use to connect to the dat
   abase</description>
       </property>
       <property>
           <name>jdbc.user</name>
           <value>oracleuser</value>
           <description>User name for connecting to the database</description>
       </property>
       <property>
           <name>jdbc.password</name>
           <value>mypassword</value>
           <description>Password for connecting to the database</description>
       </property>
   </configuration>
   ```

4. Synchronize the PXF server configuration to the Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

Read from the Oracle Table

Perform the following procedure to create a PXF external table that references the `countries` Oracle table that you created in the previous section, and reads the data in the table:

1. Create the PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE EXTERNAL TABLE oracle_countries (country_id int, country_name
varchar, population float)
          LOCATION('pxf://oracleuser.countries?PROFILE=jdbc&SERVER=oracle')
          FORMAT 'CUSTOM' (formatter='pxfwritable_import');
```

2. Display all rows of the `oracle_countries` table:

```
gpadmin=# SELECT * FROM oracle_countries ;
country_id | country_name | population
-----------+--------------+------------
         3 | Portugal     |      10.28
        24 | Zambia       |      17.86
(2 rows)
```

Write to the Oracle Table

Perform the following procedure to insert some data into the `countries` Oracle table and then read from the table. You must create a new external table for the write operation.

1. Create a writable PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE WRITABLE EXTERNAL TABLE oracle_countries_write (country_id int
, country_name varchar, population float)
          LOCATION('pxf://oracleuser.countries?PROFILE=jdbc&SERVER=oracle')
          FORMAT 'CUSTOM' (formatter='pxfwritable_export');
```

2. Insert some data into the `oracle_countries_write` table. For example:

```
gpadmin=# INSERT INTO oracle_countries_write VALUES (66, 'Colombia', 50.34);
```

3. Use the `oracle_countries` readable external table that you created in the previous section to view the new data in the `countries` Oracle table:

```
gpadmin=#  SELECT * FROM oracle_countries;
country_id | country_name | population
------------+--------------+------------
         3 | Portugal     |      10.28
        24 | Zambia       |      17.86
        66 | Colombia     |      50.34
(3 rows)
```

Example: Reading From and Writing to a Trino (formerly Presto SQL) Table

In this example, you:

- Create an in-memory Trino table and insert data into the table
- Configure the PXF JDBC connector to access the Trino database
- Create a PXF readable external table that references the Trino table

- Read the data in the Trino table using PXF
- Create a PXF writable external table the references the Trino table
- Write data to the Trino table using PXF
- Read the data in the Trino table again

Create a Trino Table

This example assumes that your Trino server has been configured with the included `memory` connector. See Trino Documentation - Memory Connector for instructions on configuring this connector.

Create a Trino table named `names` and insert some data into this table:

```
> CREATE TABLE memory.default.names(id int, name varchar, last varchar);
> INSERT INTO memory.default.names(1, 'John', 'Smith'), (2, 'Mary', 'Blake');
```

Configure the Trino Connector

You must create a JDBC server configuration for Trino, download the Trino driver JAR file to your system, copy the JAR file to the PXF user configuration directory, synchronize the PXF configuration, and then restart PXF.

This procedure will typically be performed by the Greenplum Database administrator.

1. Log in to the Greenplum Database master node:

   ```
   $ ssh gpadmin@<gpmaster>
   ```

2. Download the Trino JDBC driver and place it under `$PXF_BASE/lib`. If you relocated $PXF_BASE, make sure you use the updated location. See Trino Documentation - JDBC Driver for instructions on downloading the Trino JDBC driver. The following example downloads the driver and places it under `$PXF_BASE/lib`:

   1. If you did not relocate `$PXF_BASE`, run the following from the Greenplum master:

      ```
      gpadmin@gpmaster$ cd /usr/local/pxf-gp<version>/lib
      gpadmin@gpmaster$ wget <url-to-trino-jdbc-driver>
      ```

   2. If you relocated `$PXF_BASE`, run the following from the Greenplum master:

      ```
      gpadmin@gpmaster$ cd $PXF_BASE/lib
      gpadmin@gpmaster$ wget <url-to-trino-jdbc-driver>
      ```

3. Synchronize the PXF configuration, and then restart PXF:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   gpadmin@gpmaster$ pxf cluster restart
   ```

4. Create a JDBC server configuration for Trino as described in Example Configuration Procedure, naming the server directory `trino`. The `jdbc-site.xml` file contents should look similar to the following (substitute your Trino host system for `trinoserverhost`):

