

Performance Tuning and Sizing Guidelines for PowerExchange for Amazon Redshift on the Spark Engine

Abstract

When you use PowerExchange for Amazon Redshift on Spark engine, multiple factors such as data set size, hardware parameters, and mapping parameters, impact the adapter performance. You can optimize the performance by analyzing your data set size, using the recommended hardware, and tuning these parameters appropriately. This article describes general reference guidelines to help you analyze the data set and tune the performance of PowerExchange for Amazon Redshift on the Spark engine.

Supported Versions

- PowerExchange® for Amazon Redshift 10.2

Table of Contents

Overview.	2
Performance Sizing Recommendations.	2
Performance Tuning Parameters.	4
Spark Tuning Parameters.	4
Tune the Amazon Redshift Cluster Properties.	6
Optimizing the Amazon Redshift Mapping.	6
Troubleshooting Spark Tuning Parameters.	7

Overview

Performance tuning is an iterative process where you analyze the performance, use guidelines to estimate and define parameters that impact the performance, monitor, and adjust the results as required. Sizing is a process where you analyze the data set and use the recommended hardware to improve the performance.

This document describes the key data size, hardware, and Spark tuning parameters that you can tune to optimize the performance of PowerExchange for Amazon Redshift on the Spark engine.

Note: The performance testing results listed in this article are based on observations in an internal Informatica environment using data from real-world scenarios. The performance of PowerExchange for Amazon Redshift might vary based on individual environments and other parameters even when you use the same data.

Performance Sizing Recommendations

You can tune the hardware parameters based on the data set and using the recommended hardware capacities to optimize the performance of the machine where the Data Integration Service runs.

Performance sizing recommendations depend on the size of the data set. Consider the sizing recommendations for the following types of data set:

Sandbox data set

A data set is considered to be a sandbox data set if you read from or write to less than 1 GB of data.

For a sandbox data set, you can read 1 GB of data from an Amazon Redshift source to an HDFS database and write data to an Amazon Redshift target.

The following table describes the sizing recommendations for the system that processes sandbox data sets:

Requirement	Sizing Recommendation	Hadoop Services Recommendation
VCPUs	2	$4 * 1 = 4$
Memory	8 GB	$2 * 1 = 2$ GB
Storage	14 GB	$1 * 2 = 2$ GB

Small data set

A data set is considered to be a small data set if you read from or write to less than 100 GB of data.

For a small data set, you can read 100 GB of data from an Amazon Redshift source to an HDFS database in 5:00:00 (HH:MM:SS) hours and write data to an Amazon Redshift target in 5:00:00 (HH:MM:SS) hours.

The following table describes the sizing recommendations for the system that processes small data sets:

Requirement	Sizing Recommendation	Hadoop Services Recommendation
VCPUs	2	$5 * 1 = 5$
Memory	8 GB	$6 * 1 = 6$ GB
Storage	14 GB	$100 * 3 = 300$ GB

Medium data set

A data set is considered to be a medium data set if you read from or write to less than 1 TB of data.

For a medium data set, you can read 1 TB of data from an Amazon Redshift source to an HDFS database in 5:00:00 (HH:MM:SS) hours and write data to an Amazon Redshift target in 5:00:00 (HH:MM:SS) hours.

The following table describes the sizing recommendations for the system that processes medium data sets:

Requirement	Sizing Recommendation	Hadoop Services Recommendation
VCPUs	2	$5 * 6 = 30$
Memory	8 GB	$6 * 6 = 36$ GB
Storage	14 GB	$1 * 3 = 3$ TB

Large data set

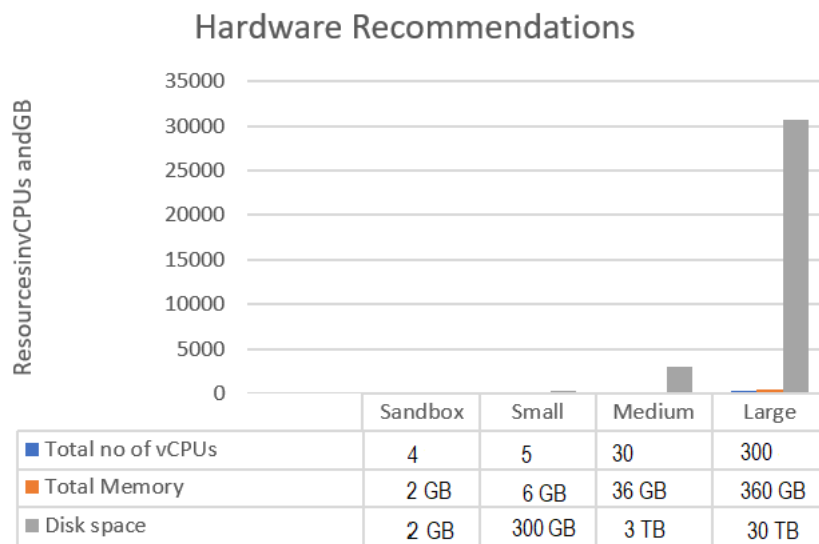
A data set is considered to be a large data set if you read from or write to less than 10 TB of data.

For a medium data set, you can read 10 TB of data from an Amazon Redshift source to an HDFS database in 5:00:00 (HH:MM:SS) hours and write data to an Amazon Redshift target in 5:00:00 (HH:MM:SS) hours.

The following table describes the sizing recommendations for the system that processes large data sets:

Requirement	Sizing Recommendation	Hadoop Services Recommendation
VCPUs	2	5 * 60 = 300
Memory	8 GB	6 * 60 = 360 GB
Storage	14 GB	10 * 3 = 30 TB

The following image represents the hardware requirements along with the change in data set and change in individual component settings:



Performance Tuning Parameters

You can optimize the performance of PowerExchange for Amazon Redshift mappings by tuning the following parameters:

- Spark Tuning Parameters
- Amazon Redshift cluster properties

Note: The final result of the Amazon Redshift mapping performance depends on the mapping complexity as well.

Spark Tuning Parameters

You can tune the following Spark parameters to optimize the performance:

spark.executor.memory

The amount of memory that is used in each executor process.

The following table describes the tuning recommendations for the Spark parameters according to the size of the data set:

Data Set	Recommended Value
Sandbox	1 GB
Small	1 GB
Medium	1 GB
Large	1 GB

spark.executor.cores

The number of cores that is used in each executor process.

The following table describes the tuning recommendations for the Spark parameters according to the size of the data set:

Data Set	Recommended Value
Sandbox	1
Small	1
Medium	1
Large	1

spark.executor.instances

The count of executors to run a process if dynamic allocation is enabled.

The following table describes the tuning recommendations for the Spark parameters according to the size of the data set:

Data Set	Recommended Value
Sandbox	10
Small	3
Medium	30
Large	300

spark.driver.memory

The amount of memory that is used for the driver process when the `SparkContext` parameter is initialized.

The following table describes the tuning recommendations for the Spark parameters according to the size of the data set:

Data Set	Recommended Value
Sandbox	1 GB
Small	2 GB
Medium	2 GB
Large	4 GB

The Spark configuration properties are defined in the `hadoopEnv.properties` file in the following directory:
`<Informatica installation directory>/services/shared/hadoop/<DIST>/infaConf/</>`

Note: You must restart the Data Integration Service after you make changes in the `hadoopEnv.properties` file.

Tune the Amazon Redshift Cluster Properties

You can tune the following additional properties at the mapping level for the Amazon Redshift cluster:

Properties	Sandbox Data Set	Small Data Set	Medium Data Set	Large Data Set
fs.s3a.connection.maximum	Default	3	30	300
Redshift compute node type	dc1.large	dc1.8xlarge	dc1.8xlarge	dc1.8xlarge
Number of Redshift compute nodes	1	2	8	64

Optimizing the Amazon Redshift Mapping

You can configure the following parameters to optimize the Amazon Redshift mapping performance:

Instance size for Informatica Data Integration Service in an Amazon EMR cluster

The Amazon EMR cluster configures each instances with the appropriate Hadoop configuration settings, such as Java memory size, number of mappers, and number of reducers for each instances. When you provision an Amazon EMR cluster instance, you must choose the instance size of your nodes appropriately as some workloads are CPU intensive while others are disk I/O or memory intensive.

Use the following guidelines to choose the right instance size for Informatica Data Integration Service in an Amazon EMR cluster:

- For memory intensive jobs, you must ensure that the `m2` family instance size have enough memory and CPU power to perform the task.
- For CPU intensive jobs, you must ensure that you have a `c2-series` instance.
- If the jobs are both memory and CPU intensive, you must ensure that you use the `cg1.4xlarge` or `cc2.8xlarge` instance as they have suitable memory and CPU power to handle a heavy workload.

Note: Informatica recommend that you use the pre-configured Amazon EMR cluster parameters even though you can change your Hadoop configuration by bootstrapping the Amazon EMR cluster.

Amazon Redshift cluster

When you provision an Amazon Redshift cluster, you must choose the **Redshift Compute Node Type** and **Number of Redshift Compute Nodes** parameters appropriately.

Note: Choose the Amazon Redshift cluster in the same region where the Data Integration Service machine runs.

Disk space

Amazon Redshift recommends that you use the `cc2`, `cg1`, or `D2` series that are ideal for the Hadoop distributed storage and MapR distributed HDD storage. You must have two disks each of 2 TB instead of two disks with 1 TB and 4 TB to reduce the disk I/O, improve write performance, and minimize the downtime.

Informatica recommend that you use multiple disks with smaller capacity instead of a single disk with larger capacity.

CPU clock speed

If your mappings consist of multiple CPU bound transformations in a medium or large data sets, Informatica recommends that you use high clock speed processor.

Network

If you need to shuffle or sort large amount data, you must use the `cg1.4xlarge` and `cc2.8xlarge` instances with 10 Gb of network speed. When the data is in memory, multiple applications are network bound.

To make the **Distributed Reduce** applications such as group-bys, reduce-bys, and SQL joins to perform faster, use 10 Gb or higher network.

Note: You can see how much data shuffles or sorts across the network from the **UI Monitoring Application** .

Troubleshooting Spark Tuning Parameters

Use the following section to troubleshoot errors while tuning the Spark parameters:

ERROR: "Container is running beyond physical memory limits in EMR cluster"

To resolve this error, increase the `spark.executor.memory` parameter to 2 GB for a small, medium, or large data sets.

ERROR: "java.lang.IllegalArgumentException: Required executor memory (6144+614 MB) is above the max threshold (6144 MB) of this cluster! Please check the values of 'yarn.scheduler.maximum-allocation-mb' and/or 'yarn.nodemanager.resource.memory-mb'"

To resolve this error, increase the `yarn.scheduler.maximum-allocation-mb` parameter to 10 GB or decrease the `spark.executor.memory` parameter from 6 GB to 4 GB for a default data set.

Authors

Krishna Prabhakar Devarakonda

Siddiq Hussain

Subhashree Salam