



Solace PubSub+ Connector for GCP Pub/Sub *User Guide*

Solace Corporation

Version 1.1.0

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
GCP Pub/Sub Connection Details	8
Using <code>gcloud cli</code> generated user credentials	9
Using Service Account Key	10
Connecting to Multiple Systems	10
User-configured Header Transforms	12
User-configured Payload Transforms	13
Registered Functions	13
Message Headers	15
Solace Headers	15
Message Headers Mapping	15
Reserved Message Headers	16
Dynamic Producer Destinations	17
Asynchronous Publishing	18
Management and Monitoring Connector	19
Monitoring Connector's States	19
Health	21
Workflow Health	21
Solace Binder Health	22
GCP Pub/Sub Binder Health	22
Leader Election	24
Leader Election Modes: Standalone / Active-Active	24
Leader Election Mode: Active-Standby	24
Leader Election Management Endpoint	25
Workflow Management	26
Workflow Management Endpoint	26
Workflow States	27
Metrics	28

Connector Meters	28
Add a Monitoring System	29
Security	30
Securing Endpoints	30
Exposed Management Web Endpoints	30
Authentication & Authorization	30
CSRF Protection	31
TLS	31
Adding External Libraries	32
Configuration	33
Providing Configuration	33
Converting Canonical Spring Property Names to Environment Variables	33
Spring Profiles	33
Configure Locations to Find Spring Property Files	33
Spring Configuration Options	34
GCP Pub/Sub Binder Configuration Options	34
GCP Pub/Sub Consumer Options	35
GCP Pub/Sub Producer Options	36
Connector Configuration Options	36
Workflow Configuration Options	38
License	40
Support	40

Preface

Solace PubSub+ Connector for GCP Pub/Sub bridges data between the Solace PubSub+ Event Broker and GCP Pub/Sub providing a flexible and efficient way to integrate GCP Pub/Sub application data with your Solace-backed, event-driven architecture and the Event Mesh. The connector is deployable standalone or in redundancy modes of “active-standby” or “active-active” to allow for high-availability and horizontal scaling of your data movement. Each connector instance supports up to 20 individual workflows (source-to-target 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 GCP Pub/Sub using default credentials, with 2 workflows enabled, workflow 0 and workflow 1. Where:

- Workflow 0 is consuming messages from the Solace PubSub+ queue, `Solace/Queue/0`, and publishing them to the GCP Pub/Sub topic, `producer-destination`.
- Workflow 1 is consuming messages from the GCP Pub/Sub subscription, 'consumer-destination-sub' of topic `consumer-destination`, and publishing them to the Solace PubSub+ topic, `Solace/Topic/1`.

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 queues which do not exist.

Prerequisites

- [Solace PubSub+ Event Broker](#)
- [GCP Pub/Sub](#)

Quick Start common steps

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

1. Decide which [GCP Pub/Sub authentication](#) mechanism will be used and accordingly have `application_default_credentials.json` or service account key JSON etc. available for use.
2. Update the provided `samples/config/application.yml` with the values for your deployment.