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
    <property>
        <name>jdbc.driver</name>
        <value>io.trino.jdbc.TrinoDriver</value>
        <description>Class name of the JDBC driver</description>
    </property>
    <property>
        <name>jdbc.url</name>
        <value>jdbc:trino://trinoserverhost:8443</value>
        <description>The URL that the JDBC driver can use to connect to the dat
abase</description>
    </property>
    <property>
        <name>jdbc.user</name>
        <value>trino-user</value>
        <description>User name for connecting to the database</description>
    </property>
    <property>
        <name>jdbc.password</name>
        <value>trino-pw</value>
        <description>Password for connecting to the database</description>
    </property>

    <!-- Connection properties -->
    <property>
        <name>jdbc.connection.property.SSL</name>
        <value>true</value>
        <description>Use HTTPS for connections; authentication using username/p
assword requires SSL to be enabled.</description>
    </property>
</configuration>
```

5. If your Trino server has been configured with a Globally Trusted Certificate, you can skip this step. If your Trino server has been configured to use Corporate trusted certificates or Generated self-signed certificates, PXF will need a copy of the server's certificate in a PEM-encoded file or a Java Keystore (JKS) file.

   Note: You do not need the Trino server's private key.

   Copy the certificate to `$PXF_BASE/servers/trino`; storing the server's certificate inside `$PXF_BASE/servers/trino` ensures that `pxf cluster sync` copies the certificate to all segment hosts.

   ```
   $ cp <path-to-trino-server-certificate> /usr/local/pxf-gp<version>/servers/trin
   o
   ```

   Add the following connection properties to the `jdbc-site.xml` file that you created in the previous step. Here, `trino.cert` is the name of the certificate file that you copied into `$PXF_BASE/servers/trino`:

   ```
   <configuration>
   ...
       <property>
           <name>jdbc.connection.property.SSLTrustStorePath</name>
           <value>/usr/local/pxf-gp<version>/servers/trino/trino.cert</value>
           <description>The location of the Java TrustStore file that will be used
   ```

```
 to validate HTTPS server certificates.</description>
    </property>
    <!-- the following property is only required if the server's certificate is
 stored in a JKS file; if using a PEM-encoded file, it should be omitted.-->
    <!--
    <property>
        <name>jdbc.connection.property.SSLTrustStorePassword</name>
        <value>java-keystore-password</value>
        <description>The password for the TrustStore.</description>
    </property>
    -->
</configuration>
```

6. Synchronize the PXF server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Read from a Trino Table

Perform the following procedure to create a PXF external table that references the `names` Trino table and reads the data in the table:

1. Create the PXF external table specifying the `jdbc` profile. Specify the Trino catalog and schema in the `LOCATION` URL. The following example reads the `names` table located in the `default` schema of the `memory` catalog:

```
CREATE EXTERNAL TABLE pxf_trino_memory_names (id int, name text, last text)
LOCATION('pxf://memory.default.names?PROFILE=jdbc&SERVER=trino')
FORMAT 'CUSTOM' (formatter='pxfwritable_import');
```

2. Display all rows of the `pxf_trino_memory_names` table:

```
gpadmin=# SELECT * FROM pxf_trino_memory_names;
 id | name | last
----+------+-------
  1 | John | Smith
  2 | Mary | Blake
(2 rows)
```

Write to the Trino Table

Perform the following procedure to insert some data into the `names` Trino table and then read from the table. You must create a new external table for the write operation.

1. Create a writable PXF external table specifying the `jdbc` profile. For example:

```
gpadmin=# CREATE WRITABLE EXTERNAL TABLE pxf_trino_memory_names_w (id int, name
 text, last text)
          LOCATION('pxf://memory.default.names?PROFILE=jdbc&SERVER=trino')
          FORMAT 'CUSTOM' (formatter='pxfwritable_export');
```

2. Insert some data into the `pxf_trino_memory_names_w` table. For example:

```
gpadmin=# INSERT INTO pxf_trino_memory_names_w VALUES (3, 'Muhammad', 'Ali');
```

3. Use the `pxf_trino_memory_names` readable external table that you created in the previous section to view the new data in the `names` Trino table:

```
gpadmin=# SELECT * FROM pxf_trino_memory_names;
 id |   name    | last
----+----------+-------
  1 | John      | Smith
  2 | Mary      | Blake
  3 | Muhammad | Ali
(3 rows)
```

Example: Using a Named Query with PostgreSQL

In this example, you:

- Use the PostgreSQL database `pgtestdb`, user `pxfuser1`, and PXF JDBC connector server configuration `pgsrvcfg` that you created in Example: Reading From and Writing to a PostgreSQL Database.
- Create two PostgreSQL tables and insert data into the tables.
- Assign all privileges on the tables to `pxfuser1`.
- Define a named query that performs a complex SQL statement on the two PostgreSQL tables, and add the query to the `pgsrvcfg` JDBC server configuration.
- Create a PXF readable external table definition that matches the query result tuple and also specifies read partitioning options.
- Read the query results, making use of PXF column projection and filter pushdown.

Create the PostgreSQL Tables and Assign Permissions

Perform the following procedure to create PostgreSQL tables named `customers` and `orders` in the `public` schema of the database named `pgtestdb`, and grant the user named `pxfuser1` all privileges on these tables:

1. Identify the host name and port of your PostgreSQL server.
2. Connect to the `pgtestdb` PostgreSQL database as the `postgres` user. For example, if your PostgreSQL server is running on the default port on the host named `pserver`:

```
$ psql -U postgres -h pserver -d pgtestdb
```

3. Create a table named `customers` and insert some data into this table:

```
CREATE TABLE customers(id int, name text, city text, state text);
INSERT INTO customers VALUES (111, 'Bill', 'Helena', 'MT');
INSERT INTO customers VALUES (222, 'Mary', 'Athens', 'OH');
INSERT INTO customers VALUES (333, 'Tom', 'Denver', 'CO');
INSERT INTO customers VALUES (444, 'Kate', 'Helena', 'MT');
INSERT INTO customers VALUES (555, 'Harry', 'Columbus', 'OH');
INSERT INTO customers VALUES (666, 'Kim', 'Denver', 'CO');
INSERT INTO customers VALUES (777, 'Erik', 'Missoula', 'MT');
INSERT INTO customers VALUES (888, 'Laura', 'Athens', 'OH');
```

```
INSERT INTO customers VALUES (999, 'Matt', 'Aurora', 'CO');
```

4. Create a table named `orders` and insert some data into this table:

```
CREATE TABLE orders(customer_id int, amount int, month int, year int);
INSERT INTO orders VALUES (111, 12, 12, 2018);
INSERT INTO orders VALUES (222, 234, 11, 2018);
INSERT INTO orders VALUES (333, 34, 7, 2018);
INSERT INTO orders VALUES (444, 456, 111, 2018);
INSERT INTO orders VALUES (555, 56, 11, 2018);
INSERT INTO orders VALUES (666, 678, 12, 2018);
INSERT INTO orders VALUES (777, 12, 9, 2018);
INSERT INTO orders VALUES (888, 120, 10, 2018);
INSERT INTO orders VALUES (999, 120, 11, 2018);
```

5. Assign user `pxfuser1` all privileges on tables `customers` and `orders`, and then exit the `psql` subsystem:

```
GRANT ALL ON customers TO pxfuser1;
GRANT ALL ON orders TO pxfuser1;
\q
```

Configure the Named Query

In this procedure you create a named query text file, add it to the `pgsrvcfg` JDBC server configuration, and synchronize the PXF configuration to the Greenplum Database cluster.

This procedure will typically be performed by the Greenplum Database administrator.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Navigate to the JDBC server configuration directory `pgsrvcfg`. For example:

```
gpadmin@gpmaster$ cd $PXF_BASE/servers/pgsrvcfg
```

3. Open a query text file named `pg_order_report.sql` in a text editor and copy/paste the following query into the file:

```
SELECT c.name, c.city, sum(o.amount) AS total, o.month
  FROM customers c JOIN orders o ON c.id = o.customer_id
  WHERE c.state = 'CO'
GROUP BY c.name, c.city, o.month
```

4. Save the file and exit the editor.

5. Synchronize these changes to the PXF configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Read the Query Results

Perform the following procedure on your Greenplum Database cluster to create a PXF external table

that references the query file that you created in the previous section, and then reads the query result data:

1. Create the PXF external table specifying the `jdbc` profile. For example:

   ```
   CREATE EXTERNAL TABLE pxf_queryres_frompg(name text, city text, total int, mont
   h int)
     LOCATION ('pxf://query:pg_order_report?PROFILE=jdbc&SERVER=pgsrvcfg&PARTITION
   _BY=month:int&RANGE=1:13&INTERVAL=3')
   FORMAT 'CUSTOM' (FORMATTER='pxfwritable_import');
   ```

   With this partitioning scheme, PXF will issue 4 queries to the remote SQL database, one query per quarter. Each query will return customer names and the total amount of all of their orders in a given month, aggregated per customer, per month, for each month of the target quarter. Greenplum Database will then combine the data into a single result set for you when you query the external table.

2. Display all rows of the query result:

   ```
   SELECT * FROM pxf_queryres_frompg ORDER BY city, total;

    name |  city  | total | month
   ------+--------+-------+-------
    Matt | Aurora |   120 |    11
    Tom  | Denver |    34 |     7
    Kim  | Denver |   678 |    12
   (3 rows)
   ```

3. Use column projection to display the order total per city:

   ```
   SELECT city, sum(total) FROM pxf_queryres_frompg GROUP BY city;

