

# Meshlet Configurations for Snowflake

To enhance your DataSync integration for **Snowflake**, you can configure the Snowflake Bulk Load [Meshlet](#) to the configurations below.

**⚠ WARNING:**

- If you will be setting up the [DataSync Agent](#) to also connect to the same database, you will want to configure the Agent to use the same [column cases](#). By default, the Snowflake Bulk Load Meshlet uses **lowercase** for the column names. Also, the [QUOTED\\_IDENTIFIERS\\_IGNORE\\_CASE](#) setting in Snowflake should be kept with its default value of **false**.
- The meshlet uses the [TIMEZONE](#) session parameter to connect to Snowflake in order to save all timestamps in UTC . By default, **glide\_date\_time** field is mapped to the [TIMESTAMP\\_LTZ\(9\)](#) column type. This allows you to query using a local timezone as needed.

**What's on this page?**

- [List of Meshlet Configurations for Snowflake](#)
- [Debugging Logs](#)

To check out the general meshlet configurations, see [General Meshlet Configurations for DataSync](#).

## List of Meshlet Configurations for Snowflake

Directive	Default Value	Description
<b>maxFileSize</b>		<p><b>Required Configuration.</b> This configuration specifies the max size for temporary files created as the meshlet pushes records to Snowflake. <b>10000</b> will be used if input is over <b>10000</b> to prevent possible performance and memory issues. A suggested value is <b>5000</b>.</p> <pre>perspectium:   filesubscriber:     maxFileSize: 5000</pre>
<b>maxString Size</b>	500	<p>Specifies the maximum string size for the entire collection. The maximum value you can use is <b>1000</b>, if input is greater than that value, the input will default to <b>1000</b>.</p> <pre>perspectium:   filesubscriber:     maxStringSize: 500</pre>
<b>customFileName</b>	\$table-\$randomid	<p>Names file with format of table - random id. File names <b>MUST</b> be unique.</p> <pre>perspectium:   filesubscriber:     customFileName: \$table-\$randomid</pre>
<b>fileDirectory</b>	/files	<p>Directory where the locally created files get made. (In respects to where application is running)</p> <pre>perspectium:   filesubscriber:     fileDirectory: /files</pre>

<b>postInterval</b>	2	<p>Minutes to check dead periods. Check every x minutes to compare if the in memory collection is the same as the last x minutes. If so, write records to file and push to Snowflake</p> <pre>perspectium:   snowflake:     postInterval: 2</pre>
<b>deleteFiles</b>	true	<p>Indicates whether you want to keep or delete locally created CSV files. Will not have a noticeable performance hit.</p> <pre>perspectium:   filessubscriber:     deleteFiles: true</pre>
<b>fileCleanerInterval</b>	1	<p>How often (in minutes) the file cleaner job will run to clean up local files created by the meshlet. This job is similar to the ServiceNow Data Cleaner Scheduled Job. For example, value of <b>4</b> will run the job every four minutes. The default value is <b>1</b>.</p> <pre>perspectium:   filessubscriber:     fileCleanerInterval: 4</pre>
<b>fileMaxAge</b>	1	<p>Any csv file in the filesDirectory older than fileMaxAge in minutes, will be automatically deleted. The default value is <b>1</b>.</p> <pre>perspectium:   filessubscriber:     fileMaxAge: 1</pre>
<b>legacyDbViewId</b>	false	<p>Derivative to use the legacy version for database view tables of concatenating GUID values into a sys_id field. If false, meshlet will use the pre-constructed encoded sys_id created by ServiceNow.</p> <pre>perspectium:   filessubscriber:     legacyDbViewId: false</pre> <p><b>NOTE:</b> If using the Share only selected fields option on your bulk share and you don't want the encoded sys_id to be used, you will want to set all three of the legacyDb properties to be true:</p> <pre>perspectium:   filessubscriber:     legacyDbViewId: true     legacyDbViewAlphaOrder: true     legacyDbViewSkipNull: true</pre>

