

# Solace PubSub+ Connector for Snowflake *User Guide*

Solace Corporation

Version 2.0.0

**solace.**

# Table of Contents

|   |    |
|---|----|
| Preface .....   | 1  |
| Getting Started .....   | 2  |
| Prerequisites .....   | 2  |
| Quick Start common steps .....  | 2  |
| Quick Start: Running the connector via command line .....                 | 2  |
| Quick Start: Running the connector via <code>start.sh</code> script ..... | 3  |
| Quick Start: Running the connector as a Container .....                   | 5  |
| Enabling Workflows .....  | 7  |
| Configuring Connection Details .....                                      | 8  |
| Solace PubSub+ Connection Details .....                                   | 8  |
| Preventing Message Loss when Publishing to Topic-to-Queue Mappings .....  | 8  |
| Snowflake Connection Details .....  | 8  |
| Connecting to Multiple Systems .....                                      | 9  |
| User-configured Header Transforms .....                                   | 11 |
| User-configured Payload Transforms .....                                  | 12 |
| Registered Functions .....  | 12 |
| Message Headers .....   | 14 |
| Solace Headers .....  | 14 |
| Reserved Message Headers .....  | 14 |
| Management and Monitoring Connector .....                                 | 15 |
| Monitoring Connector's States .....                                       | 15 |
| Health .....  | 17 |
| Workflow Health .....   | 17 |
| Solace Binder Health .....  | 18 |
| Leader Election .....   | 19 |
| Leader Election Modes: Standalone / Active-Active .....                   | 19 |
| Leader Election Mode: Active-Standby .....                                | 19 |
| Leader Election Management Endpoint .....                                 | 20 |
| Workflow Management .....   | 21 |
| Workflow Management Endpoint .....  | 21 |
| Workflow States .....   | 22 |
| Metrics .....   | 23 |
| Connector Meters .....  | 23 |
| Add a Monitoring System .....   | 24 |
| Security .....  | 25 |
| Securing Endpoints .....  | 25 |
| Exposed Management Web Endpoints .....                                    | 25 |
| Authentication & Authorization .....                                      | 25 |

|  |    |
|--|----|
| TLS.....   | 26 |
| Configuration.....   | 27 |
| Providing Configuration.....   | 27 |
| Converting Canonical Spring Property Names to Environment Variables..... | 27 |
| Spring Profiles.....   | 27 |
| Configure Locations to Find Spring Property Files.....                   | 27 |
| Obtaining Build Information.....   | 28 |
| Spring Configuration Options.....  | 29 |
| Snowflake Binder Configuration Options.....                              | 29 |
| Snowflake Producer Options.....  | 30 |
| Connector Configuration Options.....                                     | 32 |
| Workflow Configuration Options.....                                      | 34 |
| License.....   | 36 |
| Support.....   | 36 |

# Preface

Solace PubSub+ Connector for Snowflake streams data from the Solace PubSub+ Event Broker to Snowflake, providing a flexible and efficient way to load Snowflake data from your Solace-backed, event driven architecture and the Event Mesh. The connector is deployable standalone or in a redundancy mode of "active-standby" to allow for high availability of our data movement. Each connector instance supports up to 20 individual workflows (queue-to-table pipeline), minimizing the number of connector instances deployed and managed. The use of various Spring Framework technologies allows for easy configuration of the connector, advanced logging capabilities, and export of live metrics data to external monitoring solutions.

# Getting Started

Assuming you're using the default `application.yml` within this package, following one of the below quick start guides will result in a connector that will connect to the PubSub+ broker and a Snowflake service instance using specified credentials, with 2 workflows enabled, workflow 0 and workflow 1. Where:

- Workflow 0 is consuming messages from the Solace PubSub+ queue, `Queue-A`, and publishing them to the Snowflake table, `DEMO.PUBLIC.TABLEA`.
- Workflow 1 is consuming messages from the Solace PubSub+ queue, `Queue-B`, and publishing them to the Snowflake table, `DEMO.PUBLIC.TABLEB`.

A workflow is the configuration of a flow of messages from a source to a target. The connector supports up to 20 concurrent workflows per instance.



The connector will not provision Solace queues which do not exist, but can provision the Snowflake tables based on configuration.

## Prerequisites

- [Solace PubSub+ Event Broker](#)
- [Snowflake](#)

## Quick Start common steps

These are the steps that are required to run all quick-start examples:

1. Configure a Snowflake user account with a [public/private key pair](#) for authentication.
2. Grant the necessary permissions to the user account to allow for table creation and data insertion.
3. (Optional) Create the desired target test tables in accordance with the standard connector schema. // TODO: Add link when this is defined
4. Update the provided `samples/config/application.yml` with the credentials and entities for your deployment.

## Quick Start: Running the connector via command line

Run:

```
java -jar pubsubplus-connector-snowflake-2.0.0.jar --add-opens=java.base/java.nio=ALL-UNNAMED --spring.config.additional-location=file:samples/config/
```



By default, this command will detect any Spring Boot config files as per [Spring Boot's default locations](#).

For more info, see [Configure Locations to Find Spring Property Files](#).

## Quick Start: Running the connector via **start.sh** script

For convenience purposes, connector may be also started through the shell script using following syntax:

```
chmod 744 ./bin/start.sh
./bin/start.sh [-n NAME] [-l FOLDER] [-p PROFILE] [-c FOLDER] [-ch HOST] [-cp PORT] [-j FILE] [-cm] [-cmh HOST] [-cmp PORT] [-mh HOST] [-mp PORT] [-o OPTIONS] [-b]
```

In case of invalid parameters:

```
./bin/start.sh -l dummy_folder -c dummy_folder -j dummy_file.jar
```

the script shows you all errors at once:

Solace PubSub+ Connector for Snowflake

Connector startup failed:

```
Following folder doesn't exists on your filesystem: 'dummy_folder'
Following folder doesn't exists on your filesystem: 'dummy_folder'
Following file doesn't exists on your filesystem: 'dummy_file.jar'
```

In case of missing parameters, script will run with predefined values, which are the following:

| Parameter         | Default Value      | Description  |
|-------------------|--------------------|--|
| <b>-n, --name</b> | <b>application</b> | Name of the Connector instance, set up in [spring.application.name]. This name impacts on grouping connectors only   |
| <b>-l, --libs</b> | <b>./libs</b>      | Directory containing required and optional dependency jars. Such as Micrometer metrics export dependencies (if configured). If not specified, it will use the current <b>./libs/</b> folder. |

| Parameter                     | Default Value                            | Description   |
|-------------------------------|--|---|
| <code>-p, --profile</code>    | empty, no profile is used                | Profile to be used with connector's configuration. The configuration file named 'application-<profile>.yml' will be used. Default value is set to use no profile.   |
| <code>-c, --config</code>     | <code>./</code> or current folder        | Path to the folder containing the configuration files to be applied during connector startup for chosen profile.  |
| <code>-H, --host</code>       | 127.0.0.1                                | Provides host where Connector will be running on.   |
| <code>-P, --port</code>       | 8090                                     | Provides port where Connector will be running on.   |
| <code>-mp, --mgmt_port</code> | 9009                                     | Provides management port for back calls of current Connector from Connector Manager. Ignored if <code>-cm</code> argument is missing  |
| <code>-j, --jar</code>        | pubsubplus-connector-snowflake-2.0.0.jar | Path to specified JAR file to start connector. If not specified, the default jar file is used from the current folder.  |
| <code>-cm, --manager</code>   | application                              | Allow to enable cloud configuration for Connector by using Connector Manager configuration storage. Enabling this option also requires to set <code>-mp</code> or <code>--mgmt_port</code> , <code>-H</code> or <code>--host</code> and <code>-cmh</code> with <code>-cmp</code> unless you want to use defaults. Be aware, this option disable listed parameters to be read from configuration file. In this case operator must specify them explicitly as parameters for this script or use default values. |
| <code>-cmh, --cm_host</code>  | 127.0.0.1                                | Specifies the host where Connector Manager is running. Ignored if <code>-cm</code> argument is missing.   |

| Parameter                     | Default Value     | Description  |
|-------------------------------|-------------------|--|
| <code>-cmp, --cm_port</code>  | 9500              | Specifies the port where Connector Manager is running. Ignored if <code>-cm</code> argument is missing.  |
| <code>-o, --options</code>    | no default values | Specifies JVM options used on Connector start.   |
| <code>-tls</code>             | N/A               | Use HTTPS instead of HTTP. Configuration file in this case should contain additional section with preconfigured paths for key store and trust store files. |
| <code>-s, --show</code>       | N/A               | This option prints the start CLI command in raw and exit, doing effectively nothing, just a dry run.   |
| <code>-b, --background</code> | N/A               | Runs Connector in the background. No logs will be displayed and Connector continues running in detached mode   |
| <code>-h, --help</code>       | N/A               | Prints some help information and exits   |

Script also provides that help information from command line using parameter `-h`.

More configuration example of starting Connector together with Connector Manager are provided by the Connector Manager samples.

## Quick Start: Running the connector as a Container

A sample docker compose file has been packaged for your convenience:

1. Change to the `docker` directory:

```
cd samples/docker
```

This directory contains both the `docker-compose.yml` file as well as a `.env` file that contains environment secrets required for the container's health-check.

2. Run the connector:

```
docker-compose up -d
```



This sample docker compose file will:

- Expose the connector's **8090** web port to **8090** on the host.
- Connect to PubSub+ and Snowflake exposed on the host using default ports.
- Mount the **samples/config** directory.
- Mount the previously defined **libs** directory.
- Create a **healthcheck** user with read-only permissions.
  - The default username and password for this user can be found within the **.env** file.
  - This will override any users you have defined in your **application.yml**. See [here](#) for more info.
- Uses the connector's management health endpoint as the container's healthcheck.

For more info about how to use and configure this container, see [the connector's container documentation](#).

# Enabling Workflows

The provided `application.yml` enables workflow 0 and 1. To enable additional workflows, define the following properties in the `application.yml`, where `<workflow-id>` is a value between `[0-19]`:

```
spring:
  cloud:
    stream:
      bindings: # Workflow bindings
        input-<workflow-id>:
          destination: <input-destination> # Queue name
          binder: solace # Input system
        output-<workflow-id>:
          destination: <output-destination> # Table name
          binder: snowflake # Output system

solace:
  connector:
    workflows:
      <workflow-id>:
        enabled: true
```



The connector only supports workflows in the directions of:

- `solace` → `snowflake`

For more info about Spring Cloud Stream and the Solace PubSub+ binder:

- [Spring Cloud Stream Reference Guide](#)
- [Spring Cloud Stream Binder for Solace PubSub+](#)

# Configuring Connection Details

## Solace PubSub+ Connection Details

The Spring Cloud Stream Binder for Solace PubSub+ uses [Spring Boot Auto-Configuration for the Solace Java API](#) to configure its session.

In the `application.yml`, this typically would be configured as follows:

```
solace:
  java:
    host: tcp://localhost:55555
    msg-vpn: default
    client-username: default
    client-password: default
```

For more info and options to configure the PubSub+ session, see [Spring Boot Auto-Configuration for the Solace Java API](#).

## Preventing Message Loss when Publishing to Topic-to-Queue Mappings

If the connector is publishing to a topic that is subscribed to by a queue, messages may be lost if they are rejected (e.g. if queue ingress is shutdown).

To prevent message loss, configure `reject-msg-to-sender-on-discard` with the `including-when-shutdown` flag.

## Snowflake Connection Details

To manually configure Snowflake connection details, set the following in `application.yml`:

```
snowflake-binder:
  url: <locator>.<region>.snowflakecomputing.com:443
  username: snowflake_user
  role: ROLE
  private-key: "
    MIIEvwIBADANBgkqhkiG9w0BAQEFAASCBAkwggSIAgEAAoIBAQMdGWWuig8kvnE2
    HST+o4cGZi0k4vhMjLt3YbA2STw1P992NGh1sKGYVNR6ss4iOUb3B3LXCzeqNYpw
    .....
    lFWgnb6kKLqvC0jmQK+7yWR+mVULyIoRhYwJJtnDIRJyE8cJ9WNLyyAG1ISYfmS0
    8viks1H0zUSgnNii9zmigHLvyA==
  "
```

# Connecting to Multiple Systems

To connect to multiple systems of a same type, use the [multiple binder syntax](#).

For instance:

```
spring:
  cloud:
    stream:

    binders:
      solace1: # 1st solace binder in this example
        type: solace
        environment:
          solace:
            java:
              host: tcp://localhost:55555
      solace2: # 2nd solace binder in this example
        type: solace
        environment:
          solace:
            java:
              host: tcp://other-host:55555
      snowflake1: # The only snowflake binder
        type: snowflake
        environment:
          snowflake-binder:
            # Add connection and binder properties here to customize behavior.
# Required for internal use
undefined:
  type: undefined

  bindings:
    input-0:
      destination: <input-destination>
      binder: solace1 # Reference 1st solace binder
    output-0:
      destination: <output-destination>
      binder: snowflake1
    input-1:
      destination: <input-destination>
      binder: solace2 # Reference 2nd solace binder
    output-1:
      destination: <output-destination>
      binder: snowflake1
```

Defines two binders of type `solace` and one binder of type `snowflake` which are then referenced within bindings.

Note that each binder is configured independently under `spring.cloud.stream.binders.<binder-name>.environment`.



When connecting to multiple systems, all binder configuration must be specified using the multiple binder syntax for all binders.

Do not use single-binder configuration (e.g. `solace.java.*` at the root of your `application.yml`) while using the multiple binder syntax.

# User-configured Header Transforms

Generally, the consumed message's headers are propagated through the connector to the output message. If you want to transform the headers, then you may do so as follows:

```
# <workflow-id> : The workflow ID ([0-19])
# <header> : The key for the outbound header
# <expression> : A SpEL expression which has "headers" as parameters

solace.connector.workflows.<workflow-id>.transform-headers.expressions.<header>=<expression>
```

**Example 1:** To create a new header, `new_header`, for workflow `0` that is derived from the headers `foo` & `bar`:

```
solace.connector.workflows.0.transform-headers.expressions.new_header
="T(String).format('%s/abc/%s', headers.foo, headers.bar)"
```

**Example 2:** To remove the header, `delete_me`, for workflow `0`, set the header transform expression to `null`:

```
solace.connector.workflows.0.transform-headers.expressions.delete_me="null"
```

For more info about Spring Expression Language (SpEL) expressions:

- [Spring Expression Language \(SpEL\)](#)

# User-configured Payload Transforms

Message payloads going through a workflow can be transformed using a Spring Expression Language (SpEL) expression as follows:

```
# <workflow-id> : The workflow ID ([0-19])
# <expression> : A SpEL expression

solace.connector.workflows.<workflow-id>.transform-payloads.expressions[0].transform
=<expression>
```

A SpEL expression may reference:

- `payload`: to access the message payload
- `headers.<header_name>`: to access a message header value
- Registered functions



While the syntax uses an array of expressions, only a single transform expression is supported in this release. Multiple transform expressions may be supported in future versions.

## Registered Functions

[Registered functions](#) are built-in and can be called directly from SpEL expressions. To call a registered function, use the `#` character followed by the function name. The following table describes available registered functions:

| Registered Function Signature                   | Description   |
|---|---|
| <code>boolean isPayloadBytes(Object obj)</code> | <p>Returns whether the object <code>obj</code> is an instanceof <code>byte[]</code> or not.</p> <p>Sample usage of this function within a SpEL expression: <code>"#isPayloadBytes(payload) ? true : false"</code></p> |

**Example 1:** To normalize `byte[]` and String payloads as upper-cased String payloads or leave payloads unchanged when of different types.

```
solace.connector.workflows.0.transform-payloads.expressions[0].transform
="#isPayloadBytes(payload) ? new String(payload).toUpperCase() : payload instanceof
T(String) ? payload.toUpperCase() : payload"
```

**Example 2:** To convert String payloads to `byte[]` payloads using a charset retrieved from a message header or leave payloads unchanged when of different types.

```
solace.connector.workflows.0.transform-payloads.expressions[0].transform="payload  
instanceof T(String) ?  
payload.getBytes(T(java.nio.charset.Charset).forName(headers.charset)) : payload"
```

For more info about Spring Expression Language (SpEL) expressions:

- [Spring Expression Language \(SpEL\)](#)



# Message Headers

Solace and JMS headers can be created or manipulated using the [User-configured Header Transforms](#) feature described above.

## Solace Headers

Solace headers exposed to the connector are documented within the [Spring Cloud Stream Binder for Solace PubSub+](#) documentation.

## Reserved Message Headers

The following are reserved header spaces:

- `solace_`
- `scst_`
- Any headers defined by the core Spring messaging framework. See [Spring Integration: Message Headers](#) for more info.

Any headers with these prefixes, that are not defined by the connector or any technology used by the connector, may not be backwards compatible in a future version of the connector.

# Management and Monitoring Connector

## Monitoring Connector's States

The Connector provides an ability to monitor its internal states through exposed endpoints provided by [Spring Boot Actuator](#).

Actuator shares information through the endpoints reachable over HTTP(s). What endpoints are available is configured in the connector configuration file:

```
management:
  simple:
    metrics:
      export:
        enabled: true
    endpoints:
      web:
        exposure:
          include:
            "health,metrics,loggers,logfile,channels,env,workflows,leaderelection,bindings,info"
```

The above example configuration enables metrics collection through the configuration parameter of `management.simple.metrics.export.enabled` set to `true` and then shares them through the HTTP(s) endpoint together with other sections configured for the current Connector.

The set of endpoints exposed through the HTTP(s) endpoint. Exposed endpoints are available in the connector UI and are also available to the PubSub+ Connector Manager. The operator may choose to not expose all or some of these endpoints. Endpoints not exposed will not be available in the connector web UI nor the PubSub+ Connector Manager.



The simple metrics registry is only to be used for testing. It is not a production-ready means of collecting metrics. In production, use a dedicated monitoring system (e.g. Datadog, Prometheus, etc) to collect metrics.

The Actuator endpoint now contains information about Connector's internal states shared over the following HTTP(s) endpoint:

```
GET: /actuator/
```

Here is the example of the data shared with the configuration above:

```
{
  "_links": {
    "self": {
      "href": "/actuator",
      "templated": false
    }
  }
}
```

```

},
"workflows": {
  "href": "/actuator/workflows",
  "templated": false
},
"workflows-workflowId": {
  "href": "/actuator/workflows/{workflowId}",
  "templated": true
},
"leaderelection": {
  "href": "/actuator/leaderelection",
  "templated": false
},
"health-path": {
  "href": "/actuator/health/{*path}",
  "templated": true
},
"health": {
  "href": "/actuator/health",
  "templated": false
},
"metrics": {
  "href": "/actuator/metrics",
  "templated": false
},
"metrics-requiredMetricName": {
  "href": "/actuator/metrics/{requiredMetricName}",
  "templated": true
}
}
}

```

# Health

The connector reports its health status via the [Spring Boot Actuator health endpoint](#).