Quick Start: Running the connector via command line

Run:

```
java -jar pubsubplus-connector-gcppubsub-1.1.0.jar --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](#).



By default, GCP Pub/Sub client library looks for the authentication credential as per [Application Default Credentials \(ADC\) Search Order](#).

For more info, see [GCP Pub/Sub Connection Details](#).

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 [-l FOLDER] [-p PROFILE] [-c FOLDER] [-j FILE] [-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 GCP Pub/Sub
```

```
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
-l, --libs	./libs	Directory containing required and optional dependency jars. Such as Micrometer metrics export dependencies (if configured). If not specified, it will use the current ./libs/ folder.
-p, --profile	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.

Parameter	Default Value	Description
<code>-c, --config</code>	<code>current folder</code>	Path to the folder containing the configuration files to be applied during connector startup for chosen profile. By default it is set to the current folder.
<code>-j, --jar</code>	<code>pubsubplus-connector-gcppubsub-1.1.0.jar</code>	Path to specified JAR file to start connector. If not specified, the default jar file is used from the current folder.
<code>-o, --options</code>	<code>no default values</code>	Specifies JVM options used on Connector start.
<code>-b, --background</code>	<code>N/A</code>	Runs Connector in the background. No logs will be displayed and Connector continues running in detached mode
<code>-h, --help</code>	<code>N/A</code>	Prints some help information and exits

Script also provides some help information from command line:

Solace PubSub+ Connector for GCP Pub/Sub

Usage: `start.sh [-l FOLDER] [-p PROFILE] [-c FOLDER] [-j FILE] [-o OPTIONS] [-b]`

To start connector use following parameters:

<code>-l --libs</code>	Directory containing required and optional dependency jars. Such as Micrometer metrics export dependencies (if configured). If not specified, it will use the current <code>./libs/</code> folder
<code>-p --profile</code>	Profile to be used with connector's configuration. The configuration file named <code>'application-<profile>.yaml'</code> will be used Default value is empty, no profile is used.
<code>-c --config</code>	Path to the folder containing the configuration files to be applied during connector startup for chosen profile. By default it is set to the current folder.
<code>-j --jar</code>	Path to specified JAR file to start connector. if not specified, the default jar file is used,

from the current folder

- o | --options Specifies JVM options used on Connector start.
example: '-Xms64M -Xmx1G'
- b | --background Runs the Connector in the background. No logs will be
displayed and Connector continues running in detached
mode.
- h | --help Print this help page and exit



By default, GCP Pub/Sub client library looks for the authentication credential as per [Application Default Credentials \(ADC\) Search Order](#).

For more info, see [GCP Pub/Sub Connection Details](#).

Quick Start: Running the connector as a Container

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

1. Add volume mapping in the sample `docker-compose.yml` for `application_default_credentials.json` from local to target path `/app/external/gcp/config/application_default_credentials.json`

```
solace-pubsub-connector-gcppubsub:
  volumes:
    - type: bind
      source: /path/to/local/gcloud/application_default_credentials.json
      target: /app/external/gcp/config/application_default_credentials.json
      read_only: true
```

2. 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.

3. 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 GCP Pub/Sub 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](#).



By default, GCP Pub/Sub client library looks for the authentication credential as per [Application Default Credentials \(ADC\) Search Order](#) within the running container.

For more info, see [GCP Pub/Sub Connection Details](#).

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|pubsub) # Input system
        output-<workflow-id>:
          destination: <output-destination> # Topic name
          binder: (solace|pubsub) # Output system

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



The connector only supports workflows in the directions of:

- `solace` → `GCP Pub/Sub`
- `GCP Pub/Sub` → `solace`

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.

GCP Pub/Sub Connection Details

This document assumes that the user is familiar with the Google Cloud [Application Default Credentials \(ADC\)](#). The Application Default Credentials (ADC) is a strategy used by the Google authentication libraries to automatically find credentials based on the application environment (local, google cloud, on-perm etc). Listing all possible authentication options is beyond the scope of this document.

The Google authentication libraries searches the application's environment in [Application Default Credentials \(ADC\) Search Order](#).

In addition to authentication credentials the GCP Pub/Sub client libraries also requires `spring.cloud.gcp.pubsub.project-id` for establishing connection.

The `spring.cloud.gcp.pubsub.project-id` can be set in `application.yml`:

```
spring:
  cloud:
    gcp:
      pubsub:
```

```
project-id: gcp-pubsub-project-id
```

For more information about other configurable properties and their defaults, see [Spring GCP Pub/Sub Configuration](#).

The Pub/Sub API uses the GRPC protocol to send API requests to the Pub/Sub service. See [GRPC Connection Settings](#) for configuration options for customizing the GRPC behavior.

Rest of this section documents various ways to provide the authentication credentials.

Using `gcloud cli` generated user credentials

If the `user credentials by using the gcloud cli` is already setup on user's system and `application_default_credentials.json` is in default location, simply using `samples/config/application.yml` should work.

```
java -jar pubsubplus-connector-gcppubsub-1.1.0.jar --spring.config.additional-
location=file:samples/config/
```

In case the `application_default_credentials.json` is not in the default path, or user wants to explicitly configure a path, it could be configured in `samples/config/application.yml` or passed through CLI.

Configure `spring.cloud.gcp.credentials.location` in `samples/config/application.yml`

```
spring:
  cloud:
    gcp:
      credentials:
        location: file:/path/to/gcloud/application_default_credentials.json
```

Using `spring.cloud.gcp.credentials.location` property with CLI

```
java -jar pubsubplus-connector-gcppubsub-1.1.0.jar --spring.config.additional-
location=file:samples/config/ --spring.cloud.gcp.credentials.location
=file:/path/to/gcloud/application_default_credentials.json
```

Using `GOOGLE_APPLICATION_CREDENTIALS` environment variable

```
GOOGLE_APPLICATION_CREDENTIALS=/path/to/gcloud/application_default_credentials.json &&
java -jar pubsubplus-connector-gcppubsub-1.1.0.jar --spring.config.additional-
location=file:samples/config/
```

Using Service Account Key

Alternatively, user can also use the [service account key](#) with `GOOGLE_APPLICATION_CREDENTIALS` environment variable.

```
GOOGLE_APPLICATION_CREDENTIALS=/path/to/gcloud/service_account_key.json && java -jar
pubsubplus-connector-gcpubsub-1.1.0.jar --spring.config.additional-location
=file:samples/config/
```

Connecting to Multiple Systems

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

For instance:

```
spring:
  cloud:
    stream:
      binders:

        # 1st solace binder in this example
        solace1:
          type: solace
          environment:
            solace:
              java:
                host: tcp://localhost:55555

        # 2nd solace binder in this example
        solace2:
          type: solace
          environment:
            solace:
              java:
                host: tcp://other-host:55555

        # The only pubsub binder
        pubsub1:
          type: pubsub
          # The GCP Pub/Sub binder doesn't support multi-binder system
          # The GCP Pub/Sub binder must be configured as it is configured in single-
binder configuration

        # Required for internal use
        undefined:
          type: undefined

      bindings:
        input-0:
```

```

destination: <input-destination>
binder: pubsub1
output-0:
  destination: <output-destination>
  binder: solace1 # Reference 1st solace binder
input-1:
  destination: <input-destination>
  binder: pubsub1
output-1:
  destination: <output-destination>
  binder: solace2 # Reference 2nd solace binder

```

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

Note that each binder must be defined under `spring.cloud.stream.binders`.



When connecting to multiple systems, all `solace` binder configuration must be specified using the multiple binder syntax i.e. under `spring.cloud.stream.binders.<binder-name>.environment`.

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

The GCP Pub/Sub binder doesn't support multi-binder system. The `pubsub` binder must be configured as it is configured in single-binder configuration.

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 GCP Pub/Sub 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.

Message Headers Mapping

Following GCP Pub/Sub `PubsubMessage` attributes are mapped to the Solace message properties.

Header Name	Type	Access	Description
<code>gcp_pubsub_message_id</code>	<code>string</code>	Read	The ID of the message, assigned by the GCP Pub/Sub service when the message is published. See PubsubMessage.getMessageId()
<code>gcp_pubsub_publish_time</code>	<code>long (seconds)</code>	Read	The time at which the message was published, populated by the GCP Pub/Sub service when it receives the Publish call. See PubsubMessage.getPublishTime()
<code>gcp_pubsub_ordering_key</code>	<code>string</code>	Read	The Pub/Sub message's orderingKey. See PubsubMessage.getOrderingKey()

Following Solace message headers marked with `Read` access are automatically forwarded to GCP Pub/Sub messages while headers with `Write` access can be set using the [User-configured Header Transforms](#) feature.

Header Name	Type	Access	Description
<code>solace_applicationMessageId</code>	<code>string</code>	Read/Write	The application provided message ID.
<code>solace_applicationMessageType</code>	<code>string</code>	Read/Write	The application provided message type. This value is used by applications only.
<code>solace_correlationId</code>	<code>string</code>	Read/Write	The correlation ID for the message.
<code>solace_dmqlEligible</code>	<code>boolean</code>	Write	Set the Solace message's dmqlEligible header.
<code>solace_expiration</code>	<code>long</code>	Write	Set the time at which the Solace message is set to expire.
<code>solace_priority</code>	<code>int</code>	Write	Set the Solace message's priority.
<code>solace_senderId</code>	<code>string</code>	Read	The Sender ID for the Solace message.

Header Name	Type	Access	Description
<code>solace_senderTimestamp</code>	<code>long</code>	Read	The time at which the Solace message was published.
<code>solace_sequenceNumber</code>	<code>string</code>	Read	The application-defined sequence number set on the Solace message.

Reserved Message Headers

The following are reserved header spaces:

- `solace_`
- `scst_`
- `gcp_pubsub|goog`
- 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.

Dynamic Producer Destinations

To route messages to dynamic destinations at runtime, use the [User-configured Header Transforms](#) feature described above to set the following headers:

Header Name	Type	Values	Applies To	Description
<code>scst_targetDestination</code>	<code>string</code>	Any valid destination name	Solace	Specifies the name of the dynamic destination to publish to. Setting this header overrides the configured destination.
<code>gcp_pubsub_topic</code>	<code>string</code>	Any valid destination name	GCP Pub/Sub	Specifies the name of the dynamic destination to publish to. Setting this header overrides the configured destination.
<code>solace_scst_targetDestinationType</code>	<code>string</code>	<code>(queue topic)</code>	Solace	Specifies the destination type of the dynamic destination. When unspecified, the configured or default destination type is used.



Setting the `scst_targetDestination` header under `solace.connector.default.workflow.transform-headers` may not be viable if not all workflows follow the same direction.

Asynchronous Publishing

Asynchronous publishing allows the connector to process new messages without the need to wait for the previous message's publish acknowledgment. While this can boost performance, it does come with a trade-off, as it may increase the chances and volume of duplicate messages.

A workflow can be configured to wait for publisher acknowledgements asynchronously with the following config options:

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

```
solace.connector.workflows.<workflow-id>.acknowledgment.publish-async
```

```
solace.connector.workflows.<workflow-id>.acknowledgment.back-pressure-threshold
```

```
solace.connector.workflows.<workflow-id>.acknowledgment.publish-timeout
```

Consult the [Workflow Configuration Options](#) table for details and default values of those config options.

In general, reducing the `publish-timeout` increases the probability of duplicate message deliveries, while increasing the `back-pressure-threshold` is likely to result in a higher occurrence of duplicate messages.

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

The above example configuration enables metrics collection through the configuration parameter of `management.metrics.export.simple.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 pubsub 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.

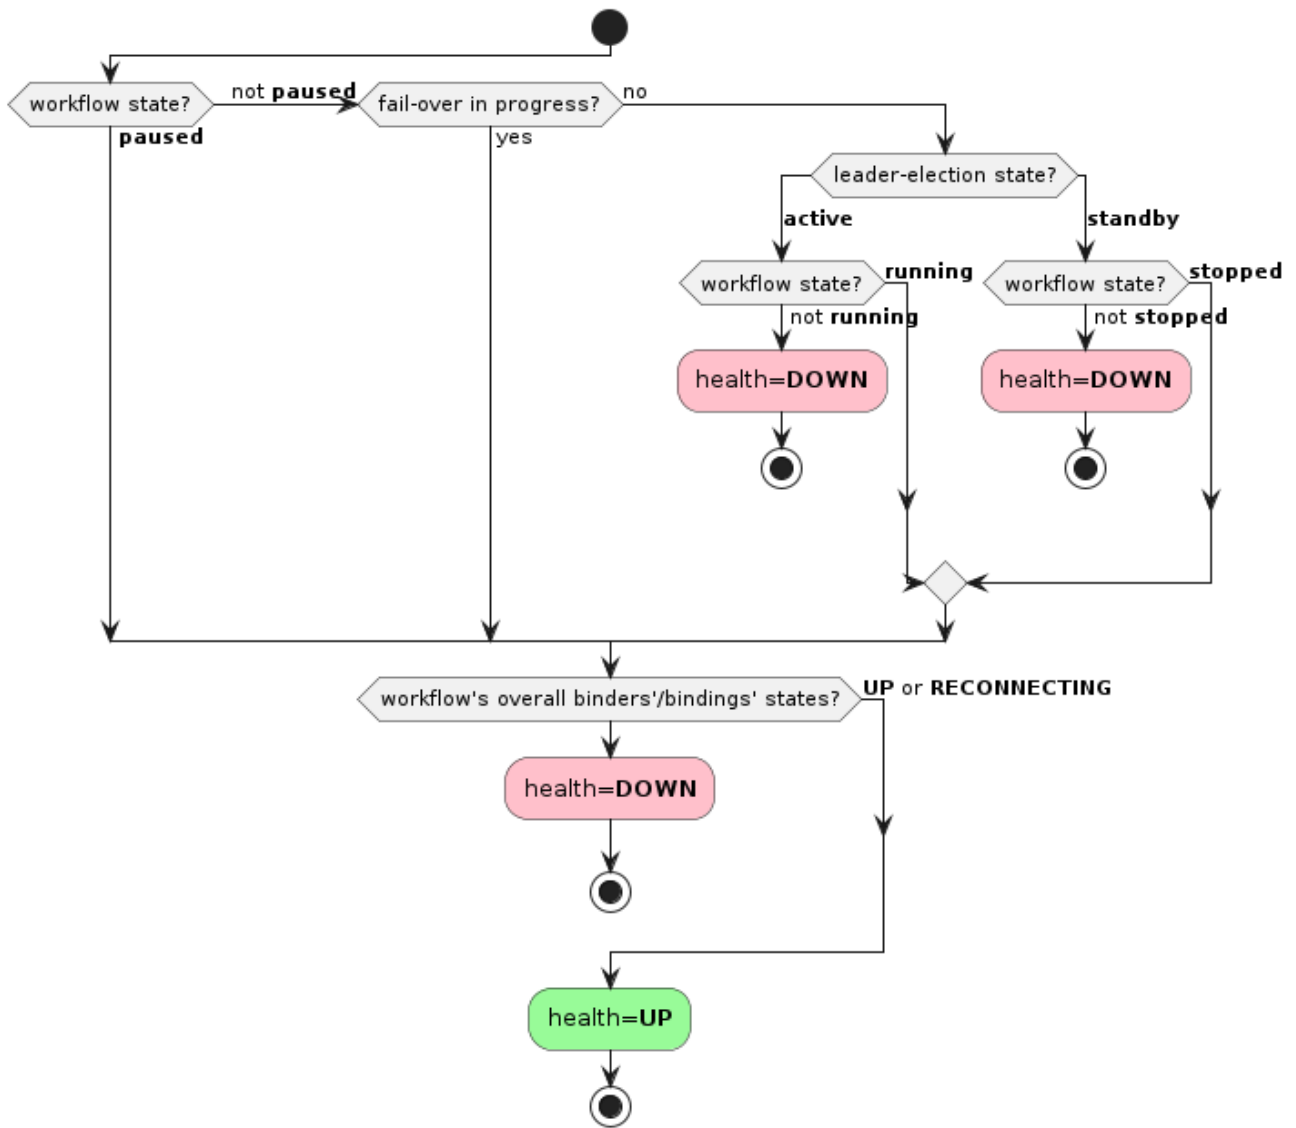


Figure 1. Workflow Health Resolution Diagram

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.

GCP Pub/Sub Binder Health

Health Status	Description
UP	Status indicating that the binder is functioning as expected.
UNKNOWN	Status indicating that the binder is in an unknown state. The binder will automatically recover when underlying connectivity issues are resolved. User intervention may be required.

Health Status	Description
DOWN	Status indicating that the binder is having difficulties reconnecting to the GCP Pub/Sub service. The binder will automatically recover when underlying connectivity issues are resolved. User intervention may be required.

For details, see [Google Cloud Pub/Sub Health Indicator](#).

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 (HTTP <code>GET</code>)	Request: None. Response: <pre>{ "state": "(active standby)", "mode": "(standalone active_active active_standby)" }</pre>

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

Endpoint	Operation	Payloads
/workflows/{workflowId}	Write (HTTP POST)	Request: <pre>{ "state": "STARTED STOPPED PAUSED RESUMED" }</pre> Response: 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