     city  | sum
   --------+-----
    Aurora | 120
    Denver | 712
   (2 rows)
   ```

   When you execute this query, PXF requests and retrieves query results for only the `city` and `total` columns, reducing the amount of data sent back to Greenplum Database.

4. Provide additional filters and aggregations to filter the `total` in PostgreSQL:

   ```
   SELECT city, sum(total) FROM pxf_queryres_frompg
               WHERE total > 100
               GROUP BY city;

     city  | sum
   --------+-----
    Denver | 678
    Aurora | 120
   (2 rows)
   ```

   In this example, PXF will add the `WHERE` filter to the subquery. This filter is pushed to and executed on the remote database system, reducing the amount of data that PXF sends back to Greenplum Database. The `GROUP BY` aggregation, however, is not pushed to the remote

and is performed by Greenplum.

Accessing Files on a Network File System with PXF

You can use PXF to read data that resides on a network file system mounted on your Greenplum Database hosts. PXF supports reading and writing the following file types from a network file system:

File Type	Profile Name	Operations Supported
delimited single line text	file:text	read, write
delimited single line comma-separated values of text	file:csv	read, write
delimited text with quoted linefeeds	file:text:multi	read
Avro	file:avro	read, write
JSON	file:json	read
ORC	file:orc	read, write
Parquet	file:parquet	read, write

PXF does not support user impersonation when you access a network file system. PXF accesses a file as the operating system user that started the PXF process, usually `gpadmin`.

Reading from, and writing to (where supported), a file of these types on a network file system is similar to reading/writing the file type on Hadoop.

Prerequisites

Before you use PXF to access files on a network file system, ensure that:

- You can identify the PXF runtime configuration directory (`$PXF_BASE`).
- You have configured PXF, and PXF is running on each Greenplum Database host. See Configuring PXF for additional information.
- All files are accessible by `gpadmin` or by the operating system user that started the PXF process.
- The network file system is correctly mounted at the same local mount point on every Greenplum Database host.
- You can identify the mount or share point of the network file system.
- You have created one or more named PXF server configurations as described in Configuring a PXF Network File System Server.

Configuring a PXF Network File System Server

Before you use PXF to access a file on a network file system, you must create a server configuration and then synchronize the PXF configuration to all Greenplum hosts. This procedure will typically be performed by the Greenplum Database administrator.

Use the server template configuration file `$PXF_HOME/templates/pxf-site.xml` when you configure a network file system server for PXF. This template file includes the mandatory property `pxf.fs.basePath` that you configure to identify the network file system share path. PXF considers the file path that you specify in a `CREATE EXTERNAL TABLE LOCATION` clause that uses this server to be relative to this share path.

PXF does not support user impersonation when you access a network file system; you must explicitly turn off user impersonation in a network file system server configuration.

1. Log in to the Greenplum Database master node:

```
$ ssh gpadmin@<gpmaster>
```

2. Choose a name for the file system server. You will provide the name to Greenplum users that you choose to allow to read from or write to files on the network file system.

   Note: The server name `default` is reserved.

3. Create the `$PXF_BASE/servers/<server_name>` directory. For example, use the following command to create a file system server configuration named `nfssrvcfg`:

```
gpadmin@gpmaster$ mkdir $PXF_BASE/servers/nfssrvcfg
```

4. Copy the PXF `pxf-site.xml` template file to the `nfssrvcfg` server configuration directory. For example:

```
gpadmin@gpmaster$ cp $PXF_HOME/templates/pxf-site.xml $PXF_BASE/servers/nfssrvc
fg/
```

5. Open the template server configuration file in the editor of your choice, and uncomment and provide property values appropriate for your environment. For example, if the file system share point is the directory named `/mnt/extdata/pxffs`, uncomment and set these server properties:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
...
    <property>
        <name>pxf.service.user.impersonation</name>
        <value>false</value>
    </property>

    <property>
        <name>pxf.fs.basePath</name>
        <value>/mnt/extdata/pxffs</value>
    </property>
...
</configuration>
```

6. Save your changes and exit the editor.

7. Synchronize the PXF server configuration to the Greenplum Database cluster:

```
gpadmin@gpmaster$ pxf cluster sync
```

Creating the External Table

The following syntax creates a Greenplum Database external table that references a file on a network file system. Use the appropriate `file:*` profile for the file type that you want to access.

```
CREATE [READABLE | WRITABLE] EXTERNAL TABLE <table_name>
    ( <column_name> <data_type> [, ...] | LIKE <other_table> )
LOCATION ('pxf://<file-path>?PROFILE=file:<file-type>[&SERVER=<server_name>][&<custom-option>=<value>[...]]')
FORMAT '[TEXT|CSV|CUSTOM]' (<formatting-properties>);
```

The specific keywords and values used in the Greenplum Database CREATE EXTERNAL TABLE command are described in the table below.

