



Solace PubSub+ Connector for MQTT

User Guide

Solace Corporation

Version 3.15.8

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	6
Enabling Workflows	7
Configuring Connection Details	8
Solace PubSub+ Connection Details	8
Preventing Message Loss when Publishing to Topic-to-Queue Mappings	8
MQTT Connection Details	8
Connecting to Multiple Systems	9
User-configured Header Transforms	12
User-configured Payload Transforms	14
Registered Functions	14
Message Headers	16
Solace Headers	16
Reserved Message Headers	16
Dynamic Producer Destinations	17
Asynchronous Publishing	18
Management and Monitoring Connector	19
Monitoring Connector's States	19
Exposed HTTP/HTTPS Endpoints	19
Health	21
Workflow Health	21
Solace Binder Health	22
Leader Election	23
Leader Election Modes: Standalone / Active-Active	23
Leader Election Mode: Active-Standby	23
Leader Election Management Endpoint	24
Workflow Management	26
Workflow Management Endpoint	26
Workflow States	27
Metrics	28
Connector Meters	28
Add a Monitoring System	29
Security	31

Securing Endpoints	31
Exposed Management Web Endpoints.....	31
Authentication & Authorization.....	31
TLS.....	32
Consuming Object Messages	33
Adding External Libraries	34
Configuration.....	35
Providing Configuration.....	35
Converting Canonical Spring Property Names to Environment Variables	35
Spring Profiles	35
Configure Locations to Find Spring Property Files	35
Obtaining Build Information	36
Spring Configuration Options.....	36
Connector Configuration Options	37
Workflow Configuration Options.....	38
MQTT Producer Configuration Options	40
MQTT Consumer Configuration Options	41
Logging	42
Configuring Logback.....	42
License	43
Support.....	43

Preface

Solace PubSub+ Connector for MQTT bridges data between the Solace PubSub+ Event Broker and MQTT providing a flexible and efficient way to integrate MQTT 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

Presuming you're using the default `application.yml` provided with the download, follow one of the below quick start below to connect to a PubSub+ event broker with the MQTT. The quick starts use default credentials as examples to get started with two workflows enabled, workflow 0 and workflow 1.

Where:

- Workflow 0 consumes messages from the Solace PubSub+ queue, `Solace/Queue/0`, and publishes them to the MQTT producer destination, `producer-destination`.
- Workflow 1 consumes messages from the MQTT consumer destination, `consumer-destination`, and publishes 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 does not provision queues that do not exist.

Prerequisites

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

Quick Start common steps

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

1. 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-mqtt-3.15.8.jar --spring.config.additional-location  
=file:samples/config/
```



By default, this command detects any Spring Boot configuration files as per the [Spring Boot's default locations](#).

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

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

For convenience, you can start the connector through the shell script using the following command:

```
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]
```

The script shows you all errors at the same time:

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

The script shows you all errors at the same time:

Solace PubSub+ Connector for MQTT

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 situations where you don't provide a parameter, the script runs with the predefined values as follows:

Parameter	Default Value	Description
-n, --name	application	The name of the connector instance, that is configured in [spring.application.name]. This name impacts on grouping connectors only.
-l, --libs	./libs	The directory that contains the required and optional dependency JAR files, such as Micrometer metrics export dependencies (if configured). If this option is not specified, it will use the current ./libs/ directory.

Parameter	Default Value	Description
<code>-p, --profile</code>	empty, no profile is used	The profile to be used with the connector's configuration. The configuration file named 'application-<profile>.yaml' is used. If this option is not specified, no profile is used.
<code>-c, --config</code>	<code>./</code> or current folder	The path to the folder containing the configuration files to be applied when the connector starts up the chosen profile. If not specified, the current directory is used.
<code>-H, --host</code>	127.0.0.1	Specifies the host where the connector runs.
<code>-P, --port</code>	8090	Specifies the port where connector runs.
<code>-mp, --mgmt_port</code>	9009	Specifies the management port for back calls of current connector from PubSub+ Connector Manager. This parameter is ignored if the <code>-cm</code> parameter is not provided.
<code>-j, --jar</code>	pubsubplus-connector-mqtt-3.15.8.jar	The path to the specified JAR file to start the connector. If the option is not specified, the default JAR file is used from the current directory.

Parameter	Default Value	Description
<code>-cm, --manager</code>	<code>application</code>	Specifies PubSub+ Connector Manager to use the configuration storage and allows you to enable the cloud configuration for the connector. When this parameter is enabled, you can specify the <code>-mp</code> or <code>--mgmt_port</code> , <code>-H</code> or <code>--host</code> , and <code>-cmh</code> with the <code>-cmp</code> parameters, unless you want to use default values for those parameters. Be aware, this option disable listed parameters to be read from configuration file. In this case, the operator must explicitly specify the parameters for the script, otherwise defaultdefault values are used.
<code>-cmh, --cm_host</code>	<code>127.0.0.1</code>	Specifies the host where Connector Manager is running. This parameter is ignored if the <code>-cm</code> parameter is not provided.
<code>-cmp, --cm_port</code>	<code>9500</code>	Specifies the port where Connector Manager is running. This parameter is ignored if <code>-cm</code> parameter is not provided.
<code>-o, --options</code>	<code>no default values</code>	Specifies the JVM options used on when the connector starts. For example, <code>-Xms64M -Xmx1G</code> .
<code>-tls</code>	<code>N/A</code>	Specifies to use HTTPS instead of HTTP. . When this parameter is used, the configuration file must contain an additional section with the preconfigured paths for the key store and trust store files.
<code>-s, --show</code>	<code>N/A</code>	Performs a dry run (does nothing). The output prints the start CLI command and its raw output and exits. This parameter is useful to check your parameters without running the connector.

Parameter	Default Value	Description
<code>-b, --background</code>	N/A	Runs the connector in the background. No logs are shown and the connector continues running in detached mode.
<code>-h, --help</code>	N/A	Prints the 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

The following steps show how to use the sample docker compose file that has been included in the package:

1. Change to the `docker` directory:

```
cd samples/docker
```

This directory contains both the `docker-compose.yml` file as well as an `.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:

- Exposes the connector's `8090` web port to `8090` on the host.
- Connects a PubSub+ event broker and MQTT exposed on the host using default ports.
- Mounts the `samples/config` directory.
- Mounts the previously defined `libs` directory.
- Creates a `healthcheck` user with read-only permissions.
 - The default username and password for this user can be found within the `.env` file.
 - This user overrides any users you have defined in your `application.yml`. See [here](#) for more information.
- Uses the connector's management health endpoint as the container's health check.

For more information 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|mqtt) # Input system
        output-<workflow-id>:
          destination: <output-destination> # Topic name
          binder: (solace|mqtt) # Output system

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



The connector only supports workflows in the directions of:

- `solace` → `MQTT`
- `MQTT` → `solace`

For more information about Spring Cloud Stream and the Solace PubSub+ binder, see:

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

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

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

For more information 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. For example, if queue ingress is shutdown.

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

MQTT Connection Details

The Spring Cloud Stream Binder for MQTT uses the following configuration to configure MQTT Connector connection.

```
spring:
  cloud:
    stream:
      bindings:
        input-0:
          destination: # configure MQTT subscribe topic
          binder: mqtt
        output-0:
          destination: # configure solace topic where MQTT events are published
          binder: solace
  mqtt:
    java:
      url: # MQTT Host
      version: # v3 or v5 (default is set to v5)
      username: # username to connect to MQTT broker
```

```

password: # password to connect to MQTT broker
cleanSession: # default is set to true
connectionTimeout: # default is set to 30
keepAliveInterval: # default is set to 60
persistence: # 'memory' or 'file'. default is memory
persistenceDirectory: # default is /tmp/mqtt
automaticReconnect: # default is false
automaticReconnectDelay: # default is 0
sessionExpiryInterval: # default is null
receiveMaximum: # default is null
maximumPacketSize: # default is null
topicAliasMaximum: # default is null
requestResponse: # default is false
requestProblemInfo: # default is false
userProperties: # key-value pair
  property-1 : # value
  property-2 : # value
ssl:
  caFile: # absolute path to file. default is empty
  clientCertificateFile: # absolute path to file. default is empty
  clientKeyFile: # absolute path to file. default is empty
  password: # client key file password if set. default is empty
  sslHostNameVerification: # default is false
willAndTestament:
  lastWillTopic: # default is null
  lastWillQos: # default is 0
  lastWillRetain: # default is false
  payload: # default is null
  payloadFormatIndicator: # default is false
  willDelayInterval: # default is null
  messageExpiryInterval: # default is null
  contentType: # default is empty
  responseTopic: # default is empty
  correlationData: # default is empty

```

Connecting to Multiple Systems

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

For example:

```

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

# 1st mqtt binder
mqtt1:
  type: mqtt
  # Add `environment` property map here if you need to customize this binder.
  environment:
    mqtt:
      java:
        url: tcp://localhost:1883

# 2nd mqtt binder
mqtt2:
  type: mqtt
  # Add `environment` property map here if you need to customize this binder.
  environment:
    mqtt:
      java:
        url: tcp://other-host:1883

# Required for internal use
undefined:
  type: undefined

bindings:
  input-0:
    destination: <input-destination>
    binder: mqtt1
  output-0:
    destination: <output-destination>
    binder: solace1 # Reference 1st solace binder
  input-1:
    destination: <input-destination>
    binder: mqtt2
  output-1:
    destination: <output-destination>
    binder: solace2 # Reference 2nd solace binder

```

The configuration above defines two binders of type `solace` and two binder of type `mqtt`, which are

then referenced within bindings.

Each binder above 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. For example, under the `spring.cloud.stream.binders.<binder-name>.environment`.
- Do not use single-binder configuration (for example, `solace.java.*` or `mqtt.java.*` at the root of your `application.yml`) while using the multiple binder syntax.

User-configured Header Transforms

Deprecated

This feature is deprecated and will be removed in a future release. Use [Message Transforms](#) instead.

Generic Migration Rules:

1. Configuration Structure Changes:

- a. Replace `solace.connector.workflows.<n>.transform-headers` with `solace.connector.workflows.<n>.transform` and within it:
 - i. Add `enabled: true` to enable transforms
 - ii. Move `transform-headers.expressions` to the `transform.expressions` list property

2. Expression Syntax Changes:

- Replace direct header references, `headers.<headerName>`, with `source['headers']['<headerName>']` for reading
- Use `target['headers']['<headerName>'] = ...` for writing
- Replace Java methods (`T(String)`, etc.) with built-in functions:
 - `T(String).join()` → `#joinString()`
 - `split()` → `#splitString()`
 - `toUpperCase()` → `#upperCaseString()`
 - `toLowerCase()` → `#lowerCaseString()`
 - etc



Example Migration:

Old Configuration:

```
solace:
  connector:
    workflows:
      0:
        transform-headers:
          expressions:
            route: "T(String).format('%s/%s', headers.region,
headers.status)"
            count: "headers.count.toString()"

```

New Configuration:

```
solace:

```

```
connector:
  workflows:
    0:
      transform:
        enabled: true
        expressions:
          - transform: "target['headers']['route'] = #joinString('/',
source['headers']['region'], source['headers']['status'])"
          - transform: "target['headers']['count'] =
#convertNumberToString(source['headers']['count'])"
```

Generally, the consumed message's headers are propagated through the connector to the output message. If you want to transform the headers, then you can 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 information about Spring Expression Language (SpEL) expressions, see [Spring Expression Language \(SpEL\)](#).

User-configured Payload Transforms



Deprecated

This feature is deprecated and will be removed in a future release. Use [Message Transforms](#) instead.

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 the future.

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 the available registered functions:

Registered Function Signature	Description
<code>boolean isPayloadBytes(Object obj)</code>	<p>Returns whether the object <code>obj</code> is an instance of <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 information about Spring Expression Language (SpEL) expressions, see [Spring Expression Language \(SpEL\)](#).

Message Headers

Solace and MQTT 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 in 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 future releases of this connector.

Dynamic Producer Destinations

To route messages to dynamic destinations at runtime, use the [Message Transform](#) feature to set the following headers:

Header Name	Type	Values	Applies To	Description
<code>scst_targetDestination</code>	<code>string</code>	Any valid destination name	Solace, MQTT	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.*` 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 the publisher's acknowledgments, asynchronously with the following configuration 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
```

See the [Workflow Configuration Options](#) table for details and defaults values of those configuration options.

In general, reducing the `publish-timeout` option increases the probability of duplicate message deliveries, while increasing the `back-pressure-threshold` option 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](#).

An Actuator shares information through the endpoints reachable over HTTP/HTTPS. The endpoints that are available are configured in the connector configuration file.

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 sample configuration enables metrics collection through the configuration parameter of `management.simple.metrics.export.enabled` set to `true` and then shares them through the HTTP/HTTPS endpoint together with other sections configured for the current connector.

Exposed HTTP/HTTPS Endpoints

The set of endpoints exposed through the HTTP/HTTPS endpoint.

- Exposed endpoints are available if you query the endpoints using the web interface (for example `https://localhost:8090/actuator/<some_endpoint>`) and also available in PubSub+ Connector Manager.
- The operator may choose to not expose all or some of these endpoints. If so, the Actuator endpoints that are not exposed are not visible if you query the endpoints (for example, `https://localhost:8090/actuator/<some_endpoint>`) nor in 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 (for example, Datadog, Prometheus, etc.) to collect metrics.

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

```
GET: /actuator/
```

The following shows an 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 using the [Spring Boot Actuator health endpoint](#).

To configure the information returned by the `health` endpoint, configure the following properties:

- `management.endpoint.health.show-details`
- `management.endpoint.health.show-components`

For more information, about health endpoints, see [Spring Boot documentation](#).

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

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

This configuration would always show the full details 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	A status that indicates the workflow is functioning as expected.
DOWN	A status that indicates the workflow is unhealthy. Operator intervention may be required.




Figure 1. Workflow Health Resolution Diagram

This health indicator is enabled default. To disable it, set the property as follows:

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

Solace Binder Health

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

Leader Election

The connector has three leader election modes for redundancy:

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 the following configuration:

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

Leader Election Modes: Standalone / Active-Active

When the connector starts, all enabled workflows start at the same time. The connector itself is considered as always active.

Leader Election Mode: Active-Standby

If the connector is in active-standby mode, 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 creates 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, failover only happens when there are infrastructure failures (for example, the JVM goes down or networking failures to the management queue).

If a workflow fails to start or stop during failover, it will retry up to some maximum defined by the configuration option, `solace.connector.management.leader-election.fail-over.max-attempts`.

During failover, the connector attempts to start or stop all enabled workflows. After an attempt has been made to start or stop each workflow, the connector transitions to the active/standby mode regardless of the status of the workflows.

Leader Election Management Endpoint

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

Operators can 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>)	<p>Request: None.</p> <p>Response:</p> <pre> { "mode": { "type": "(standalone active_active ① active_standby)", "state": "(active standby)", ② "source": { ③ "queue": "<management-queue-name>", "host": "<management-host>", "msgVpn": "<management-vpn>" } } } </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>

The MQTT v3 consumer flow (where the `mqtt` binder is defined in the `input-#`) supports `standalone` and `active_standby` modes only.



The MQTT v5 consumer flow supports `standalone`, `active_active` and `active_standby` modes. When using the `active_active` mode, the consumer binding's `destination` property should use the MQTT shared subscription syntax in

order to distribute the messages among the multiple instances, otherwise each instance will receive the copy of the message.

Workflow Management

Workflow Management Endpoint

A custom `workflows` management endpoint using [Spring Actuator](#) is 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>
<code>/workflows/{workflowId}</code>	Write (HTTP <code>POST</code>)	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 continue to persist when write operations are used to change the 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 failover 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 the producer or consumer binding is not implementing Spring's Lifecycle interface, Spring always reports the bindings as **state=N/A**. The **state=N/A** is ignored when deciding the overall state of the workflow. For example, if the consumer's binding is **state=running** and producer's binding **state=N/A** (or vice-versa), the workflow state would be **running**.

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

Metrics

This connector uses [Spring Boot Metrics](#) that 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 name: <bindingName> result: (success failure) exception: (none exception simple class name)	The 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 name: <bindingNames> result: (success failure) exception: (none exception simple class name)	The 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 (for example, bindings can have a matching <code>destination</code> , <code>group</code> , or both). NOTE: Setting a binding's <code>group</code> is not supported.
<code>solace.connector.message.size.payload</code>	DistributionSummary Base Units: bytes	name: <bindingName>	The message payload size.	

Name	Type	Tags	Description	Notes
<code>solace.connector.message.size.total</code>	DistributionSummary Base Units: bytes	name: <bindingName>	The total message size.	
<code>solace.connector.publish.ack</code>	Counter Base Units: acknowledgments	name: <bindingName> result: (success failure) exception: (none exception simple class name)	The publish acknowledgment count.	
<code>solace.connector.transform.expressions.count</code>	Gauge Base Units: expressions	workflow.id: <workflowId>	Transformation expressions count	
<code>solace.connector.transform.time</code>	Timer	workflow.id: <workflowId> result: (success failure)	Transformation execution time	This is the aggregate time for all expressions used to transform a single message.



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 (or batch of messages), how long it took to process that message, and if that step completed successfully.

Instead, we recommend that you 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 includes the following monitoring systems:

- [Datadog](#)
- [Dynatrace](#)
- [Influx](#)
- [JMX](#)
- [OpenTelemetry \(OTLP\)](#)

- [StatsD](#)

To add additional monitoring systems, add the system's `micrometer-registry-<system>` JAR file and its dependency JAR files 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

There are many endpoints that are automatically enabled for this connector. For a comprehensive list, see [Management and Monitoring 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

This release of the connector only supports basic HTTP authentication.

By default, no users are created unless the operator configures them in their configuration file. The configuration parameters responsible for security are as follows:

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

In the above example, we have created two 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 `solace.connector.security.enabled` parameter `false`.



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

The following example shows you 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 two 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, and so on), overriding works by replacing the entire list. In other words, you must pick one place to define all your users, whether in a **single** application properties file or as environment variables.

For more information, see [Spring Boot - Merging Complex Types](#).

TLS

TLS is disabled by default.

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

Consuming Object Messages

For the connector to process object messages, it needs access to the classes which define the object payloads.

Assuming that your payload classes are in their own project(s) and are packaged into their own jar(s), place these jar(s) and their dependencies (if any) onto [the connector's classpath](#).



It is recommended that these jars only contain the relevant payload classes to prevent any oddities.

In the jar(s), your class files must be archived in the same directory/classpath as the application that publishes them.



e.g. If the source application is publishing a message with payload type, `MySerializablePayload`, defined under classpath `com.sample.payload`, then when packaging the payload jar for the connector, the `MySerializablePayload` class must still be accessible under the `com.sample.payload` classpath.

Typically, build tools such as Maven or Gradle will handle this when packaging jars.

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

For information about about how the connector detects configuration properties, see [Spring Boot: Externalized Configuration](#).

Converting Canonical Spring Property Names to Environment Variables

For information about converting the Spring property names to environment variables, see the [Spring documentation](#).

Spring Profiles

If multiple configuration files exist within the same configuration directory for use in different environments (development, production, etc.), use Spring profiles.

Using Spring profiles allow you to define different application property files under the same directory using the filename format, `application-{profile}.yml`.

For example:

- `application.yml`: The properties in non-specific files that always apply. Its properties are overridden by the properties defined in profile-specific files.
- `application-dev.yml`: Defines properties specific to the development environment.
- `application-prod.yml`: Defines properties specific to the production environment.

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

See [Spring Boot: Profile-Specific Files](#) for more information and an example.

Configure Locations to Find Spring Property Files

By default, the connector detects any Spring property files as described in the [Spring Boot's default locations](#).

- If you want to add additional locations, add `--spring.config.additional-location=file:<custom-config-dir>` (This parameter is 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>`.

For more information about configuring locations to find Spring property files, see [Spring Boot documentation](#).



If you want configuration files for multiple, different connectors within the same `config` directory for use in different environments (such as development, production, etc.), we recommend that you use [Spring Boot Profiles](#) instead of child directories. For example:

- Set up your configuration like this:
 - `config/application-prod.yml`
 - `config/application-dev.yml`
- Do not do this:
 - `config/prod/application.yml`
 - `config/dev/application.yml`

Child directories are intended to be used for merging configuration from multiple sources of configuration properties. For more information and an example of when you might want to use multiple child directories to compose your application's configuration, 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 want to exclude build data from the output of the `info` endpoint, set `management.info.build.enabled` to `false`.

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

Spring Configuration Options

This connector packages many libraries for you 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](#)

Connector Configuration Options

The following table lists the configuration options. The following options in **Config Option** are prefixed with `solace.connector.:`


Config Option	Type	Default Value	Description
<code>management.leader-election.fail-over.max-attempts</code>	<code>int</code> Constraint: <code>> 0</code>	<code>3</code>	The maximum number of attempts to perform a fail-over.
<code>management.leader-election.fail-over.back-off-initial-interval</code>	<code>long</code> Constraint: <code>> 0</code>	<code>1000</code>	The initial interval (milliseconds) to back-off when retrying a fail-over.
<code>management.leader-election.fail-over.back-off-max-interval</code>	<code>long</code> Constraint: <code>> 0</code>	<code>10000</code>	The maximum interval (milliseconds) to back-off when retrying a fail-over.
<code>management.leader-election.fail-over.back-off-multiplier</code>	<code>double</code> Constraint: <code>>= 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>	<code>enum</code> One of: <ul style="list-style-type: none"> • <code>standalone</code> • <code>active_active</code> • <code>active_standby</code> 	<code>standalone</code>	<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>	<code>string</code>	<code>null</code>	The management queue name.

Config Option	Type	Default Value	Description
<code>management.session.*</code>	See Spring Boot Auto-Configuration for the Solace Java API		Defines the management session. This has the same interface as that used by <code>solace.java.*</code> . See Spring Boot Auto-Configuration for the Solace Java API for more info.
<code>security.enabled</code>	boolean	true	If <code>true</code> , security is enabled. Otherwise, anyone has access to the connector's endpoints.
<code>security.users[<index>].name</code>	string	null	The name of the user.
<code>security.users[<index>].password</code>	string	null	The password for the user.
<code>security.users[<index>].roles</code>	list<string> Valid values: • <code>admin</code>	empty list (i.e. read-only)	The list of roles that the specified user has. It has read-only access if no roles are returned.

Workflow Configuration Options

These configuration options are defined under the following prefixes:

- `solace.connector.workflows.<workflow-id>.`: If the options support per-workflow configuration and the default prefixes.
- `solace.connector.default.workflow.`: If the options support default workflow configuration.

Config Option	Type	Default Value	Description
<code>enabled</code>	boolean	false	If <code>true</code> , the workflow is enabled. <div>  Cannot be set at the default workflow level. </div>
<code>transform.enabled</code>	boolean	false	If <code>true</code> , message transformation is enabled for the workflow. Disables the legacy <code>transform-headers.*</code> and <code>transform-payload.*</code> options. See Message Transforms for more info.

Config Option	Type	Default Value	Description
<code>transform.source-payload.content-type</code>	string One of: <ul style="list-style-type: none"> <code>application/vnd.solace.micro-integration.unspecified</code> <code>application/json</code> 	<code>application/vnd.solace.micro-integration.unspecified</code>	The content type to interpret the source payload as. See Content Type Interpretation for more info.
<code>transform.target-payload.content-type</code>	string One of: <ul style="list-style-type: none"> <code>application/vnd.solace.micro-integration.unspecified</code> <code>application/json</code> 	<code>application/vnd.solace.micro-integration.unspecified</code>	The content type to interpret and serialize the target payload as. See Content Type Interpretation for more info.
<code>transform.expressions[<index>].transform</code>	string A SpEL expression		A SpEL expression at some <code><index></code> in the ordered list of expressions to transform the message. See Message Transform for more info.
<code>transform-headers.expressions</code>	Map<string, string> Key: A header name. Value: A SpEL string that accepts <code>headers</code> as parameters.	empty map	<div>  <div> Deprecated. Use Message Transforms (<code>transform.*</code>) instead. </div> </div> <p>A mapping of header names to header value SpEL expressions.</p> <p>The SpEL context contains the <code>headers</code> parameter that can be used to read the input message's headers.</p>
<code>acknowledgment.publish-async</code>	boolean	<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 the publisher acknowledgments are processed synchronously regardless of this setting.

Config Option	Type	Default Value	Description
<code>acknowledgment.back-pressure-threshold</code>	int Constraint: ≥ 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>	int Constraint: ≥ -1	600000	The maximum amount of time (in millisecond) to wait for asynchronous publisher acknowledgments before considering a message as failed. A value of -1 means to wait indefinitely for publisher acknowledgments.

MQTT Producer Configuration Options

These configuration options are all prefixed by `spring.cloud.stream.mqtt.bindings.<channel-name>.producer:`

In case if these properties need to be applied by default to all producers `spring.cloud.stream.mqtt.default.producer`

Config Option	Type	Valid Values	Default Value	Description
<code>clientId</code>	string	any	Randomly generated UUID	Set the client-id for the producer
<code>qos</code>	integer	0,1,2	1	Set the QOS level for the publish event
<code>retained</code>	boolean	true or false	false	Whether to retain published messages
<code>async</code>	boolean	true or false	false	Whether to use async publish
<code>allowedHeaders</code>	string[]		empty	MQTT V5 only. List of headers that need to be included or excluded as part of user properties. To exclude a header prefix it with ! . e.g. !timestamp , * would exclude only timestamp and include all other headers.

MQTT Consumer Configuration Options

These configuration options are all prefixed by `spring.cloud.stream.mqtt.bindings.<channel-name>.consumer`:

In case if these properties need to be applied by default to all consumers `spring.cloud.stream.mqtt.default.consumer`.

Config Option	Type	Valid Values	Default Value	Description
<code>clientId</code>	<code>string</code>	<code>any</code>	<code>Randomly generated UUID</code>	Set the client-id for the consumer
<code>qos</code>	<code>integer</code>	<code>0,1,2</code>	<code>1</code>	Set the QOS level for the subscribed event
<code>binary</code>	<code>boolean</code>	<code>true or false</code>	<code>false</code>	Set to true to convert payload to bytes
<code>allowedHeaders</code>	<code>string[]</code>		<code>empty</code>	MQTT V5 only. List of headers that need to be included or excluded as part of user properties. To exclude a header prefix it with <code>!</code> (e.g. <code>!timestamp</code> , <code>*</code> would exclude only timestamp and include all other headers).
<code>dltEnabled</code>	<code>boolean</code>	<code>true or false</code>	<code>false</code>	Set to true to enable DLT (dead letter topic) for the failed consumer messages. If <code>dltName</code> is not configured explicitly, failed messages are forwarded to the <code>dlt/{source-topic-name}</code> topic.
<code>dltName</code>	<code>string</code>	<code>valid MQTT topic name</code>	<code>empty</code>	When <code>dltEnabled</code> is set to <code>true</code> , failed consumer messages will be forwarded to the topic configured with <code>dltName</code>

Logging

Configuring Logback

If you require Logback configuration beyond what is already available through Spring config properties, then you can decide to run your connector with a `logback-spring.xml` file. For information about using the logback, see:

- [logging.config](#) spring property.
- [Logging](#)
- [Configure Logback for Logging](#) sections in the Spring documentation.

The main difference compared to [what Spring Boot provides](#), is that this connector provides its own alternative logback configuration files that can be [included](#) into your `logback-spring.xml` file.

These new files can be included and are found at `com/solace/connector/core/logging/logback/`:

`defaults.xml`

Provides conversion rules, pattern properties and common logger configurations.



Always include `defaults.xml`

We recommend that you always include this file in your `logback-spring.xml` file as it includes a `%sanitize` `<conversionRule>` that's applied to the default `CONSOLE_LOG_PATTERN` and `FILE_LOG_PATTERN`.

This conversion rule does some filtering to obfuscate potentially sensitive data from the logs.

`console-appender.xml`

Adds a `ConsoleAppender` using the `CONSOLE_LOG_PATTERN`.

`file-appender.xml`

Adds a `RollingFileAppender` using the `FILE_LOG_PATTERN` and `ROLLING_FILE_NAME_PATTERN` with appropriate settings.

Aside from these new include-able files, this connector still supports all the same logging options that regular Spring Boot does.

For more information and examples, see the [Configure Logback for Logging](#) section in the Spring documentation.

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](#).