running	All bindings have state="running" .
stopped	All bindings have state="stopped" .
paused	All consumer bindings and all pausable producer bindings have state="paused" .
unknown	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.processor</code>	Timer	type: channel name: <bindingName> result: (success failure) exception: (none exception 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.processor</code>	Timer	type: channel name: <bindingNames> result: (success failure) exception: (none exception 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><destination>.<group>.errors</code>). 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>). Though a reminder that setting a binding's <code>group</code> is not supported.
<code>solace.connector.message.size.payload</code>	DistributionSummary Base Units: bytes	name: <bindingName>	Message payload size	

Name	Type	Tags	Description	Notes
<code>solace.connector.message.size.total</code>	DistributionSummary Base Units: bytes	name: <bindingName>	Total message size	
<code>solace.connector.publish.ack</code>	Counter Base Units: acknowledgments	name: <bindingName> result: (success failure) 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.metrics.export.<system>.enabled=true` in the `application.yml`.

Security

Securing Endpoints

Exposed Management Web Endpoints

By default, this connector only enables the `health` and `leaderelection` management endpoints. Where for the `health` endpoint, only the root status is returned 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 define 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.

CSRF Protection

Spring Boot enables CSRF protection by default on all management endpoints (see [Spring Cross Site Request Forgery Protection](#)). Though this connector disables CSRF protection for all POST requests on actuator endpoints so that users with write permissions (those with the `admin` role) can perform POST requests.

To fully disable CSRF protection, set the following config option:

```
solace.connector.security.csrf-enabled=false
```

TLS

TLS is disabled by default.

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

Adding External Libraries

The connector jar uses the `loader.path` property as the recommended mechanism for adding external libraries to the connector's classpath.

See [Spring Boot - PropertiesLauncher Features](#) for more info.

To add libraries to the connector's container image, see [the connector's container documentation](#).

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}.yaml`.

eg:

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

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](#)

GCP Pub/Sub Binder Configuration Options

The following properties are available at the GCP Pub/Sub binder level.

These properties are to be prefixed with `spring.cloud.stream.gcp.pubsub.default`.

For more info, See [GCP Pub/Sub binder configuration](#) and [GCP Pub/Sub configuration](#) in Spring Cloud GCP reference document.

Config Option	Type	Valid Values	Default Value	Description
<code>consumer.auto-create-resources</code>	boolean	(true false)	false	To automatically generate the underlying resources, like the GCP Pub/Sub topics and subscriptions for consumers. Note: In production environment, it is recommended to pre create a resources.
<code>producer.auto-create-resources</code>	boolean	(true false)	false	To automatically generate the underlying resources, like the GCP Pub/Sub topics for producers. Note: In production environment, it is recommended to pre create a resources.

GCP Pub/Sub Consumer Options

The following config options are available to customize the behavior of the GCP Pub/Sub consumers.

The GCP Pub/Sub consumers settings can be either global (when options prefixed with `spring.cloud.gcp.pubsub.subscriber.`) or subscription-specific (when options prefixed with `spring.cloud.gcp.pubsub.subscription.[subscription-name].`).

Config Option	Type	Valid Values	Default Value	Description
<code>parallel-pull-count</code>	int	> 0	1	The number of pull worker threads. Note: The GCP Pub/Sub binder creates thread pool with $\max(6, 2 * \text{parallel-pull-count})$ threads for pulling messages (i.e. if <code>parallel-pull-count</code> <= 3 still pool with 6 threads will be created).
<code>executor-threads</code>	int	> 0	4	The number of consumer threads used to process the pulled messages.
<code>max-ack-extension-period</code>	int	>= 0	0	The maximum period a message ack deadline will be extended, in seconds. See GCP Pub/Sub Consumer ACK Extension for more information.

Config Option	Type	Valid Values	Default Value	Description
<code>min-duration-per-ack-extension</code>	int	≥ 0	0	The lower bound for a single modify ack extension period, in seconds.
<code>max-duration-per-ack-extension</code>	int	≥ 0	0	The upper bound for a single modify ack extension period, in seconds.

See [global consumer configuration](#) and [subscription specific consumer configuration](#) in Spring Cloud GCP reference document.

GCP Pub/Sub Producer Options

See [producer configuration](#) in Spring Cloud GCP reference document.

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.
<code>management.leader-election.fail-over.back-off-multiplier</code>	double	≥ 1.0	2.0	The multiplier to apply to the back-off interval between each retry of a fail-over.

Config Option	Type	Valid Values	Default Value	Description
<code>management.leader-election.mode</code>	enum	(standalone active_active active_standby)	standalone	<p>The connector's leader election mode.</p> <p>standalone: A single instance of a connector without any leader election capabilities.</p> <p>active_active: A participant in a cluster of connector instances where all instances are active.</p> <p>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.</p>
<code>management.queue</code>	string	any	null	The management queue name.
<code>management.session.*</code>		See Spring Boot Auto-Configuration for the Solace Java API		<p>Defines the management session. This has the same interface as that used by <code>solace.java.*</code>.</p> <p>See Spring Boot Auto-Configuration for the Solace Java API for more info.</p>
<code>security.enabled</code>	boolean	(true false)	true	If <code>true</code> , security is enabled. Otherwise, anyone has access to the connector's endpoints.
<code>security.csrf-enabled</code>	boolean	(true false)	true	If <code>true</code> , CSRF protection is enabled. Makes sense only if <code>solace.connector.security.enabled</code> is <code>true</code> .
<code>security.users[<index>].name</code>	string	any	null	The name of this user.
<code>security.users[<index>].password</code>	string	any	null	The password of this user.

Config Option	Type	Valid Values	Default Value	Description
<code>security.users[<index>].roles</code>	<code>list<string></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 Default	<code>Map<string, string></code>	Key: A header name. Value: A SpEL string which accepts <code>headers</code> as parameters.	empty map	A mapping of header names to header value SpEL expressions. 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 Default	<code>boolean</code>	<code>(true false)</code>	<code>false</code>	If <code>true</code> , publisher acknowledgment processing is done asynchronously. 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 Default	int	≥ 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 Default	int	≥ -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](#).