<b>legacyDbViewAlphaOrder</b>	false	<p><b>Supported in Helium 6.1.7 Release and Newer</b></p> <p>By default, the Snowflake meshlet will concatenate GUID values for database view tables into a <code>sys_id</code> field in the order it reads other <code>*_sys_id</code> fields (such as the <code>fc_sys_id</code> and <code>jc_sys_id</code> fields in the <code>flow_context_chunk</code> database view table) which is not guaranteed to be in alphabetical order.</p> <p>With this configuration set to <b>true</b>, the <code>*_sys_id</code> fields will be read and concatenated in alphabetical order i.e. <code>fc_sys_id</code> + <code>jc_sys_id</code> for <code>flow_context_chunk</code>.</p> <p>This feature requires <code>legacyDbViewId</code> be set to <b>true</b>. If <code>legacyDbViewId</code> is not configured to <b>true</b>, this configuration will not be enabled.</p> <pre>perspectium:   filesubscriber:     legacyDbViewAlphaOrder: false</pre> <p>If this configuration is not specified, the meshlet will concatenate GUID values in the order it reads other <code>*_sys_id</code> fields.</p>
<b>legacyDbViewSkipNull</b>	false	<p><b>Supported in Helium 6.1.7 Release and Newer</b></p> <p>By default, the Snowflake meshlet will concatenate all GUID values for database view tables into a <code>sys_id</code> field including GUID values that are <b>null</b> (i.e. have no value), matching how ServiceNow appends <code>sys_id</code> values. For example if the <code>fc_sys_id</code> field has a value of <b>12345</b> and <code>jc_sys_id</code> has a value of <b>null</b>, the <code>sys_id</code> field will be <b>12345null</b> (in ServiceNow this would show as <code>__ENC__ABCDEER=NULL</code> as ServiceNow does encoding on <code>*_sys_id</code> values).</p> <p>With this configuration set to <b>true</b>, <code>*_sys_id</code> fields with values of <b>null</b> will be skipped and not added to the concatenated value. In the above example, the <code>sys_id</code> field would be <b>12345</b>.</p> <p>This feature requires <code>legacyDbViewId</code> be set to <b>true</b>. If <code>legacyDbViewId</code> is not configured to <b>true</b>, this configuration will not be enabled.</p> <pre>perspectium:   filesubscriber:     legacyDbViewSkipNull: false</pre> <p>If this configuration is not specified, the meshlet will concatenate all GUID values including <b>null</b> values.</p>
<b>useReceipts</b>	true	<p>Enable data guarantee.</p> <pre>perspectium:   message:     receipts:       enable: true</pre>
<b>receiptOutboundQueue</b>		<p>Queue receipts are returned to, typically default subscribed to on a ServiceNow instance.</p> <pre>perspectium:   message:     receiptOutboundQueue: psp.out.servicenow.dev1234</pre>
<b>useConnectionProperties</b>	false	<p>When enabled, the user, password, and role properties on the yml are no longer being used to create the database connection.</p> <pre>perspectium:   snowflake:     useConnectionProperties: true</pre> <p>This will allow you to add these or other parameters to your <code>connectionUrl</code> to leverage other forms of authentication. For example:</p> <pre>connectionUrl: jdbc:snowflake://hga12424.us-east-1.snowflakecomputing.com/? USER=xxxx&amp;PASSWORD=xxxx&amp;ROLE=xxxx</pre>

<b>updateTimeOnInsert</b>	true	<p>In conjunction with <a href="#">Add custom fields in Snowflake</a>, if you want the IO update column to be NULL on record inserts, set the value to <b>false</b>.</p> <pre>perspectium:   snowflake:     updateTimeOnInsert: true</pre> <p>By default, updateTimeOnInsert is true, and IO update columns are populated on record inserts.</p>
<b>useLocalSchema</b>	false	<p>Enabling this directive will force the snowflake meshlet to read schemas via files exclusively from the local file storage and disable all calls to ServiceNow. Any new tables shared will require table schemas to be added to the local file storage. See <a href="#">Load Schema Files</a> on how to get started.</p> <pre>perspectium:   auth:     useLocalSchema: true</pre>
<b>loadSchemasOnStart</b>	true	<p><b>Supported in Helium 6.1.11 Release and Newer</b></p> <p>In conjunction with the <b>useLocalSchema</b> configuration above, this configuration controls whether the schemas stored in the local file storage are loaded when the meshlet is started to immediately create the tables in the Snowflake database.</p> <p>By default, this value is <b>true</b> meaning that when the meshlet starts up, it will load all the schemas stored in local file storage and then use them to create the tables in Snowflake when the meshlet is starting up. This will ensure faster performance for when records are received by the meshlet to insert/update into the database table since it doesn't have to create tables at that time.</p> <p>Set this value to <b>false</b> if you want the meshlet to create the table only when it receives a record for the table.</p> <pre>perspectium:   auth:     loadSchemasOnStart: true</pre>
<b>softDelete</b>	false	<p>This feature allows you to add columns to a database table that gets updated when the meshlet receives a message indicating a record is to be deleted. The meshlet will update said columns appropriately and will not delete the record in Snowflake. For more information, see <a href="#">Soft Deletes</a>.</p> <pre>perspectium:   snowflake:     softDelete: true</pre>
<b>quotedColumns</b>	false	<p>Enabling this directive will add quotations around column names in the database.</p> <pre>perspectium:   snowflake:     quotedColumns: true</pre>
<b>translateBooleans</b>	false	<p>To translate boolean values such as <b>true</b> and <b>false</b> to <b>1</b> (true) and <b>0</b> (false). For the case where you've modified the <a href="#">data baseConfig.json</a> so that boolean fields are created as NUMBER fields in Snowflake, this configuration will translate values in boolean fields of the incoming ServiceNow records from <b>true/false</b> to <b>1/0</b> so they can be saved into Snowflake.</p> <p>By default, this configuration is <b>false</b> and boolean values are <b>not translated</b>. To enable this configuration, set it <b>true</b>:</p> <pre>perspectium:   mapping:     translateBooleans: true</pre>

<b>truncateColumns</b>	true	<p><b>Supported in Helium 6.1.4 Release and Newer</b></p> <p>By default, the Snowflake meshlet will create fields without specifying a length (the one exception is <a href="#">timestamps</a>) so that these fields are created with the maximum size i.e. VARCHAR(16777216) to avoid issues with the content being too long for the column.</p> <p>If columns are created with a size where content being inserted is longer than the column's maximum length, by default the meshlet will truncate the text to fit into the column using Snowflake's <a href="#">TRUNCATECOLUMNS</a> option. To change it so that the meshlet does not truncate the column (and instead will throw an error), change the configuration to <b>false</b>:</p> <pre>perspectium:   snowflake:     truncateColumns: false</pre> <p>If this configuration is not specified, the meshlet will treat this configuration as true and truncate content to fit into a column's length.</p>
<b>extendColumnsMaxLength</b>	true	<p><b>Supported in Helium 6.1.5 Release and Newer</b></p> <p>With this configuration, if the meshlet sees content that is greater than the current column's length, it will extend the column to the max size supported for Snowflake i.e. VARCHAR(16777216) for string (VARCHAR) fields.</p> <p>Using this configuration along with the above <b>truncateColumns</b> will allow columns to be extended to their maximum size and then only truncate when the content is longer than the maximum size for the column so that there is no loss of records.</p> <p>For example, with both the <b>extendColumnsMaxLength</b> and <b>truncateColumns</b> configurations set to true, you will see:</p> <ol style="list-style-type: none"> <li>1) The Snowflake meshlet receives a record of 5,000 characters for a table column that is defined as VARCHAR(4000)</li> <li>2) The column is extended to VARCHAR(16777216)</li> <li>3) The record is saved successfully</li> <li>4) A record with 18,000,000 characters is received and is truncated to 16,777,216 characters to fit into this column and save the record into Snowflake</li> </ol> <p>If you prefer that columns not be extended, change the configuration to <b>false</b>. Note that by default all columns are created with their lengths as maximum size as mentioned in the <b>truncateColumns</b> configuration. So this is only applicable if you've configured it or created tables with smaller sizes.</p> <pre>perspectium:   snowflake:     extendColumnsMaxLength: false</pre>

[Go to top of page](#)

## Debugging Logs

To enable debugging logs for the Snowflake Meshlet, you can configure your connection url in the **application.yml** with one the following:

```
jdbc:snowflake://xxxxxx.snowflakecomputing.com/?TRACING=CONFIG
```

```
jdbc:snowflake://xxxxxx.snowflakecomputing.com/?TRACING=FINEST
```

```
jdbc:snowflake://xxxxxx.snowflakecomputing.com/?TRACING=ALL
```

Where **ALL** will provide the most detailed level of additional logging. See [tracing=<string>](#) for more information on the different logging levels available.

A tmp log file will be created. See [How to generate log file on Snowflake connectors](#).

