



PUBLIC

SAP HANA Streaming Analytics 2.0 SP 04

Document Version: 1.0.2 – 2022-09-22

SAP HANA Streaming Analytics: Configuration and Administration Guide

Content

- 1 SAP HANA Streaming Analytics: Configuration and Administration Guide. 7**
- 2 Introduction. 8**
 - 2.1 SAP HANA Streaming Analytics Overview. 8
 - 2.2 Streaming Cluster Nodes. 9
 - 2.3 Tools for Administration and Monitoring. 10
 - 2.4 SAP HANA Streaming Analytics Licenses. 10
- 3 Internal and External Ports for SAP HANA Streaming Analytics. 12**
- 4 Configuring SAP HANA Streaming Analytics for the First Time. 13**
 - 4.1 Configuring an External unixODBC Driver Manager on Linux. 14
 - Configuring an ODBC Driver Manager for a Streaming Host with a Proxy. 16
 - 4.2 Authorization for Administration Tasks. 17
 - 4.3 Configure Web IDE in XS Advanced Cockpit. 20
 - Set Up the Space 20
 - Set Up Role Collections. 21
 - Enable Spaces for Development. 23
 - Configure the SAP HANA Streaming Analytics Service Broker. 23
- 5 Managing the Streaming Service and Hosts. 26**
 - 5.1 Verify that the Streaming Service is Installed and Running. 27
 - Verify Using HANA Studio. 27
 - Verify Using the Streaming Host. 27
 - 5.2 Start the Streaming Service Manually. 28
 - 5.3 Stop the Streaming Service Manually. 28
 - 5.4 Auto-Restart the Streaming Service. 29
 - Stop-Restart Using SAP HANA cockpit. 29
 - Stop-Restart Using SAP HANA Studio. 30
- 6 Cluster Configuration. 31**
 - 6.1 File and Directory Infrastructure. 31
- 7 High Availability. 36**
- 8 Set Up System Replication. 39**
 - 8.1 Configure the Primary and Secondary Systems. 41
 - 8.2 Modify the Assigned Worker Group. 42
 - Modify Using SAP HANA Studio. 42

8.3	Enabling the Activation of Content in the SAP HANA Repository.	43
9	Deploying an Adapter to a Cluster.	44
9.1	Application Deployment Parameters.	45
9.2	Sample Application Deployment Configuration File.	47
9.3	Adapter Warm Standby.	48
9.4	Adapter Affinity.	48
9.5	Running Custom Adapters on a Remote Server.	48
10	Deploying a Project to a Cluster.	50
10.1	Deploying a Project to a Cluster from Studio.	50
10.2	Deploying a Project to a Cluster Using streamingclusteradmin.	51
10.3	Editing Cluster Parameters in Project Configuration.	52
10.4	Accessing Project Contents using SAP HANA Smart Data Access	53
11	Monitoring Streaming Analytics.	54
11.1	Monitoring a Project with Utilities.	55
	Monitoring Project Status.	56
11.2	Streaming Analytics Statistics in SAP HANA Studio.	57
11.3	Editing the XS Job for SAP HANA Streaming Analytics.	60
11.4	Monitoring with Metadata Streams.	61
	_ESP_Adapter_Statistics.	63
	_ESP_Clients.	64
	_ESP_Clients_Monitor.	64
	_ESP_Clockupdates.	66
	_ESP_Columns.	66
	_ESP_Config.	67
	_ESP_Connectors.	67
	_ESP_GD_Sessions.	68
	_ESP_Keycolumns.	69
	_ESP_Project_Monitor.	69
	_ESP_Resource_Info.	70
	_ESP_RunUpdates.	72
	_ESP_Streams.	72
	_ESP_Streams_Monitor.	73
	_ESP_Streams_Topology.	74
	_ESP_Subscriptions.	74
	_ESP_Subscriptions_Ext.	75
12	Opening SAP HANA Cockpit.	76
13	Managing Streaming Analytics with SAP HANA Cockpit.	77
13.1	Managing Cluster Settings.	78

	Store an Encrypted Password for a Key Store.	79
	Adjust Heartbeat Settings.	80
	Edit the Streaming Node Name.	81
	Edit Application Restart Settings.	84
	Setting Disallowed Users.	85
	Configuring the Streaming Web Service.	85
	Configuring the Web Services Provider.	87
13.2	Managing Streaming Nodes.	91
	Start or Stop the Streaming Web Service.	93
	Start or Stop the Web Services Provider.	94
13.3	Managing Streaming Workspaces.	94
	Add a Streaming Workspace.	95
	Remove a Streaming Workspace.	96
13.4	Managing Streaming Projects.	97
	Add a Streaming Project.	99
	Remove a Streaming Project.	100
	Start or Stop a Streaming Project.	101
13.5	Managing Streams.	102
13.6	Managing Streaming Adapters.	103
	Add an Adapter.	105
	Remove an Adapter.	106
	Start or Stop an Adapter.	107
13.7	Managing Streaming Publishers.	108
13.8	Managing Streaming Subscribers.	108
13.9	Monitoring Alerts.	109
	Alert Details.	109
	Streaming Analytics Alerts in SAP HANA Cockpit.	110
13.10	Managing Streaming Permissions.	118
	Grant a Permission to a User.	123
	Revoke a Permission from a User.	124
	Create a Role.	125
	Remove a Role.	126
	Grant a Permission to a Role.	126
	Revoke a Permission from a Role.	127
	Grant a Role to a User.	128
	Revoke a Role from a User.	129
13.11	Managing Streaming Data Services.	129
	Adding a JDBC Connection to an External Database.	131
	Adding an ODBC Connection to an External Database.	132
	Adding an OCS Connection to an External Database.	133
	Adding a Connection to an R Data Service.	134

	Adding a Connection to an SAP HANA Database.	135
	Updating or Deleting a Data Service.	136
14	Administer a Cluster.	137
14.1	Safeguarding Your Data.	137
	Backing Up SAP HANA with SAP HANA Streaming Analytics.	138
	Recovering SAP HANA with SAP HANA Streaming Analytics.	140
14.2	Memory Usage.	141
14.3	Log Files.	142
	Installation Log Files.	143
	Cluster Log Configuration File.	144
	Project Log Files.	145
	Logging Level.	146
	Locating Log and Trace Files for Troubleshooting Adapters.	147
14.4	Cluster Administrative Tool.	148
15	Copy or Move a Tenant Database.	152
16	Provisioning the Streaming Analytics Service to a Tenant Database.	154
17	Unprovision Streaming Analytics from a Tenant Database.	156
18	Initializing the Connection Between a Tenant Database and a Streaming Service.	158
19	Uninitializing the Connection Between a Tenant Database and a Streaming Service.	159
20	Activating Content in the SAP HANA Repository.	160
20.1	Disabling the Activation of Content in the SAP HANA Repository.	161
21	Troubleshooting.	163
21.1	Understanding Server Messages.	163
21.2	Login and Connection Problems.	164
	Error: Invalid Login Credentials.	165
	Cannot Connect to the Cluster.	165
	A Studio Project Does Not Run, Reports Login Failure.	166
	A Utility Fails to Connect to the Server.	166
	A Utility Fails to Connect to a Project.	167
	An External Adapter Fails to Start.	167
	An Unmanaged External Adapter Fails to Connect to a Project.	168
	ODBC Data Source Name Not Found.	168
	Nodes Are Missing from Cockpit and streamingclusteradmin.	168
21.3	Project Problems.	168
	A Utility Fails to Connect to a Project.	169
	A Project Triggers Java Out-of-Memory Errors.	169
	Published Data Lost When Project Fails.	170

	User Cannot Perform Tasks Despite Being Granted Permissions.	170
	Error: Could Not Allocate Controller.	171
21.4	Data_streaming Directory Already Exists During Installation.	171
21.5	Java Error When Initializing the streamingserver Service.	172
21.6	Cannot Access Streaming Analytics Pages in SAP HANA Cockpit.	172
21.7	Cannot Start Up Streaming Server.	172

1 **SAP HANA Streaming Analytics: Configuration and Administration Guide**

Configure clusters and external database access, deploy adapters and projects to a cluster, and monitor and administer SAP HANA streaming analytics using SAP HANA cockpit.

2 Introduction

Administer SAP HANA streaming analytics as part of your overall SAP HANA landscape.

This guide describes configuration and administration tasks specific to streaming analytics, and discusses overall administration for your SAP HANA system.

For details on administering SAP HANA, see the *SAP HANA Administration Guide*.

In this section:

[SAP HANA Streaming Analytics Overview \[page 8\]](#)

SAP HANA streaming analytics lets you analyze events as they arrive, enabling immediate response to the information.

[Streaming Cluster Nodes \[page 9\]](#)

An SAP HANA streaming analytics cluster consists of one or more nodes, with one node per host. Cluster nodes manage the streaming analytics projects and adapters running on the streaming analytics cluster.

[Tools for Administration and Monitoring \[page 10\]](#)

Use the SAP HANA studio and the SAP HANA cockpit to administer and monitor SAP HANA streaming analytics.

[SAP HANA Streaming Analytics Licenses \[page 10\]](#)

Streaming analytics requires its own license, separate from the SAP HANA license.

Related Information

[SAP HANA Administration Guide](#)

2.1 SAP HANA Streaming Analytics Overview

SAP HANA streaming analytics lets you analyze events as they arrive, enabling immediate response to the information.

You can filter and normalize raw data before capturing the desired data in the SAP HANA database. You can apply event processing logic to monitor the incoming data, generating alerts and notifications—or initiating an immediate automatic response—when specific conditions are detected. Adapters and programming interfaces are available to connect streaming projects to data sources and destinations.

Streaming data models are contained in projects that run on the streaming cluster. One streaming cluster is supported per SAP HANA system (SID). You can run multiple projects simultaneously, providing resiliency. The streaming cluster runs on one or more dedicated hosts in the SAP HANA landscape. You may add hosts during or after installation, using the SAP HANA Lifecycle Management (LCM) tools.

You can install SAP HANA and streaming analytics on a dedicated server (referred to as a `dedicated host deployment`), or you can install them on the same server (referred to as a `same host deployment`).

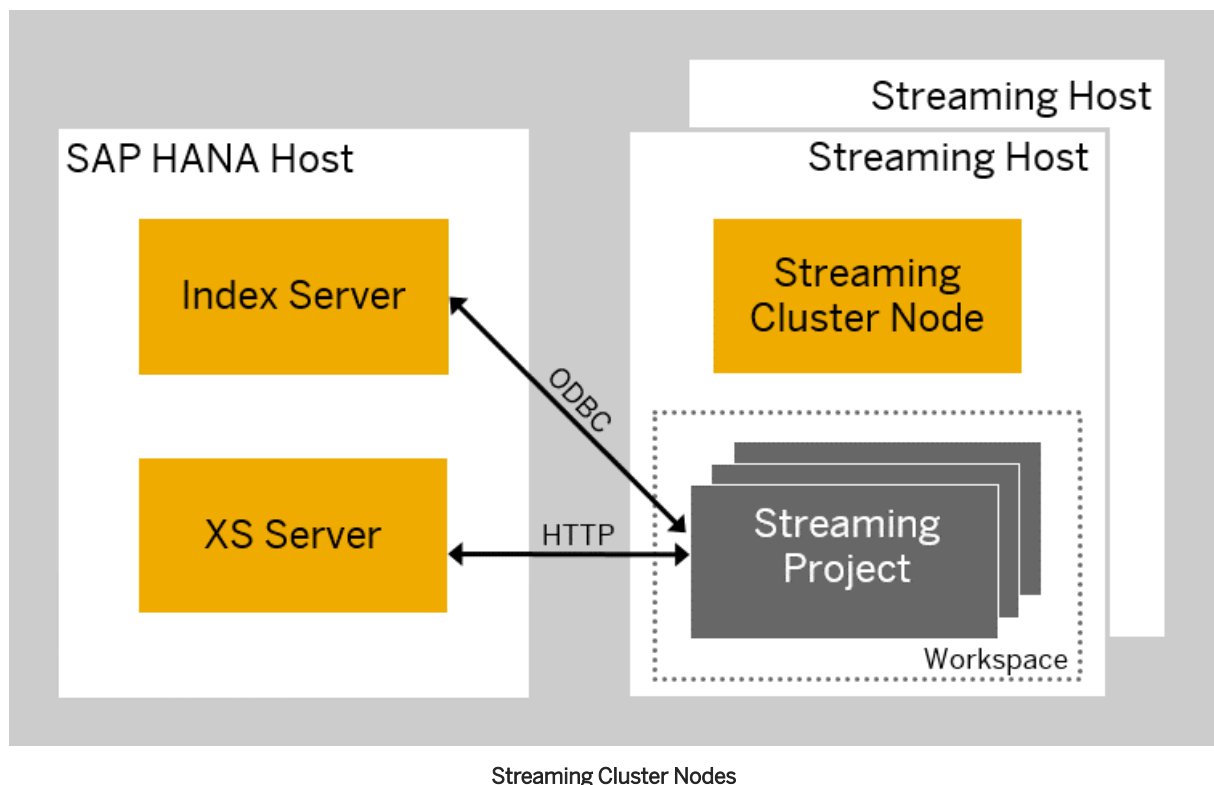
2.2 Streaming Cluster Nodes

An SAP HANA streaming analytics cluster consists of one or more nodes, with one node per host. Cluster nodes manage the streaming analytics projects and adapters running on the streaming analytics cluster.

A streaming cluster can run on one or more hosts. If a project fails, the single or multihost cluster node automatically attempts to restart the project. A multihost cluster provides several additional benefits:

- Eliminates the single point of failure risk
- Provides recovery if a project or node fails
- Distributes the processing load across multiple hosts

In a multihost cluster, streaming cluster nodes run on each host as a network of peers, eliminating a single point of failure. The cluster remains functional as long as at least one cluster node is running. To learn how to add a new streaming host, see the *SAP HANA Streaming Analytics: Installation and Update Guide* for detailed instructions.



Related Information

[SAP HANA Streaming Analytics: Installation and Update Guide](#)

2.3 Tools for Administration and Monitoring

Use the SAP HANA studio and the SAP HANA cockpit to administer and monitor SAP HANA streaming analytics.

The SAP HANA cockpit is an SAP Fiori Launchpad site that provides you with a single point-of-access to a range of Web-based applications for the administration of SAP HANA. The SAP HANA cockpit provides an enhanced view of your streaming information within the context of your SAP HANA system. Role-based security in SAP HANA cockpit ensures that users only have access to pages for which they are authorized. For an overview, see *SAP HANA Cockpit* in the *SAP HANA Administration Guide*. For details on streaming analytics cockpit links, see [Managing Streaming Analytics with SAP HANA Cockpit \[page 77\]](#).

The SAP HANA studio runs on the Eclipse platform and is both the central development environment and the main administration tool for SAP HANA. Use the SAP HANA studio to start and stop services, to monitor the system, to configure system settings, and to manage users and authorizations. The SAP HANA studio presents its various tools in the form of perspectives. Administration and monitoring features are available primarily within the SAP HANA Administration Console perspective. For an overview, see *SAP HANA Studio* in the *SAP HANA Administration Guide*.

Related Information

[SAP HANA Administration Guide](#)

2.4 SAP HANA Streaming Analytics Licenses

Streaming analytics requires its own license, separate from the SAP HANA license.

You need only one SAP HANA streaming analytics license per installation of streaming analytics, regardless of the number of streaming hosts within the single SAP HANA system. All tenant databases running streaming analytics share the single streaming analytics license.

Consistent with SAP HANA, a new installation of streaming analytics includes a temporary license allowing streaming analytics to run without a permanent license for 90 days. After 90 days, a permanent license is required. If a temporary license expires without a permanent license being installed, streaming analytics ceases to function until you install a valid permanent license. If an installed permanent license expires, a temporary license is issued, valid for 28 days

Depending on whether SAP HANA and streaming analytics were installed together or at different times, the expiration dates of the individual temporary licenses may differ.

For more information on license types, checking the current license key, and installing or deleting a permanent license, see *System Administration > Managing SAP HANA Licenses* in the *SAP HANA Administration Guide*.

Related Information

[Managing SAP HANA Licenses](#)

3 Internal and External Ports for SAP HANA Streaming Analytics

Learn which ports you can assign (or be assigned) for different components based on your system.

For default (single) tenant databases, internal and external ports are fixed values that are automatically assigned. For multiple tenant systems, you can still use the previous default values for one of your tenants but can also choose from six other port ranges for your other tenants. If you specify 3xx16 as the internal port of one of your tenants, the other default values are automatically used for the rest of the ports (external, Web Services, Streaming Web Service). If you don't specify the previous default value or don't specify a port number at all, one of the other ports are automatically assigned based on availability. Once they're assigned, the port numbers remain the same and don't change when you start and stop the streaming analytics server.

The default port ranges for multiple tenants are 3xx40 to 3xx99, with nine ports reserved for each range. For example, one range is 3xx40 to 3xx48, the next 3xx49 to 3xx57, and so on. This is similar to other SAP HANA services, the only difference being that streaming analytics reserves nine ports while other services reserve three. With these ports, seven tenants can run on the same SAP HANA host. If you would like to have additional tenants on the same host, you can reserve additional ports by editing the `[multidb] reserved_instance_numbers` property in the `global.ini` file.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='streamingserver' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view. The external port number is always one number higher than the internal port. For example, if the internal port is 3xx40 then the external port is 3xx41.

Port Use	Values Used in Default (Single) Tenant Systems	Default Values in Multiple Tenant Systems
Internal	3xx16	3xx40, 3xx49, 3xx58, 3xx67, 3xx76, and 3xx85
External	3xx26	3xx41, 3xx50, 3xx59, 3xx68, 3xx77, and 3xx86
Web Services Provider (SOAP port)	9090	3xx42, 3xx51, 3xx60, 3xx69, 3xx78, and 3xx87
Web Services Provider (REST port)	9091	3xx43, 3xx52, 3xx61, 3xx70, 3xx79, and 3xx88
Web Services Provider (WebSocket port)	9092	3xx44, 3xx53, 3xx62, 3xx71, 3xx80, and 3xx89
Streaming Web Service	9093	3xx45, 3xx54, 3xx63, 3xx72, 3xx81, and 3xx90

4 Configuring SAP HANA Streaming Analytics for the First Time

After installation, some configuration is required before you can use SAP HANA streaming analytics.

1. Add a streaming host.
Each cluster node requires one streaming host. If you did not add a streaming host during installation, or need to add additional hosts, see the *SAP HANA Streaming Analytics: Installation and Update Guide* for detailed instructions.
2. (Optional; for Linux only) Migrate unixODBC connections.
Beginning with version SPS 10, the SAP HANA streaming analytics installation package includes a unixODBC driver manager. If you are using Linux and have an existing ODBC driver manager set up, migrate the old `.odbc.ini` file to the `$HOME` directory of the user running streaming analytics.
3. (Required for Linux users with only SAP HANA studio installed) Configure an external unixODBC driver manager, including any proxy settings. See [Configuring an External unixODBC Driver Manager on Linux \[page 14\]](#) and [Configuring an ODBC Driver Manager for a Streaming Host with a Proxy \[page 16\]](#).
4. Create users and assign permissions.
The SAP HANA user **SYSTEM**, by default, has the permissions to deploy and run a project on the streaming cluster.
Use the Manage Streaming Permissions link in the SAP HANA cockpit to manage roles and user and role permissions. For information on how to manage user authorization from the SAP HANA cockpit, see [Managing Streaming Permissions \[page 118\]](#).
You can also use the `streamingclusteradmin` utility from the terminal to grant these permissions to other users. For information on granting permissions using this utility, see the [User Authorization Policies](#) section of the *SAP HANA Streaming Analytics: Security Guide*.
5. (Required for running sample projects) Migrate the streaming analytics example data from `STREAMING_HOME/examples/ccl/exampledata` to your workspace, or the workspace of any user that requires it.
6. Test your setup by loading, compiling, and running a test project.
For a guided explanation, you can follow the tutorial in [Appendix: Tutorial for Building and Testing a Project in Studio](#) in the *SAP HANA Streaming Analytics: Developer Guide*. For a guided explanation, you can follow the tutorial in [Appendix: Tutorial for Building and Testing a Project in Studio](#) in the *SAP HANA Streaming Analytics: Developer Guide*.

In this section:

[Configuring an External unixODBC Driver Manager on Linux \[page 14\]](#)

You can use ODBC to connect to streaming analytics from a machine where streaming analytics is not installed. For ODBC connections from Linux, install and configure a unixODBC driver manager.

[Authorization for Administration Tasks \[page 17\]](#)

Administration tasks require certain privileges and roles, which are granted by a user administrator.

[Configure Web IDE in XS Advanced Cockpit \[page 20\]](#)

Perform these configuration tasks after installing SAP Web IDE to begin using SAP HANA streaming analytics.

Related Information

[SAP HANA Streaming Analytics: Installation and Update Guide](#)
[User Authorization Policies](#)
[Appendix: Tutorial for Building and Testing a Project in Studio](#)

4.1 Configuring an External unixODBC Driver Manager on Linux

You can use ODBC to connect to streaming analytics from a machine where streaming analytics is not installed. For ODBC connections from Linux, install and configure a unixODBC driver manager.

Prerequisites

Install the SAP HANA client. See the [SAP HANA Client Installation and Update Guide](#) for instructions.

Context

ODBC connections may be necessary for streaming SDK applications, other client applications, and smart data access. You don't need to register the ODBC driver on Linux. As well, the ODBC driver manager (unixODBC 2.3.1) is included with the streaming analytics server installation package, so in most cases, you don't need to set up your own. However, if your client application is on an external system, you need to download and set up an external unixODBC driver manager.

Procedure

1. Download and install the unixODBC driver manager, which is available from your Linux provider, or from <http://www.unixodbc.org/download.html>.
2. Navigate to the location of the unixODBC driver manager. The default location is `/usr/local/lib`.
3. Locate `libodbc.so.1` in the default folder. If `libodbc.so.1` does not exist, execute the following to create a symbolic link from `libodbc.so.1` to `libodbc.so.2` (or to `libodbc.so.3`, depending on which version of unixODBC you installed):

```
ln -s libodbc.so.2.0.0 libodbc.so.1
```

Streaming analytics connects to `libodbc.so.1`, which unixODBC 2.3.1 has renamed `libodbc.so.2`. Creating the symbolic link allows streaming analytics to use `libodbc.so.2` (or `libodbc.so.3`).

4. Create the `.odbc.ini` file at `/usr/sap/<SID>`.
5. In `/hana/shared/<SID>/streaming/STREAMING.sh`, update `LD_LIBRARY_PATH` to:

```
LD_LIBRARY_PATH="/usr/local/lib":$LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

`LD_LIBRARY_PATH` points to where `libodbc.so.1` is available. In this example it is located at `/usr/sap/home/unixodbc2.3.1/lib`. If using `libodbc.so.2` or `libodbc.so.3`, create a file called `libodbc.so.1`, link it to `libodbc.so.2` (or `libodbc.so.3`), and set `LD_LIBRARY_PATH` to:

```
LD_LIBRARY_PATH=/usr/sap/home/unixodbc2.3.1/lib:$LD_LIBRARY_PATH export
LD_LIBRARY_PATH
```

6. In the `/usr/local/lib` directory, locate `.odbcinst.ini`.
7. Add an entry in the **Driver** parameter for your SAP HANA ODBC driver:

```
[HDBODBC]
Description=64-bit HANA ODBC Driver
Driver=<install-location>/hdbclient/libodbcHDB.so
FileUsage=1
```

8. Add an entry in the **Driver** parameter for your streaming analytics ODBC driver. For example:

```
[STREAMING1]
Driver=<install-location>/libstreamingpsqlodbca_lib.so
Database=<streaming-workspace-name>/<streaming-project-name>
ServerName=<machine-name>
Port=3<instance-number><internal-port>
SSLMode=enable
```

The streaming analytics service runs on one of seven reserved internal ports: 3xx16 (default), 3xx40, 3xx49, 3xx58, 3xx67, 3xx76, or 3xx85, where xx is the instance number. If not specified, the default, or first available port is used. If you would like to have additional tenants on the same host, you can reserve additional ports by editing the `[multidb] reserved_instance_numbers` property in the `global.ini` file.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='streamingserver' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view. The external port number is always one number higher than the internal port. For example, if the internal port is 3xx40 then the external port is 3xx41.

9. Log in to the streaming host as the `<sid>adm` user, and execute:

```
HDB stop
HDB start
```

i Note

Do not kill or stop the streaming analytics service from SAP HANA studio; these are not valid alternatives to using `HDB stop` and `HDB start`.

Next Steps

If your streaming analytics host has a proxy set up, see [Configuring an ODBC Driver Manager for a Streaming Host with a Proxy \[page 16\]](#) for further configuration steps.

In this section:

[Configuring an ODBC Driver Manager for a Streaming Host with a Proxy \[page 16\]](#)

If your streaming analytics host has a proxy set up, perform additional configuration tasks.

4.1.1 Configuring an ODBC Driver Manager for a Streaming Host with a Proxy

If your streaming analytics host has a proxy set up, perform additional configuration tasks.

Prerequisites

- Install the SAP HANA client. See the [SAP HANA Client Installation and Update Guide](#) for instructions.
- If you are using SAP HANA studio on Linux, and have no other components of streaming analytics installed on the machine, install and configure the unixODBC driver manager, version 2.3.1. See [Configuring an External unixODBC Driver Manager on Linux \[page 14\]](#) for more information.

Procedure

1. Navigate to `/hana/shared/<SID>/streaming` and open `STREAMING.sh`.
2. Update `no_proxy`:

```
export no_proxy=<last-two-words-of-domain-name>,$no_proxy
```

3. Log in to the streaming host as the `<sid>adm` user, and issue the command:

```
HDB stop
```

4. Kill the `<PID>` for `sapstartsrv`:

```
HDB info  
kill -9 <PID>
```

5. Start the streaming server:

```
HDB start
```


Related Information

[SAP HANA Client Installation and Update Guide](#)

4.2 Authorization for Administration Tasks

Administration tasks require certain privileges and roles, which are granted by a user administrator.

System Monitoring Tasks

Task	Required Authorization
View information in the system views of the SYS schema	System privilege CATALOG READ
View alert information	Object privilege SELECT on the schema _SYS_STATISTICS

System Administration Tasks

Task	Required Authorization
Stop, start, restart systems	Operating system user (<sid>adm) credentials
Create and drop tenant databases	System privilege DATABASE ADMIN
Stop and start tenant databases	i Note These operations can only be performed from the system database
Stop and start database services	System privilege SERVICE ADMIN
Cancel operations	System privilege SESSION ADMIN
Configure system checks: configure e-mail notification, alter alert thresholds, configure check interval	<ul style="list-style-type: none">System privilege INIFILE ADMIN if the statistics server is running as a separate server process in your system, system privilege INIFILE ADMINObject privileges INSERT and UPDATE on the schema _SYS_STATISTICS if the statistics service is running as an embedded service

System Properties Configuration

Task	Required Authorization
Change the properties of a system as defined in the parameters of the configuration (*.ini) files	System privilege INIFILE ADMIN
Change the system-level properties of a tenant database	System privilege DATABASE ADMIN and INIFILE ADMIN

i Note
System-level properties can only be changed in the system database.

License Management

Task	Required Authorization
Check license information	System privilege LICENSE ADMIN
Install and delete license keys	System privilege LICENSE ADMIN

Troubleshooting Tasks

Task	Required Authorization
Configure traces	System privilege TRACE ADMIN
Delete trace files	System privilege TRACE ADMIN
Troubleshoot offline system (including access diagnosis files) using Administrator editor diagnosis mode	Operating system user (<sid>adm) credentials

User and Role Management

Task	Required Authorization
Create, change, and delete database users	System privilege USER ADMIN
Create, change, and delete roles in runtime	System privilege ROLE ADMIN

Task	Required Authorization
Grant and revoke privileges/roles to/from users and roles	To grant privileges and roles requires you have been granted the privilege or role with authorization to grant it to others. To grant privileges on activated repository objects requires you be authorized to execute certain stored procedures. See <i>Prerequisites for Granting and Revoking Privileges and Roles</i> in the <i>SAP HANA Administration Guide</i> .
Configure the password policy and the password blacklist	System privilege INFILE ADMIN and object privileges INSERT and DELETE for either the <code>_SYS_PASSWORD_BLACKLIST</code> table or the <code>_SYS_SECURITY</code> schema.
View procedures and calculation views in the authorization dependency viewer	System privilege CATALOG ADMIN

Additional Security Administration

Task	Required Authorization
Activate and configure auditing	System privilege AUDIT ADMIN or AUDIT OPERATOR or INFILE ADMIN
Create and manage audit policies	System privilege AUDIT ADMIN or AUDIT OPERATOR
Read the audit trail database table <code>SYS.AUDIT_LOG</code>	System privilege AUDIT OPERATOR
Delete entries from the audit trail	System privilege AUDIT OPERATOR
Enable persistence encryption and access the encryption monitoring views <code>M_PERSISTENCE_ENCRYPTION_STATUS</code> and <code>M_PERSISTENCE_ENCRYPTION_KEYS</code>	System privilege RESOURCE ADMIN
Change the page encryption key used for data volume encryption and re-encrypt the data	System privilege RESOURCE ADMIN

For a complete list of authorizations by task, see *SAP HANA Administration Guide*.

Related Information

[SAP HANA Administration Guide](#)

4.3 Configure Web IDE in XS Advanced Cockpit

Perform these configuration tasks after installing SAP Web IDE to begin using SAP HANA streaming analytics.

In this section:

[Set Up the Space \[page 20\]](#)

SAP Web IDE allows developers to use designated spaces for building and running streaming projects. To configure and maintain spaces for different development teams, log on to XS Advanced (XSA) Cockpit.

[Set Up Role Collections \[page 21\]](#)

SAP Web IDE supplies predefined role templates which can be assigned to role collections. Assigning the templates to these collections will grant the relevant Web IDE permissions to the users in the respective roles for creating streaming projects.

[Enable Spaces for Development \[page 23\]](#)

Enabling spaces for development by installing the builder in each space allows you to build the projects you create.

[Configure the SAP HANA Streaming Analytics Service Broker \[page 23\]](#)

Before you can use the SAP HANA streaming analytics plugin in the SAP Web IDE, you need to register any associated tenant databases through the streaming analytics service broker.

4.3.1 Set Up the Space

SAP Web IDE allows developers to use designated spaces for building and running streaming projects. To configure and maintain spaces for different development teams, log on to XS Advanced (XSA) Cockpit.

Prerequisites

- XS Advanced (XSA), streaming analytics (with the streaming analytics plugin for SAP Web IDE), and the streaming service broker are installed. For details, see *Ensure that SAP HANA Streaming Analytics Works with Web IDE* in the *SAP HANA Streaming Analytics: Installation and Update Guide*.

Procedure

1. Log on to the XS Advanced Cockpit.
2. Select **Organizations** on the navigation bar.
3. Select your desired organization.
4. Select the PROD space.
5. Select **Members** on the navigation bar.

6. Select **Add Members**.
7. Enter XSA_ADMIN under User IDs.
8. Check **Space Developer**.
9. Click **OK** to save.

Next Steps

Configure the role collections and assign them to users. See *Set up the Role Collections* in the *SAP HANA Streaming Analytics: Configuration and Administration Guide*.

Related Information

[Ensure that SAP HANA Streaming Analytics Works with Web IDE Set Up Role Collections \[page 21\]](#)

4.3.2 Set Up Role Collections

SAP Web IDE supplies predefined role templates which can be assigned to role collections. Assigning the templates to these collections will grant the relevant Web IDE permissions to the users in the respective roles for creating streaming projects.

Prerequisites

- XS Advanced (XSA), streaming analytics (with the streaming analytics plugin for SAP Web IDE), and the streaming service broker are installed. For details, see *Ensure that SAP HANA Streaming Analytics Works with Web IDE* in the *SAP HANA Streaming Analytics: Installation and Update Guide*.

Procedure

1. Log on to the XS Advanced Cockpit.
2. Select **Security > Role Collections**.
3. Create a role collection for Web IDE administration:
 - a. Select **New Role Collection**.
 - b. Enter a name for a collection and click **Save**.
 - c. Select the new collection.

- d. Select **Add Role**.
 - e. For Application Identifier, select **webide!i1**.
 - f. For Role Template and Role, select **WebIDE_Administrator**.
 - g. Select **Save**.
4. Create a role collection for Web IDE development:
 - a. Select **New Role Collection**.
 - b. Enter a name for a collection and click **Save**.
 - c. Select the new collection.
 - d. Add the Web IDE developer role:
 1. Select **Add Role**.
 2. For Application Identifier, select **webide!i1**.
 3. For Role Template and Role, select **WebIDE_Developer**.
 - e. Add the SRTT developer role:
 1. Select **Add Role**.
 2. For Application Identifier, select **com-sap-xsa-sds-rtt**.
 3. For Role Template and Role, select **xsa_sds_rtt_developer_template**.
 5. Return to the XSA Cockpit homepage.
 6. Select **User Management**.
 7. Assign the administration role collection to XSA_ADMIN:
 - a. Under Actions, select **Assign Role Collections** for XSA_ADMIN.
 - b. Click **Add**.
 - c. Select your administration role collection and click **OK**.
 - d. Click **Save**.
 8. For each developer user, assign the developer role collection:
 - a. Under Actions, select **Assign Role Collections** for the user.
 - b. Click **Add**.
 - c. Select your developer role collection and XS_CONTROLLER_USER and click **OK**.
 - d. Click **Save**.

Next Steps

Enable the space to complete the configuration of the spaces you created.

Related Information

[Ensure that SAP HANA Streaming Analytics Works with Web IDE](#)

[Configure the SAP HANA Streaming Analytics Service Broker \[page 23\]](#)

[Enable Spaces for Development \[page 23\]](#)

4.3.3 Enable Spaces for Development

Enabling spaces for development by installing the builder in each space allows you to build the projects you create.

Prerequisites

- The XSA_ADMIN user has the Space Developer role for the PROD space.
- The XSA_ADMIN user has the WEBIDE_Administrator role.

Procedure

1. As the <sid>adm user, log in to XSA services:

```
xs-admin-login
```

2. Get the URL for the space enablement tool:

```
xs app di-space-enablement-ui --urls
```

3. In a browser, open the URL and log in as XSA_ADMIN.
4. For the PROD space, click **enable**.

Related Information

[Set Up Role Collections \[page 21\]](#)

4.3.4 Configure the SAP HANA Streaming Analytics Service Broker

Before you can use the SAP HANA streaming analytics plugin in the SAP Web IDE, you need to register any associated tenant databases through the streaming analytics service broker.

Prerequisites

- XS Advanced (XSA), streaming analytics (with the streaming analytics plugin for SAP Web IDE), and the streaming service broker are installed. For details, see *Ensure that SAP HANA Streaming Analytics Works with Web IDE* in the *SAP HANA Streaming Analytics: Installation and Update Guide*.

- You have access to the tenant database SYSTEM user, or another tenant database user with the SAP HANA role XSA_SUPPORT_VIEW, the CATALOG READ privilege, and the CREATE SCHEMA privilege.

i Note

If you have already configured the service broker and need to change the credentials it uses, the user needs to have the SAP HANA role XSA_SUPPORT_VIEW, the CATALOG READ privilege, and the SYS_STREAMING schema with the object privileges CREATE_ANY, SELECT, INSERT, UPDATE, and DELETE.

- You have the XSA application role collection XS_CONTROLLER_ADMIN.
- You have enabled the tenant database for XSA.
- If the tenant database you are using is not the default tenant database:
 - You have enabled the additional tenant database for XSA.
 - You have mapped the additional tenant database to an XSA space or organization.

Context

The streaming service broker connects streaming servers associated with tenant databases to XSA spaces, and allows you to build streaming modules through SAP Web IDE. Then, you can manage streaming analytics projects through the streaming analytics runtime tool.

Procedure

1. Log on to the XS Advanced Cockpit.
2. In the home navigation pane, choose **More...**
3. In the **SAP HANA Streaming Analytics** section, select **Service Broker Configuration**.
4. Open the **HANA System Provisioning** tab and click **Edit**.
5. Fill in the login information for the SAP HANA system database, then click **Save**.

The XSA controller fills in the host name and instance number by default, but you can also change these values to add your own SAP HANA system instead.

6. (Only applicable to systems migrated from single-container systems) Fill in the login information for the tenant database. Use either the SYSTEM user, or a user with the XSA_SUPPORT_VIEW role. Then, click **Save**.

You can also change the tenant database name if the default entry is not accurate.

i Note

The tenant database fields only appear if you upgraded your system from 2.0 SP 00 or earlier to 2.0 SP 01 or later. In direct installations of 2.0 SP 01 or later, the tenant connection information gets stored on the system database, and you do not have to enter it manually.

7. Open the **Streaming Service Provisioning** tab.
8. Choose a tenant database and click **Register**.

If the streaming service status is **Not Provisioned**, the service has not been provisioned to a tenant database, or has not been initialized.

9. Enter the streaming cluster manager password when prompted, then click **Save**.

i Note

This is the same password that you create when initializing the streaming service.

You can now use the SAP HANA streaming analytics plugin in the SAP HANA Web IDE for this database.

5 Managing the Streaming Service and Hosts

You can start or stop the streaming service manually, or allow it to auto-restart. As well, you can set up additional streaming hosts or delete existing hosts.

An SAP HANA streaming analytics cluster consists of one or more nodes, with one node per host. Each tenant is associated with only one cluster. By default, up to seven tenants can run on a streaming host (a host that has the streaming role assigned). You can add more, if needed, by reserving additional ports. To do this, edit the `[multidb] reserved_instance_numbers` property in the `global.ini` file.

i Note

Throughout the documentation, we also refer to streaming nodes as streaming servers.

A cluster contains one or more workspaces, each with one or more projects. Nodes run and manage the projects and adapters running on those workspaces. All workspaces are available to the entire cluster. You share a project within the cluster or within a workspace with other users. You can adjust the settings of the clustered projects using SAP HANA cockpit.

In this section:

[Verify that the Streaming Service is Installed and Running \[page 27\]](#)

Adding the streaming analytics host creates and starts the streaming analytics service. Verify that the service is installed and running either through SAP HANA studio or by using the `HDB info` command on the host.

[Start the Streaming Service Manually \[page 28\]](#)

If stopped, you can manually start the streaming service.

[Stop the Streaming Service Manually \[page 28\]](#)

Stop the streaming service to prevent a restart through the auto-restart mechanism.

[Auto-Restart the Streaming Service \[page 29\]](#)

Through the SAP HANA cockpit or the SAP HANA studio Administration Console, you can stop the streaming service, and allow it to restart through the auto-restart mechanism.

5.1 Verify that the Streaming Service is Installed and Running

Adding the streaming analytics host creates and starts the streaming analytics service. Verify that the service is installed and running either through SAP HANA studio or by using the `HDB info` command on the host.

Verify Using HANA Studio

Procedure

1. In the SAP HANA studio Administration console, right-click the database in the Systems panel, select **► Configuration and Monitoring ► Open Administration ►**, then select the **Landscape** tab.
2. Verify that the status of the `streaming analytics` on the streaming analytics host is active (green).

Results

- If the `streaming analytics` service is inactive (red), restart the service.
- If the service is not listed, the streaming analytics worker is not installed. See the *SAP HANA Streaming Analytics: Configuration and Administration Guide* for steps on adding the streaming analytics host.

Verify Using the Streaming Host

Procedure

1. Log on to the streaming analytics host as the `<sid>adm`, where `<sid>` is the SAP HANA ID in lowercase letters. For example, if the `<sid>` is `XY1`, log in as `xy1adm`.
2. Type the following to display a list of services that are running:

```
HDB info
```

If `hdbstreamingserver` appears on the list, the streaming analytics service is running.

```
rh1adm 13151 13107 0.0 737604 34036 \hdbstreamingserver -  
hdb:instance=01 -hdb:hostname=...
```

Results

- If the `hdbstreamingserver` service is not listed, but you know that the streaming analytics host is installed, restart the service and verify again.
- If you are unsure if the streaming analytics host is installed, use SAP HANA studio to verify that the service is installed and running.

5.2 Start the Streaming Service Manually

If stopped, you can manually start the streaming service.

Procedure

- Use either method:

Option	Description
SAP HANA streaming analytics host	<ol style="list-style-type: none">1. Log on to the SAP HANA streaming analytics host as the system administrator (<code><sid>adm</code>).2. Issue this command: <pre>HDB start</pre>
Any host	<ol style="list-style-type: none">1. Log on to any host in the landscape as the system administrator (<code><sid>adm</code>).2. Issue this command, where <code><instance-number></code> is the SAP HANA instance number and <code><streaming-host></code> is the name of the streaming analytics host: <pre>sapcontrol -nr <instance-number> -function InstanceStart <streaming-host> <instance-number></pre>

5.3 Stop the Streaming Service Manually

Stop the streaming service to prevent a restart through the auto-restart mechanism.

Context

Once stopped, the streaming service can only be restarted manually.

Procedure

- Use either method:

Option	Description
SAP HANA streaming analytics host	<ol style="list-style-type: none">1. Log on to the streaming analytics host as the system administrator (<sid>adm), where <sid> is the SAP HANA system ID, using lowercase letters. For example, if the <sid> is XY1, log in as xy1adm.2. Issue this command: <pre>HDB stop</pre>
Any host	<ol style="list-style-type: none">1. Log on to any host in the landscape as the system administrator (<sid>adm).2. Issue this command, where <instance-number> is the SAP HANA instance number and <streaming-host> is the name of the streaming analytics host: <pre>sapcontrol -nr <instance-number> -function InstanceStop <streaming-host> <instance-number></pre>

5.4 Auto-Restart the Streaming Service

Through the SAP HANA cockpit or the SAP HANA studio Administration Console, you can stop the streaming service, and allow it to restart through the auto-restart mechanism.

Stop-Restart Using SAP HANA cockpit

Prerequisites

- You have the SERVICE ADMIN system privilege.
- You have the SELECT privilege on the _SYS_STATISTICS schema.

Context

i Note

The SAP HANA auto-restart feature supports high availability. If an SAP HANA service is disabled by failure or intentional intervention, this auto-restart function detects the failure and restarts the stopped service process.

Procedure

1. In SAP HANA cockpit, click **Overall Database Status**.
2. Highlight the `hdbstreamingserver` service and click **Stop Service**.
3. Click **Yes** to confirm. A success message appears.

Stop-Restart Using SAP HANA Studio

Prerequisites

- You have the SERVICE ADMIN system privilege.
- You are logged on to the SAP HANA database in the Administration Console.

Context

i Note

The SAP HANA auto-restart feature supports high availability. If an SAP HANA service is disabled by failure or intentional intervention, this auto-restart function detects the failure and restarts the stopped service process.

Procedure

1. In the SAP HANA studio Administration console, right-click the database in the Systems panel, select **► Configuration and Monitoring ► Open Administration ▾**, then select the **Landscape** tab.
2. Right-click the streaming service and choose the required option:

Option	Description
Stop...	The service stops and then typically restarts.
Kill...	The service stops immediately and then typically restarts.

6 Cluster Configuration

The SAP HANA database maintains configuration information for all nodes in the cluster.

The SAP HANA database also holds policy configuration, and persists cluster cache data. There is one cluster per SAP HANA system.

In streaming analytics most of the cluster and node settings are configured automatically. However, there are a few settings that you can modify using the SAP HANA cockpit. The SAP HANA cockpit is a Web-based tool that allows you to manage streaming projects, adapters, and other components. For more information, see [Managing Streaming Analytics with SAP HANA Cockpit \[page 77\]](#).

In this section:

[File and Directory Infrastructure \[page 31\]](#)

Use configuration files and directories to manage project deployment, security and options like failover, affinities, and high availability.

6.1 File and Directory Infrastructure

Use configuration files and directories to manage project deployment, security and options like failover, affinities, and high availability.

Streaming file locations are relative to the SAP HANA installation, for example `/hana/shared/<SID>/streaming/STREAMING-2_0`.

The following table describes the various SAP HANA streaming analytics files and directories:

Name	Description
<code><adapter-base-directory></code>	Adapters are sandboxed. All paths provided by users are placed relative to the adapter sandbox base directory <code>/adapters/</code> .

Name	Description
sandbox-base-directory	<p>Default is <code>STREAMING_SHARED</code>. A configuration file macro for resources required to be on a shared drive that cannot be modified.</p> <p><code>STREAMING_SHARED</code> is the base sandbox directory that each project and adapter is contained within. From there, adapters can only read and write files within their own folder under the <code>STREAMING_SHARED/<tenant-db-name>/adapters/</code> directory. Put files used by your adapters into the correct directory to ensure that the adapter can access them. If using schema discovery, set your adapter's <code>dir</code> property to use an absolute path (<code>STREAMING_SHARED/<tenant-db-name>/adapters</code>) during discovery, and then a relative path (relative to the sandbox directory) at runtime. Your project will not run unless you set a relative path before you start it.</p> <p>Location</p> <p style="text-align: right;"><code>/hana/data_streaming/<SID>/</code></p>
cluster.key	<p>Used to encrypt SSL files and passwords in cluster configuration.</p> <p>Location</p> <p style="text-align: right;"><code>/hana/shared/<SID>/streaming/cluster/<db-name>/keys/</code></p>
cluster_log4j2.xml	<p>Cluster node log configuration file. <code>cluster_log4j2.xml</code> is a log4j configuration file.</p> <p>Location</p> <p style="text-align: right;"><code>STREAMING_HOME/./cluster/<db-name>/config/cluster_log4j2.xml</code>, where <code><db-name></code> is the lower case name of the tenant database.</p>
<code>STREAMING_HOME</code>	<p>Represents the SAP HANA streaming analytics installation directory. <code>STREAMING_HOME</code> is also an environment variable and a configuration file macro.</p> <p>Location</p> <p style="text-align: right;"><code>/hana/shared/<SID>/streaming/STREAMING-2_0</code></p>
<code>STREAMING_SHARED</code>	<p>A configuration file macro for resources required to be on a shared drive that cannot be modified.</p> <p>Location</p> <p style="text-align: right;"><code>/hana/data_streaming/<SID>/</code></p>

Name	Description
STREAMING_CUSTOM_ADAPTERS_HOME	<p>Represents the SAP HANA streaming analytics custom adapters directory. STREAMING_CUSTOM_ADAPTERS_HOME is also an environment variable and a configuration file macro.</p> <p>Needs a shared drive?</p> <p>Yes</p> <p>Location</p> <p>STREAMING_CUSTOM_ADAPTERS_HOME is not set by default. For information about setting this environment variable, see the Directories for Custom Adapters topic in the <i>SAP HANA Streaming Analytics: Building Custom Adapters</i>.</p>
Input files for projects	<p>Resources a project requires at runtime—for example, input CSV files.</p> <p>Recommended location</p> <p>Working directory for the project:</p> <pre data-bbox="922 1003 1299 1093"><project-base-directory>/ <workspace-name>.<project-name>.<instance-number></pre> <p>Input files for adapters are placed relative to the adapter sandbox base directory <code>/adapters/</code>.</p>
Log store	<p>Saves project input and output data in the event of failover.</p> <p>Location</p> <pre data-bbox="922 1303 1315 1429">STREAMING_SHARED/<db-name>/ store/<filename-property>/ <workspace-name>.<project-name>.<instance-number></pre>
<project-base-directory>	<p>Parent of the working directories for project processes. Also contains cluster workspaces, project instances, and log stores. The names of project working directories take the form:</p> <pre data-bbox="735 1572 1394 1639"><workspace-name>.<project-name>.<instance-number></pre> <p>where <instance-number> can be 0 or 1 if hot standby is enabled, and is always 0 if hot standby is disabled.</p> <p>Location</p> <pre data-bbox="922 1787 1283 1841">STREAMING_SHARED/cluster/ projects/<db-name></pre>

Name	Description
<project-name>.ccl	<p>Definition of a project in Continuous Computation Language (.ccl). A project has one <project-name>.ccl file. Editing the project in the SAP HANA streaming analytics plugin for SAP HANA studio modifies the CCL or an ancillary file. Ancillary files store information about the visual presentation of the project in SAP HANA studio.</p> <p>Location</p> <p>Local (studio) cluster: <user's-home-dir>/hdbstudio/workspace/<project-name>/<project-name>.ccl</p>
<project-name>.ccr	<p>The project configuration (.ccr) file is an XML file that contains runtime and deployment configuration for the project, including adapter, parameter, and binding definitions. A project can have multiple CCR files allowing you to specify different adapters in different deployment environments, or to change runtime parameter values.</p> <p>Location</p> <p>Local (studio) cluster: <user's-home-dir>/hdbstudio/<project-name>/<project-name>.ccr</p>
<project-name>.ccx	<p>The compiled .ccl file. You must specify the path to the CCX file when you deploy a project to a remote (non-studio) cluster. You do not need to view or edit the CCX file.</p> <p>Location</p> <p>Local (studio) cluster: <user's-home-dir>/hdbstudio/<project-name>/<project-name>.ccx</p>
Project working directory	<p>Contains files generated by the project and any input files required by the project.</p> <p>Location</p> <p>STREAMING_SHARED/cluster/projects/<db-name>/<workspace-name>.<project-name>.<instance-number></p>

Name	Description
server.key server.crt	<p>Unencrypted SSL files. The streaming analytics installer only creates the encrypted (.enc) versions of the SSL files.</p> <div data-bbox="735 421 1394 533" style="background-color: #f0f0f0; padding: 5px;"> <p>i Note The Streaming Web Service requires unencrypted SSL files.</p> </div> <p>Location</p> <pre data-bbox="922 589 1326 685">/hana/shared/<SID>/ streaming/cluster/<db-name>/ keys/</pre>
server.key.enc server.crt.enc	<p>Encrypted SSL files. Streaming analytics requires a set of encrypted SSL files (by default) or a set of unencrypted SSL files, but not both.</p> <p>Location</p> <pre data-bbox="922 864 1326 960">/hana/shared/<SID>/ streaming/cluster/<db-name>/ keys</pre>
streamingserver_<machine hostname>.<5 digit cluster node's rpc port>.<3 digit log serial number>.log	<p>Cluster node log. