


Set up the DataSync Agent to share to AWS S3

Version 5.0.3.1

Gold

Perspectium DataSync Agents support the replication of data from your app to an Amazon Web Services (AWS) S3 bucket, or an AWS S3 Subscriber Agent. By configuring your AWS S3 Subscriber Agent, data from your app can be replicated and then saved as either **.json** or **.xml** file(s) in your AWS S3 bucket.



Prerequisites

- ⚠ First, you will need to set up the Perspectium [DataSync Agent](#).
- ⚠ You should also [stop running your DataSync Agent](#) before making any Agent configuration changes.

Procedure

To configure your DataSync Agent to run as an AWS S3 Subscriber Agent, follow these steps:

1

Add joda-time and aws-java-sdk libraries

Add the following .jar files to your DataSync Agent's **extlib** directory:

- [joda-time](#) (v2.10.3)
- [aws-java-sdk-core](#) (v1.11.729)
- [aws-java-sdk-kms](#) (v1.11.729)
- [aws-java-sdk-s3](#) (v1.11.729)

While newer versions may work, they have not been tested and it is suggested you use the versions listed above which have been confirmed to work with this release.

2

Access your agent.xml configuration file

Navigate to the directory where you saved your **agent.xml** file when installing your DataSync Agent.

3

Delete database directives

Open your **agent.xml** file in a text editing application and delete the following directives nested within the **<task>** tag:

- **<database_type>**
- **<database_server>**
- **<database_port>**
- **<database_user>**
- **<database_password>**
- **<database_parms>**
- **<database_column_max_size>**
- **<database>**

4

Update the values for <task_name> and <handler>

Locate the **<task_name>** and **<handler>** directives nested within the **<task>** tag and update their values as follows:

Directive	Update value to...
<task_name>	s3_agent_subscribe
<handler>	com.perspectium.replicator.file.S3Subscriber

5

Add AWS directives

Within the <task> tag, nest the following directives:

Directive	Description	Required?
<access_key>	Access Key associated with your AWS account i NOTE: Cannot be used with <use_instance_credentials/> directive	Yes, ONLY when the DataSync Agent is not installed on an EC2 instance.
<secret_access_key>	Secret Access Key associated with your AWS account i NOTE: Cannot be used with <use_instance_credentials/> directive	Yes, ONLY when the DataSync Agent is not installed on an EC2 instance.
<use_instance_credentials/>	Checks the IAM roles on an EC2 instance and uploads to the S3 bucket if the instance has the correct permissions i NOTE: Cannot be used with <access_key> or <secret_access_key> directive	No
<region>	Region that your AWS S3 bucket resides in. i NOTE: If a region is specified for an Agent on an EC2 instance, then the region must match the region of the EC2 instance and the region of the S3 bucket	Yes
<s3_bucket>	Name of your AWS S3 bucket, including subdirectories if desired, to specify where the records will be uploaded e.g. bucketName/folder1/folder2. For example, with <s3_bucket>psp-bucket</s3_bucket> will save records into the psp-bucket S3 bucket. With <s3_bucket>psp-bucket/datasync-agent/tables/\$table</s3_bucket> configured, if an incident record is being processed and uploaded to the AWS S3 bucket, then the record will be saved in the psp-bucket S3 bucket and in the /datasync-agent/tables/incident directory in that bucket, creating the directories datasync-agent, tables and incident automatically. i NOTE: Adding the \$table token indicates this token will be replaced by the table name of the record.	Yes
<file_format>	Format you want to save your data records in, e.g., json or xml.	Yes

6

Save your agent.xml

Save the changes you've made to your **agent.xml** and close the file. An example **agent.xml** configuration for an AWS S3 Subscriber Agent is shown below:

```
<?xml version="1.0" encoding="UTF-8"?>
<config>
  <agent>
    <share />
    <subscribe>
      <task>
        <task_name>s3_agent_subscribe</task_name>
        <message_connection password="password" user="user">https://mesh.perspectium.net</message_connection>
        <instance_connection password="password" user="user">https://myinstance.service-now.com</instance_connection>
        <handler>com.perspectium.replicator.file.S3Subscriber</handler>
        <decryption_key>The cow jumped over the moon</decryption_key>
        <access_key>AccessKey</access_key>
        <secret_access_key>SecretAccessKey</secret_access_key>
        <region>us-west-2</region>
        <s3_bucket>examples3bucket</s3_bucket>
        <file_format>json</file_format>
      </task>
    </subscribe>
    <polling_interval>40</polling_interval>
  </agent>
</config>
```

Files saved in the AWS S3 bucket will be named **<task_name>.<randomized_unique_identifier>.<file_format>**. A randomized unique identifier is used to ensure there are no file naming collisions when saving to the S3 bucket. Using the above configuration example, a file would be named **s3_agent_subscribe.00b470b7-901c-4447-9316-023a265d632f.json**.

NOTE: In this configuration example, your data records will be saved in your AWS S3 bucket as one file. To save each record from your app as an individual file in your AWS S3 bucket, use the following **agent.xml** configuration example as a guide:

```
<?xml version="1.0" encoding="UTF-8"?>
<config>
  <agent>
    <share />
    <subscribe>
      <task>
        <task_name>s3_agent_subscribe</task_name>
        <message_connection password="password" user="user">https://<customer>.perspectium.net</message_connection>
        <instance_connection password="password" user="user">https://<instance>.service-now.com</instance_connection>
        <handler>com.perspectium.replicator.file.S3Subscriber</handler>
        <decryption_key>The cow jumped over the moon</decryption_key>
        <access_key>AccessKey</access_key>
        <secret_access_key>SecretAccessKey</secret_access_key>
        <region>us-west-2</region>
        <s3_bucket>examples3bucket</s3_bucket>
        <file_format>json</file_format>
        <file_prefix>record_</file_prefix>
        <file_suffix>.json</file_suffix>
        <one_record_per_file/>
      </task>
    </subscribe>
    <polling_interval>40</polling_interval>
  </agent>
</config>
```

Saving one record per file supports the following configuration directives:

Directive	Example	Description	Required?
-----------	---------	-------------	-----------

<file_prefix>	<code><file_prefix>record_</file_prefix></code>	A prefix for the file name of each record. If this directive is not specified, "psp.replicator." will be used as the prefix. If you choose to specify a prefix with a \$table token, e.g. <code>psp_\$table_</code> , the token will be replaced by the table name of the record. For example, if you have <code><file_prefix>psp_\$table_</file_prefix></code> configured, and an incident record is being processed, the filename will result to the following <code>psp_incident_5f82dfaf-cf30-4b37-8f02-94248ge7orvi</code> .	No
<file_suffix>	<code><file_suffix>.xml</file_suffix></code>	A suffix for the file name of each record. If this directive is not specified, ".json" will be used as the suffix.	No

In this case, each record will be saved in its own file named **<file_prefix><randomized_unique_identifier><file_suffix>**. Using the above configuration example, a file would be named **record_00b470b7-901c-4447-9316-023a265d632f.xml**.

7

Run your AWS S3 Subscriber DataSync Agent

After configuring your **agent.xml** file to support your AWS S3 Subscriber Agent, [start running your DataSync Agent](#) again.

Similar topics

- [Set up DataSync Agent to share to HP Vertica](#)
- [Database indexes](#)
- [Set up DataSync Agent to share to MySQL](#)
- [DataSync for ServiceNow \(Replicator\)](#)
- [Get started with the DataSync Agent](#)

Contact Perspective Support



US: **1 888 620 8880**

UK: **44 208 068 5953**

support@perspectrum.com