Keyword	Value
<file-path>	The path to a directory or file on the network file system. PXF considers this file or path as being relative to the `pxf.fs.basePath` property value specified in <server_name>'s server configuration. <file-path> must not specify a relative path nor include the dollar sign (`$`) character.
PROFILE	The `PROFILE` keyword value must specify a `file:<file-type>` identified in the table above.
SERVER=<server_name>	The named server configuration that PXF uses to access the network file system. PXF uses the `default` server if not specified.
<custom-option>=<value>	<custom-option> is profile-specific.
FORMAT <value>	PXF profiles support the `TEXT`, `CSV`, and `CUSTOM` formats.
<formatting-properties>	Formatting properties supported by the profile; for example, the `FORMATTER` or `delimiter`.

The <custom-option>s, FORMAT, and <formatting-properties> that you specify when accessing a file on a network file system are dependent on the <file-type>. Refer to the Hadoop documentation for the <file-type> of interest for these settings.

Example: Reading From and Writing to a CSV File on a Network File System

This example assumes that you have configured and mounted a network file system with the share point `/mnt/extdata/pxffs` on the Greenplum Database master, standby master, and on each segment host.

In this example, you:

- Create a CSV file on the network file system and add data to the file.
- Configure a network file system server for the share point.
- Create a PXF readable external table that references the directory containing the CSV file, and read the data.
- Create a PXF writable external table that references the directory containing the CSV file,

and write some data.

- Read from the original readable external table again.

Create a CSV File

1. Create a directory (relative to the network file system share point) named `/mnt/extdata/pxffs/ex1`:

```
gpadmin@gpmaster$ mkdir -p /mnt/extdata/pxffs/ex1
```

2. Create a CSV file named `somedata.csv` in the directory:

```
$ echo 'Prague,Jan,101,4875.33
Rome,Mar,87,1557.39
Bangalore,May,317,8936.99
Beijing,Jul,411,11600.67' > /mnt/extdata/pxffs/ex1/somedata.csv
```

Create the Network File System Server

Create a server configuration named `nfssrvcfg` with share point `/mnt/extdata/pxffs` as described in Configuring a PXF Network File System Server.

Read Data

Perform the following procedure to create a PXF external table that references the `ex1` directory that you created in a previous section, and then read the data in the `somedata.csv` file in that directory:

1. Create a PXF external table that references `ex1` and that specifies the `file:text` profile. For example:

```
gpadmin=# CREATE EXTERNAL TABLE pxf_read_nfs(location text, month text, num_ord
ers int, total_sales float8)
            LOCATION ('pxf://ex1/?PROFILE=file:text&SERVER=nfssrvcfg')
            FORMAT 'CSV';
```

Because the `nfssrvcfg` server configuration `pxf.fs.basePath` property value is `/mnt/exdata/pxffs`, PXF constructs the path `/mnt/extdata/pxffs/ex1` to read the file.

2. Display all rows of the `pxf_read_nfs` table:

```
gpadmin=# SELECT * FROM pxf_read_nfs ORDER_BY num_orders DESC;
 location  | month | num_orders | total_sales
-----------+-------+------------+-------------
 Beijing   | Jul   |        411 |    11600.67
 Bangalore | May   |        317 |     8936.99
 Prague    | Jan   |        101 |     4875.33
 Rome      | Mar   |         87 |     1557.39
(4 rows)
```

Write Data and Read Again

Perform the following procedure to insert some data into the `ex1` directory and then read the data again. You must create a new external table for the write operation.

1. Create a writable PXF external table that references `ex1` and that specifies the `file:text` profile. For example:

```
gpadmin=# CREATE WRITABLE EXTERNAL TABLE pxf_write_nfs(location text, month tex
t, num_orders int, total_sales float8)
            LOCATION ('pxf://ex1/?PROFILE=file:text&SERVER=nfssrvcfg')
          FORMAT 'CSV' (delimiter=',');
```

2. Insert some data into the `pxf_write_nfs` table. For example:

```
gpadmin=# INSERT INTO pxf_write_nfs VALUES ( 'Frankfurt', 'Mar', 777, 3956.98 )
;
INSERT 0 1
gpadmin=# INSERT INTO pxf_write_nfs VALUES ( 'Cleveland', 'Oct', 3812, 96645.37
 );
INSERT 0 1
```

   PXF writes one or more files to the `ex1/` directory when you insert into the `pxf_write_nfs` table.

3. Use the `pxf_read_nfs` readable external table that you created in the previous section to view the new data you inserted into the `pxf_write_nfs` table:

```
gpadmin=# SELECT * FROM pxf_read_nfs ORDER BY num_orders DESC;
 location  | month | num_orders | total_sales
-----------+-------+------------+-------------
 Cleveland | Oct   |       3812 |    96645.37
 Frankfurt | Mar   |        777 |     3956.98
 Beijing   | Jul   |        411 |    11600.67
 Bangalore | May   |        317 |     8936.99
 Prague    | Jan   |        101 |     4875.33
 Rome      | Mar   |         87 |     1557.39
(6 rows)
```

   When you select from the `pxf_read_nfs` table here, PXF reads the `somedata.csv` file and the new files that it added to the `ex1/` directory in the previous step.

Troubleshooting PXF

PXF Errors

The following table describes some errors you may encounter while using PXF:

Error Message	Discussion
Protocol “pxf” does not exist	**Cause**: The `pxf` extension was not registered. **Solution**: Create (enable) the PXF extension for the database as described in the PXF Enable Procedure.
Invalid URI pxf://<path-to-data>: missing options section	**Cause**: The `LOCATION` URI does not include the profile or other required options. **Solution**: Provide the profile and required options in the URI when you submit the `CREATE EXTERNAL TABLE` command.
PXF server error : Input path does not exist: hdfs://<namenode>:8020/<path-to-file>	**Cause**: The HDFS file that you specified in <path-to-file> does not exist. **Solution**: Provide the path to an existing HDFS file.
PXF server error : NoSuchObjectException(message: <schema>.<hivetable> table not found)	**Cause**: The Hive table that you specified with <schema>.<hivetable> does not exist. **Solution**: Provide the name of an existing Hive table.
PXF server error : Failed connect to localhost:5888; Connection refused (<segment-id> slice<N> <segment-host>:<port> pid=<process-id>) …	**Cause**: The PXF Service is not running on <segment-host>. **Solution**: Restart PXF on <segment-host>.
PXF server error: Permission denied: user=<user>, access=READ, inode=“<filepath>”:-rw——-	**Cause**: The Greenplum Database user that executed the PXF operation does not have permission to access the underlying Hadoop service (HDFS or Hive). See Configuring the Hadoop User, User Impersonation, and Proxying.
PXF server error: PXF service could not be reached. PXF is not running in the tomcat container	**Cause**: The `pxf` extension was updated to a new version but the PXF server has not been updated to a compatible version. **Solution**: Ensure that the PXF server has been updated and restarted on all hosts.
ERROR: could not load library “/usr/local/greenplum-db-x.x.x/lib/postgresql/pxf.so”	**Cause**: Some steps have not been completed after a Greenplum Database upgrade or migration, such as `pxf cluster register`. **Solution**: Make sure you follow the steps outlined for PXF Upgrade and Migration.

Most PXF error messages include a `HINT` that you can use to resolve the error, or to collect more information to identify the error.

PXF Logging

Refer to the Logging topic for more information about logging levels, configuration, and the `pxf-app.out` and `pxf-service.log` log files.

Addressing PXF JDBC Connector Time Zone Errors

You use the PXF JDBC connector to access data stored in an external SQL database. Depending upon the JDBC driver, the driver may return an error if there is a mismatch between the default time zone set for the PXF Service and the time zone set for the external SQL database.

For example, if you use the PXF JDBC connector to access an Oracle database with a conflicting time zone, PXF logs an error similar to the following:

```
java.io.IOException: ORA-00604: error occurred at recursive SQL level 1
ORA-01882: timezone region not found
```

Should you encounter this error, you can set default time zone option(s) for the PXF Service in the `$PXF_BASE/conf/pxf-env.sh` configuration file, `PXF_JVM_OPTS` property setting. For example, to set the time zone:

```
export PXF_JVM_OPTS="<current_settings> -Duser.timezone=America/Chicago"
```

You can use the `PXF_JVM_OPTS` property to set other Java options as well.

As described in previous sections, you must synchronize the updated PXF configuration to the Greenplum Database cluster and restart the PXF Service on each host.

About PXF External Table Child Partitions

Greenplum Database supports partitioned tables, and permits exchanging a leaf child partition with a PXF external table.

When you read from a partitioned Greenplum table where one or more partitions is a PXF external table and there is no data backing the external table path, PXF returns an error and the query fails. This default PXF behavior is not optimal in the partitioned table case; an empty child partition is valid and should not cause a query on the parent table to fail.

The `IGNORE_MISSING_PATH` PXF custom option is a boolean that specifies the action to take when the external table path is missing or invalid. The default value is `false`, PXF returns an error when it encounters a missing path. If the external table is a child partition of a Greenplum table, you want PXF to ignore a missing path error, so set this option to `true`.