To configure the information returned by the `health` endpoint, configure the following properties: `management.endpoint.health.show-details` and `management.endpoint.health.show-components`. Refer to [Spring Boot documentation](#) for details.

Health for the workflow, Solace binder, and snowflake binder components are exposed when `management.endpoint.health.show-components` is enabled. For example:

```
management:
  endpoint:
    health:
      show-components: always
      show-details: always
```

This would always show the full detail of the health check including the workflows and binders. The default value is `never`.

## Workflow Health

A `workflows` health indicator is provided to show the health status for each of a connector's workflows. This health indicator has the following form:

```
{
  "status": "(UP|DOWN)",
  "components": {
    "<workflow-id>": {
      "status": "(UP|DOWN)",
      "details": {
        "error": "<error message>"
      }
    }
  }
}
```

| Health Status | Description  |
|---------------|--|
| UP            | Status indicating that the workflow is functioning as expected.                      |
| DOWN          | Status indicating that the workflow is unhealthy. User intervention may be required. |

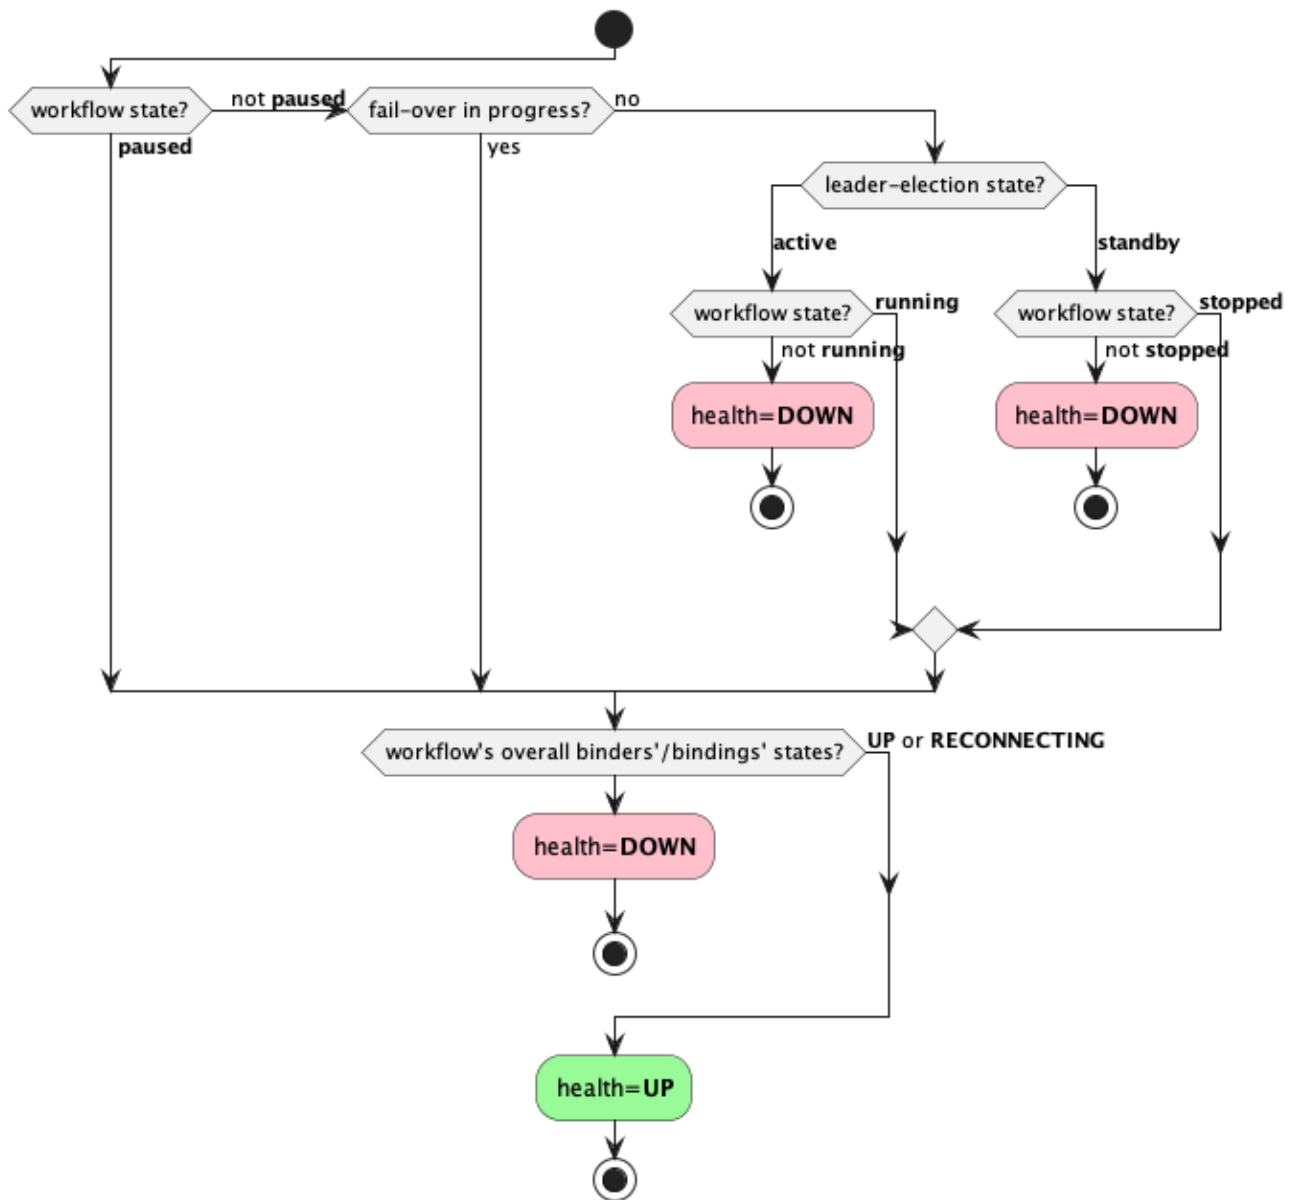


Figure 1. Workflow Health Resolution Diagram

This health indicator is enabled default. To disable it, add this to your configuration:

```
management.health.workflows.enabled=false
```

## Solace Binder Health

For details, see the [Solace binder](#) documentation.

# Leader Election

The connector has 3 leader election modes:

| Leader Election Mode | Description  |
|----------------------|--|
| Standalone (Default) | A single instance of a connector without any leader election capabilities.   |
| Active-Active        | A participant in a cluster of connector instances where all instances are active.  |
| Active-Standby       | A participant in a cluster of connector instances where only one instance is active (i.e. the leader), and the others are standby. |

Operators can configure the leader election mode by setting:

```
solace.connector.management.leader-election.mode
=(standalone|active_active|active_standby)
```

## Leader Election Modes: Standalone / Active-Active

All enabled workflows are started during connector startup and the connector is considered as always active.

## Leader Election Mode: Active-Standby

If the connector is in active-standby mode, then a PubSub+ management session and management queue must be configured as follows:

```
solace.connector.leader-election.mode=active_standby

# Management session
# Exact same interface as solace.java.*
solace.connector.management.session.host=<management-host>
solace.connector.management.session.msgVpn=<management-vpn>
solace.connector.management.session.client-username=<client-username>
solace.connector.management.session.client-password=<client-password>
solace.connector.management.session.<other-property-name>=<value>

# Management queue name accessible by the management session
# Must have exclusive access type
solace.connector.management.queue=<management-queue-name>
```

To determine if the connector is **active** or **standby**, it will create a flow to the management queue. If this flow is active, then the connector's state is **active** and will start its enabled workflows. Otherwise, if this flow is inactive, then the connector's state is **standby** and will stop its enabled workflows.

At a macro level for a cluster of connectors, fail-over only happens when there are infrastructure failures (e.g. JVM goes down or networking failures to the management queue).

If a workflow fails to be started/stopped during fail-over, it will retry up to some maximum defined by the config option, `solace.connector.management.leader-election.fail-over.max-attempts`.

During fail-over, the connector will attempt to start/stop all enabled workflows. Once an attempt has been made to start/stop each workflow, then the connector has transitioned to active/standby mode regardless of the status of the workflows.

## Leader Election Management Endpoint

A custom `leaderelection` management endpoint was provided using [Spring Actuator](#).

Navigate to the connector's `leaderelection` management endpoint to view its leader election status.

| Endpoint                     | Operation                        | Payloads  |
|------------------------------|----------------------------------|---|
| <code>/leaderelection</code> | Read<br>(HTTP <code>GET</code> ) | <p><b>Request:</b> None.</p> <p><b>Response:</b></p> <pre> {   "mode": {     "type": "(standalone                 active_active   ①               active_standby)",     "state": "(active   standby)", ②     "source": { ③       "queue": "&lt;management-queue-name&gt;",       "host": "&lt;management-host&gt;",       "msgVpn": "&lt;management-vpn&gt;"     }   } } </pre> <p>① Mandatory parameter in output</p> <p>② Mandatory parameter in output</p> <p>③ Optional section. Appears only when <code>type</code> is set to <code>active_standby</code>.</p> |

# Workflow Management

## Workflow Management Endpoint

A custom `workflows` management endpoint using `Spring Actuator` was provided to manage workflows.

To enable the `workflows` management endpoint:

```
management:
  endpoints:
    web:
      exposure:
        include: "workflows"
```

Once the `workflows` management endpoint is enabled, the following operations can be performed:

| Endpoint                             | Operation                        | Payloads  |
|--------------------------------------|----------------------------------|---|
| <code>/workflows</code>              | Read<br>(HTTP <code>GET</code> ) | <b>Request:</b> None.<br><br><b>Response:</b><br>Same payload as the <code>/workflows/{workflowId}</code> read operation, but as a list of all workflows.   |
| <code>/workflows/{workflowId}</code> | Read<br>(HTTP <code>GET</code> ) | <b>Request:</b> None.<br><br><b>Response:</b> <pre>{   "id": "&lt;workflowId&gt;",   "enabled": (true false),   "state": "(running stopped paused unknown)",   "inputBindings": [     "&lt;input-binding&gt;"   ],   "outputBindings": [     "&lt;output-binding&gt;"   ] }</pre> |



| Endpoint                | Operation            | Payloads  |
|-------------------------|----------------------|---|
| /workflows/{workflowId} | Write<br>(HTTP POST) | <b>Request:</b> <pre>{   "state": "STARTED STOPPED PAUSED RESUMED" }</pre> <b>Response:</b> None. |



Only workflows with Solace PubSub+ consumers (where the **solace** binder is defined in the **input-#**) support pause/resume.



Some features require for the connector to manage workflow lifecycles. There's no guarantee that workflow states will persist when write operations are used to change workflow states while such features are in use.

For example: When the connector is configured in the active-standby leader election mode, workflows will automatically transition from **running** to **stopped** when the connector fails over from **active** to **standby**. Vice versa for a fail-over in the opposite direction.

## Workflow States

A workflow's state is defined as the aggregate states of its bindings (see the [bindings management endpoint](#)) as follows:

| Workflow State | Condition   |
|----------------|---|
| <b>running</b> | All bindings have <b>state="running"</b> .  |
| <b>stopped</b> | All bindings have <b>state="stopped"</b> .  |
| <b>paused</b>  | All consumer bindings and all pausable producer bindings have <b>state="paused"</b> . |
| <b>unknown</b> | None of the other states. Represents an inconsistent aggregate binding state.         |



When producer/consumer binding is not implementing Spring's Lifecycle interface, Spring always reports binding's **state=N/A**. The **state=N/A** will be ignored in deciding the overall state of the workflow. So for example, if the consumer binding **state=running** and producer binding **state=N/A** or vice a versa the workflow state will be **running**.

For more info about binding states, see [Spring Cloud Stream: Binding visualization and control](#).

# Metrics

This connector uses [Spring Boot Metrics](#) which leverages Micrometer to manage its metrics.

## Connector Meters

In addition to the meters already provided by the Spring framework, this connector introduces the following custom meters:

| Name   | Type  | Tags   | Description           | Notes  |
|--|---|--|-----------------------|--|
| <code>solace.connector.process</code>              | Timer                                       | type: channel<br>name: <bindingName><br>result:<br>(success failure)<br>exception:<br>(none exception<br>simple class name)  | Processing time       | This meter is a rename of <code>spring.integration.send</code> whose <code>name</code> tag matches a binding name.   |
| <code>solace.connector.error.process</code>        | Timer                                       | type: channel<br>name: <bindingNames><br>result:<br>(success failure)<br>exception:<br>(none exception<br>simple class name) | Error processing time | This meter is a rename of <code>spring.integration.send</code> whose <code>name</code> tag matches an input binding's error channel name ( <code>&lt;destination&gt;.&lt;group&gt;.errors</code> ).<br>Meters might be merged under the same <code>name</code> tag (delimited by <code> </code> ) if multiple bindings have the same error channel name (i.e. bindings have matching <code>destination</code> and/or <code>group</code> ).<br><b>Though a reminder that setting a binding's <code>group</code> is not supported.</b> |
| <code>solace.connector.message.size.payload</code> | DistributionSummary<br>Base Units:<br>bytes | name: <bindingName>  | Message payload size  |  |

| Name   | Type   | Tags  | Description                  | Notes |
|--|--|---|------------------------------|-------|
| <code>solace.connector.message.size.total</code> | DistributionSummary<br><br>Base Units: bytes | name: <bindingName>   | Total message size           |       |
| <code>solace.connector.publish.ack</code>        | Counter<br><br>Base Units: acknowledgments   | name: <bindingName><br><br>result: (success failure)<br><br>exception: (none exception simple class name) | Publish acknowledgment count |       |



The `solace.connector.process` meter with `result=failure` is not a reliable measure of tracking the number of failed messages. It only tells you how many times a step processed a message, how long it took to process that message, and if that step completed successfully.

Instead, it's recommended to use a combination of `solace.connector.error.process` and `solace.connector.publish.ack` to track failed messages.

## Add a Monitoring System

By default, this connector only includes JMX as its supported monitoring system.

To add additional monitoring systems, add the system's `micrometer-registry-<system>` jar and its dependency jars to the connector's classpath. The included systems can then be individually enabled/disabled by setting `management.<system>.metrics.export.enabled=true` in the `application.yml`.

# Security

## Securing Endpoints

### Exposed Management Web Endpoints

Please refer to the [Management and Monitoring Connector](#) section for a comprehensive list of endpoints that are automatically enabled for this connector.

The **health** endpoint only returns the root status by default (i.e. no health details).

To enable other management endpoints, see [Spring Actuator Endpoints](#).

### Authentication & Authorization

For this release, the connector only supports basic HTTP authentication.

By default, no users will be created unless the operator configures it in their config. Configuration parameters responsible for security are:

```
solace:
  connector:
    security:
      enabled: true
      users:
        - name: user1
          password: pass
        - name: admin1
          password: admin
      roles:
        - admin
```

In the above example, we have created 2 users:

- **user1**: Has access to perform GET (Read) requests.
- **admin1**: Has access to perform GET and POST (Read & Write) requests.

To fully disable security and permit anyone to access the connector's web endpoints, operators can configure the parameter `solace.connector.security.enabled` switched to `false`.



While these properties could be defined in an `application.yml` file, we recommend that you use environment variables to set secret values.

Here is an example of how to define users using environment variables:

```
# Create user with no role (i.e. read-only)
SOLACE_CONNECTOR_SECURITY_USERS_0_NAME=user1
```

```
SOLACE_CONNECTOR_SECURITY_USERS_0_PASSWORD=pass

# Create user with admin role
SOLACE_CONNECTOR_SECURITY_USERS_1_NAME=admin1
SOLACE_CONNECTOR_SECURITY_USERS_1_PASSWORD=admin
SOLACE_CONNECTOR_SECURITY_USERS_1_ROLES_0=admin
```

In the above example, we have created 2 users:

- **user1**: Has access to perform GET (Read) requests.
- **admin1**: Has access to perform GET and POST (Read & Write) requests.



`solace.connector.security.users` is a list. When users are defined in multiple sources (different `application.yml` files, environment variables, etc), then overriding works by replacing the entire list. Meaning that you must pick one place to defined your users and define all of them there. Whether it be in a **single** application properties file or all of them in the environment variables.

See [Spring Boot - Merging Complex Types](#) for more info.

## TLS

TLS is disabled by default.

To configure TLS, see [Spring Boot - Configure SSL](#) and [TLS Setup in Spring](#).

# Configuration

## Providing Configuration

See [Spring Boot: Externalized Configuration](#) for info about how the connector will detect configuration properties.

## Converting Canonical Spring Property Names to Environment Variables

See the [Spring documentation](#) for how to provide configuration options as environment variables.

## Spring Profiles

If multiple config files will exist within the same config folder for use in different environments (dev, prod, etc), then use Spring profiles.

This will allow you to define different application property files under the same directory using the file name format, `application-{profile}.yml`.

eg:

- `application.yml`
  - Properties in non-specific files always applies. It's properties are overridden by those defined in profile-specific files.
- `application-dev.yml`
  - Defines properties specific to the `dev` environment.
- `application-prod.yml`
  - Defines properties specific to the `prod` environment.

Individual profiles can then be enabled by setting the `spring.profiles.active` property.

See [Spring Boot: Profile-Specific Files](#) for more info as well as an example.

## Configure Locations to Find Spring Property Files

By default, the connector will detect any Spring property files as per [Spring Boot's default locations](#).

- If you want to add additional locations, add `--spring.config.additional-location=file:<custom-config-dir>` (similar to the example command in [Quick Start: Running the connector via command line](#)).
- If you want to exclusively use the locations that you've defined and ignore Spring Boot's default locations, add `--spring.config.location=optional:classpath:/,optional:classpath:/config/,file:<custom-config-dir>`.

See [Spring Boot documentation](#) for more info.

If config files for multiple, different, connectors will exist within the same config folder for use in different environments (e.g. dev, prod, etc), then consider using [Spring Boot Profiles](#) instead of child directories to do this.

i.e.:



- Do this:
  - `config/application-prod.yml`
  - `config/application-dev.yml`
- Instead of this:
  - `config/prod/application.yml`
  - `config/dev/application.yml`

Child directories are intended to be used for merging configuration from multiple sources of config properties. For more information and an example of when you might want to use multiple child directories to compose your application's configuration, please see the [Spring Boot documentation](#).

## Obtaining Build Information

Build information, including version, build date, time and description is enabled by default via [Spring Boot Actuator Info Endpoint](#). By default, every connector shares all information related to its **build** only.

Below is the structure of the output data:

```
{
  "build": {
    "version": "<connector version>",
    "artifact": "<connector artifact>",
    "name": "<connector name>",
    "time": "<connector build time>",
    "group": "<connector group>",
    "description": "<connector description>",
    "support": "<support information>"
  }
}
```

If you wish to exclude build data from the output of the info endpoint, you can do so by setting `management.info.build.enabled` to `false`.

Alternatively, if you want to disable the info endpoint entirely, you should remove 'info' from the list of endpoints specified in `management.endpoints.web.exposure.include`.

## Spring Configuration Options

This connector packages a lot of libraries to customize functionality. Here are some references to get started:

- [Spring Cloud Stream](#)
- [Spring Cloud Stream Binder for Solace PubSub+](#)
- [Spring Logging](#)
- [Spring Actuator Endpoints](#)
- [Spring Metrics](#)

## Snowflake Binder Configuration Options

The following properties are available at the binder level and are complementary to the properties described in the [Snowflake Connection Details](#) section.

These properties are to be prefixed with `snowflake-binder`.

| Config Option            | Type                | Valid Values     | Default Value   | Description   |
|--------------------------|---------------------|------------------|-----------------|---|
| <code>url</code>         | <code>string</code> | <code>any</code> | <code>""</code> | Snowflake account URL usually in the form <code>&lt;locator&gt;.&lt;region&gt;.snowflakecomputing.com:443</code> .  |
| <code>username</code>    | <code>string</code> | <code>any</code> | <code>""</code> | User login name for the Snowflake account.  |
| <code>private-key</code> | <code>string</code> | <code>any</code> | <code>""</code> | The private key to authenticate the user. Include only the key, not the header or footer. Currently only supports unencrypted keys.<br><br>See <a href="#">Using Key Pair Authentication &amp; Key Pair Rotation</a> for more information.  |
| <code>role</code>        | <code>string</code> | <code>any</code> | <code>""</code> | The role the user should take while performing ingestion. Note, that while this property is used for most Snowflake actions, triggering a classic, non-streaming Snowpipe requires the user account to have a default role with an assigned <code>OPERATE</code> privilege or <code>OWNERSHIP</code> privilege on the specified pipe. |



| Config Option    | Type   | Valid Values                    | Default Value      | Description  |
|------------------|--------|---------------------------------|--------------------|--|
| warehouse        | string | any                             | ""                 | The warehouse to use when performing non-streaming operations such as object provisioning or standard Snowpipe ingestion. Can be omitted when using <code>SNOWPIPE_STREAMING</code> (default). |
| ingestion-method | string | (SNOWPIPE   SNOWPIPE_STREAMING) | SNOWPIPE_STREAMING | Specifies whether to use Snowpipe Streaming or standard Snowpipe to load data from Solace.   |

## Snowflake Producer Options

The following config options are available to Snowflake producers.

To apply these options to all output bindings, they should be prefixed with:

- `spring.cloud.stream.snowflake.default.producer.`

To apply these options to a specific binding, they should be prefixed with:

- `spring.cloud.stream.snowflake.bindings.<bindingName>.producer.`

| Config Option | Type   | Valid Values | Default Value | Description  |
|---------------|--------|--------------|---------------|--|
| name          | string | any          | ""            | Application name. This must be unique across all Solace workflows used by the customer. This name must be a valid Snowflake unquoted identifier.   |
| database      | string | any          | ""            | The name of the database that contains the table to insert rows into. This can also be specified in the binding <code>destination</code> when using a fully-qualified table name in the form <code>&lt;database&gt;.&lt;schema&gt;.&lt;table&gt;</code> .  |
| schema        | string | any          | ""            | The name of the schema that contains the table to insert rows into. This can also be specified in the binding <code>destination</code> when using a fully-qualified table name in the form <code>&lt;database&gt;.&lt;schema&gt;.&lt;table&gt;</code> or <code>&lt;schema&gt;.&lt;table&gt;</code> . |

| Config Option                         | Type                | Valid Values                         | Default Value          | Description   |
|---------------------------------------|---------------------|--------------------------------------|------------------------|---|
| <code>table</code>                    | <code>string</code> | <code>any</code>                     | <code>""</code>        | The name of the table to insert rows into. This can also be specified in the binding <code>destination</code> when using a fully-qualified table name in the form <code>&lt;database&gt;.&lt;schema&gt;.&lt;table&gt;</code> , <code>&lt;schema&gt;.&lt;table&gt;</code> , or <code>&lt;table&gt;</code> . Specifying this as a default binding setting allows multiple workflows to share the same target table without explicitly declaring the target table.   |
| <code>batch-max-size</code>           | <code>int</code>    | <code>&gt; 0</code>                  | <code>2000</code>      | Number of messages to receive before flushing the connector buffer into the Snowflake channel or stage.   |
| <code>batch-timeout</code>            | <code>int</code>    | <code>&gt; 0</code>                  | <code>2000</code>      | Number of milliseconds to wait for a full buffer before flushing into the Snowflake channel or stage.   |
| <code>acknowledgment-mode</code>      | <code>string</code> | <code>(ON_COMMIT ON_INSERT)</code>   | <code>ON_COMMIT</code> | Controls when the connector will send acknowledgments to the Solace broker. If configured for <code>ON_COMMIT</code> the connector will wait until a row has been successfully committed to the target table before sending a message acknowledgment. If configured for <code>ON_INSERT</code> the acknowledgment will be sent as soon as the message has been written to the Snowflake channel or stage.   |
| <code>provisioning-mode.tables</code> | <code>string</code> | <code>(NONE CREATE OVERWRITE)</code> | <code>CREATE</code>    | Specifies if and how the connector should provision the specified target table. If <code>NONE</code> is specified, the connector will not attempt to create the table. If <code>CREATE</code> is specified, the table will be created if it does not exist. If <code>OVERWRITE</code> is specified the table will be created always.<br><br><b>Warning:</b> Specifying <code>OVERWRITE</code> could potentially result in the loss of data and is intended primarily for test and development purposes. |

| Config Option                         | Type   | Valid Values            | Default Value | Description   |
|---------------------------------------|--------|-------------------------|---------------|---|
| <code>provisioning-mode.stages</code> | string | (NONE CREATE OVERWRITE) | CREATE        | <p>Specifies if and how the connector should provision the underlying stage when using standard <b>SNOWPIPE</b> ingestion mode. If <b>NONE</b> is specified, the connector will not attempt to create the stage. If <b>CREATE</b> is specified, the stage will be created if it does not exist. If <b>OVERWRITE</b> is specified the stage will be created always.</p> <p><b>Warning:</b> Specifying <b>OVERWRITE</b> could potentially result in the loss of data and is intended primarily for test and development purposes.</p> |
| <code>provisioning-mode.pipes</code>  | string | (NONE CREATE OVERWRITE) | CREATE        | <p>Specifies if and how the connector should provision underlying pipe when using standard <b>SNOWPIPE</b> ingestion mode. If <b>NONE</b> is specified, the connector will not attempt to create the stage. If <b>CREATE</b> is specified, the stage will be created if it does not exist. If <b>OVERWRITE</b> is specified the stage will be created always.</p> <p><b>Warning:</b> Specifying <b>OVERWRITE</b> could potentially result in the loss of data and is intended primarily for test and development purposes.</p>      |

## Connector Configuration Options

These configuration options are all prefixed by `solace.connector.:`

| Config Option   | Type | Valid Values | Default Value | Description  |
|---|------|--------------|---------------|--|
| <code>management.leader-election.fail-over.max-attempts</code>              | int  | > 0          | 3             | The maximum number of attempts to perform a fail-over.                     |
| <code>management.leader-election.fail-over.back-off-initial-interval</code> | long | > 0          | 1000          | The initial interval (milliseconds) to back-off when retrying a fail-over. |
| <code>management.leader-election.fail-over.back-off-max-interval</code>     | long | > 0          | 10000         | The maximum interval (milliseconds) to back-off when retrying a fail-over. |

| Config Option   | Type    | Valid Values   | Default Value           | Description   |
|---|---------|--|-------------------------|---|
| <code>management.leader-election.fail-over.back-off-multiplier</code> | double  | <code>&gt;= 1.0</code>   | <code>2.0</code>        | The multiplier to apply to the back-off interval between each retry of a fail-over.   |
| <code>management.leader-election.mode</code>                          | enum    | <code>(standalone active_active active_standby)</code>                     | <code>standalone</code> | <p>The connector's leader election mode.</p> <p><b>standalone:</b><br/>A single instance of a connector without any leader election capabilities.</p> <p><b>active_active:</b><br/>A participant in a cluster of connector instances where all instances are active.</p> <p><b>active_standby:</b><br/>A participant in a cluster of connector instances where only one instance is active (i.e. the leader), and the others are standby.</p> |
| <code>management.queue</code>   | string  | any  | <code>null</code>       | The management queue name.  |
| <code>management.session.*</code>                                     |         | See <a href="#">Spring Boot Auto-Configuration for the Solace Java API</a> |                         | <p>Defines the management session. This has the same interface as that used by <code>solace.java.*</code>.</p> <p>See <a href="#">Spring Boot Auto-Configuration for the Solace Java API</a> for more info.</p>   |
| <code>security.enabled</code>   | boolean | <code>(true false)</code>  | <code>true</code>       | If <code>true</code> , security is enabled. Otherwise, anyone has access to the connector's endpoints.  |
| <code>security.users[&lt;index&gt;].name</code>                       | string  | any  | <code>null</code>       | The name of this user.  |
| <code>security.users[&lt;index&gt;].password</code>                   | string  | any  | <code>null</code>       | The password of this user.  |

| Config Option                                    | Type                            | Valid Values       | Default Value               | Description  |
|--|---------------------------------|--------------------|-----------------------------|--|
| <code>security.users[&lt;index&gt;].roles</code> | <code>list&lt;string&gt;</code> | <code>admin</code> | empty list (i.e. read-only) | The list of roles which this user has. Has read-only access if no roles are given. |

## Workflow Configuration Options

These configuration options are defined under the prefix, `solace.connector.workflows.<workflow-id>`. (if they support per-workflow config), and the default prefix, `solace.connector.default.workflow`. (if they support default workflow config).

| Config Option                              | Applicable Scopes       | Type                                   | Valid Values  | Default Value      | Description  |
|--|-------------------------|--|---|--------------------|--|
| <code>enabled</code>                       | Per-Workflow            | <code>boolean</code>                   | <code>(true false)</code>   | <code>false</code> | If <code>true</code> , the workflow is enabled.  |
| <code>transform-headers.expressions</code> | Per-Workflow<br>Default | <code>Map&lt;string, string&gt;</code> | <b>Key:</b><br>A header name.<br><br><b>Value:</b><br>A SpEL string which accepts <code>headers</code> as parameters. | empty map          | A mapping of header names to header value SpEL expressions.<br><br>The SpEL context contains the <code>headers</code> parameter which can be used to read the input message's headers.   |
| <code>acknowledgment.publish-async</code>  | Per-Workflow<br>Default | <code>boolean</code>                   | <code>(true false)</code>   | <code>false</code> | If <code>true</code> , publisher acknowledgment processing is done asynchronously.<br><br>The workflow's consumer and producer bindings must support this mode, otherwise, publisher acknowledgments are processed synchronously regardless of this setting. |

| Config Option                                       | Applicable Scopes       | Type | Valid Values | Default Value | Description   |
|---|-------------------------|------|--------------|---------------|---|
| <code>acknowledgment.back-pressure-threshold</code> | Per-Workflow<br>Default | int  | $\geq 1$     | 255           | The maximum number of outstanding messages with unresolved acknowledgments. Message consumption is paused when the threshold is reached to allow for producer acknowledgments to catch up.                              |
| <code>acknowledgment.publish-timeout</code>         | Per-Workflow<br>Default | int  | $\geq -1$    | 600000        | Maximum amount of time, in millisecond, to wait for asynchronous publisher acknowledgments before considering a message as failed. A value of <code>-1</code> means to wait indefinitely for publisher acknowledgments. |

# License

This project is licensed under the Solace Community License, Version 1.0. - See the [LICENSE](#) file for details.

## Support

Support is offered best effort via our [Solace Developer Community](#).

Premium support options are available, please [Contact Solace](#).