For example, PXF ignores missing path errors generated from the following external table:

```
CREATE EXTERNAL TABLE ext_part_87 (id int, some_date date)
  LOCATION ('pxf://bucket/path/?PROFILE=s3:parquet&SERVER=s3&IGNORE_MISSING_PATH=true'
)
FORMAT 'CUSTOM' (formatter = 'pxfwritable_import');
```

The `IGNORE_MISSING_PATH` custom option applies only to file-based profiles, including `*:text`, `*:parquet`, `*:avro`, `*:json`, `*:AvroSequenceFile`, and `*:SequenceFile`. This option is *not available* when the external table specifies the `hbase`, `hive[:*]`, or `jdbc` profiles, or when reading from S3 using S3-Select.

Addressing Hive MetaStore Connection Errors

The PXF Hive connector uses the Hive MetaStore to determine the HDFS locations of Hive tables. Starting in PXF version 6.2.1, PXF retries a failed connection to the Hive MetaStore a single time. If you encounter one of the following error messages or exceptions when accessing Hive via a PXF external table, consider increasing the retry count:

- `Failed to connect to the MetaStore Server.`
- `Could not connect to meta store ...`
- `org.apache.thrift.transport.TTransportException: null`

PXF uses the `hive-site.xml` `hive.metastore.failure.retries` property setting to identify the maximum number of times it will retry a failed connection to the Hive MetaStore. The `hive-site.xml` file resides in the configuration directory of the PXF server that you use to access Hive.

Perform the following procedure to configure the number of Hive MetaStore connection retries that PXF will attempt; you may be required to add the `hive.metastore.failure.retries` property to the `hive-site.xml` file:

1. Log in to the Greenplum Database master node.
2. Identify the name of your Hive PXF server.
3. Open the `$PXF_BASE/servers/<hive-server-name>/hive-site.xml` file in the editor of your choice, add the `hive.metastore.failure.retries` property if it does not already exist in the file, and set the value. For example, to configure 5 retries:

   ```
   <property>
       <name>hive.metastore.failure.retries</name>
       <value>5</value>
   </property>
   ```

4. Save the file and exit the editor.
5. Synchronize the PXF configuration to all hosts in your Greenplum Database cluster:

   ```
   gpadmin@gpmaster$ pxf cluster sync
   ```

6. Re-run the failing SQL external table command.

PXF Utility Reference

The Greenplum Platform Extension Framework (PXF) includes the following utility reference pages:

- pxf cluster
- pxf

pxf cluster

Manage the PXF configuration and the PXF Service instance on all Greenplum Database hosts.

Synopsis

```
pxf cluster <command> [<option>]
```

where `<command>` is:

```
help
init (deprecated)
migrate
prepare
register
reset (deprecated)
restart
start
status
stop
sync
```

Description

The `pxf cluster` utility command manages PXF on the master, standby master, and on all Greenplum Database segment hosts. You can use the utility to:

- Start, stop, and restart the PXF Service instance on the master, standby master, and all segment hosts.
- Display the status of the PXF Service instance on the master, standby master, and all segment hosts.
- Synchronize the PXF configuration from the Greenplum Database master host to the standby master and to all segment hosts.
- Copy the PXF extension control file from the PXF installation on each host to the Greenplum installation on the host after a Greenplum upgrade.
- Prepare a new `$PXF_BASE` runtime configuration directory.

- Migrate PXF 5 `$PXF_CONF` configuration to `$PXF_BASE`.

`pxf cluster` requires a running Greenplum Database cluster. You must run the utility on the Greenplum Database master host.

(If you want to manage the PXF Service instance on a specific segment host, use the `pxf` utility. See `pxf`.)

Commands

help
: Display the `pxf cluster` help message and then exit.

init (deprecated)
: The command is equivalent to the `register` command.

migrate
: Migrate the configuration in a PXF 5 `$PXF_CONF` directory to `$PXF_BASE` on each Greenplum Database host. When you run the command, you must identify the PXF 5 configuration directory via an environment variable named `PXF_CONF`. PXF migrates the version 5 configuration to `$PXF_BASE`, copying and merging files and directories as necessary.
You must manually migrate any `pxf-log4j.properties` customizations to the `pxf-log4j2.xml` file.

prepare
: Prepare a new `$PXF_BASE` directory on each Greenplum Database host. When you run the command, you must identify the new PXF runtime configuration directory via an environment variable named `PXF_BASE`. PXF copies runtime configuration file templates and directories to this `$PXF_BASE`.

register
: Copy the PXF extension control file from the PXF installation on each host to the Greenplum installation on the host. This command requires that `$GPHOME` be set, and is run once after you install PXF 6.x the first time, or run after you upgrade your Greenplum Database installation.

reset (deprecated)
: The command is a no-op.

restart
: Stop, and then start, the PXF Service instance on the master, standby master, and all segment hosts.

start
: Start the PXF Service instance on the master, standby master, and all segment hosts.

status
: Display the status of the PXF Service instance on the master, standby master, and all segment hosts.

stop
: Stop the PXF Service instance on the master, standby master, and all segment hosts.

sync
: Synchronize the PXF configuration (`$PXF_BASE`) from the master to the standby master and to

all Greenplum Database segment hosts. By default, this command updates files on and copies files to the remote. You can instruct PXF to also delete files during the synchronization; see Options.

If you have updated the PXF user configuration or add new JAR or native library dependencies, you must also restart PXF after you synchronize the PXF configuration.

Options

The `pxf cluster sync` command takes the following option:

–d | ––delete

Delete any files in the PXF user configuration on the standby master and segment hosts that are not also present on the master host.

Examples

Stop the PXF Service instance on the master, standby master, and all segment hosts:

```
$ pxf cluster stop
```

Synchronize the PXF configuration to the standby and all segment hosts, deleting files that do not exist on the master host:

```
$ pxf cluster sync --delete
```

See Also

`pxf`

pxf

Manage the PXF configuration and the PXF Service instance on the local Greenplum Database host.

Synopsis

```
pxf <command> [<option>]
```

where <command> is:

```
cluster
help
init (deprecated)
migrate
prepare
register
reset (deprecated)
restart
start
status
stop
sync
version
```

Description

The `pxf` utility manages the PXF configuration and the PXF Service instance on the local Greenplum Database host. You can use the utility to:

- Synchronize the PXF configuration from the master to the standby master or to a segment host.
- Start, stop, or restart the PXF Service instance on the master, standby master, or a specific segment host, or display the status of the PXF Service instance running on the master, standby master, or a segment host.
- Copy the PXF extension control file from a PXF installation on the host to the Greenplum installation on the host after a Greenplum upgrade.
- Prepare a new `$PXF_BASE` runtime configuration directory on the host.

(Use the `pxf cluster` command to prepare a new `$PXF_BASE` on all hosts, copy the PXF extension control file to `$GPHOME` on all hosts, synchronize the PXF configuration to the Greenplum Database cluster, or to start, stop, or display the status of the PXF Service instance on all hosts in the cluster.)

Commands

cluster
: Manage the PXF configuration and the PXF Service instance on all Greenplum Database hosts. See `pxf cluster`.

help
: Display the `pxf` management utility help message and then exit.

init (deprecated)
: The command is equivalent to the `register` command.

migrate
: Migrate the configuration in a PXF 5 `$PXF_CONF` directory to `$PXF_BASE` on the host. When you run the command, you must identify the PXF 5 configuration directory via an environment variable named `PXF_CONF`. PXF migrates the version 5 configuration to the current `$PXF_BASE`, copying and merging files and directories as necessary.
You must manually migrate any `pxf-log4j.properties` customizations to the `pxf-log4j2.xml` file.

prepare
: Prepare a new `$PXF_BASE` directory on the host. When you run the command, you must identify the new PXF runtime configuration directory via an environment variable named `PXF_BASE`. PXF copies runtime configuration file templates and directories to this `$PXF_BASE`.

register
: Copy the PXF extension files from the PXF installation on the host to the Greenplum installation on the host. This command requires that `$GPHOME` be set, and is run once after you install PXF 6.x the first time, or run when you upgrade your Greenplum Database installation.

reset (deprecated)
: The command is a no-op.

restart
: Restart the PXF Service instance running on the local master, standby master, or segment host.

start
: Start the PXF Service instance on the local master, standby master, or segment host.

status
: Display the status of the PXF Service instance running on the local master, standby master, or segment host.

stop
: Stop the PXF Service instance running on the local master, standby master, or segment host.

sync
: Synchronize the PXF configuration (`$PXF_BASE`) from the master to a specific Greenplum Database standby master or segment host. You must run `pxf sync` on the master host. By default, this command updates files on and copies files to the remote. You can instruct PXF to also delete files during the synchronization; see Options.

version
: Display the PXF version and then exit.

Options

The `pxf sync` command, which you must run on the Greenplum Database master host, takes the following option and argument:

-d | --delete
: Delete any files in the PXF user configuration on `<gphost>` that are not also present on the master host. If you specify this option, you must provide it on the command line before `<gphost>`.

<gphost>
: The Greenplum Database host to which to synchronize the PXF configuration. Required. `<gphost>` must identify the standy master host or a segment host.

Examples

Start the PXF Service instance on the local Greenplum host:

```
$ pxf start
```

See Also

`pxf cluster`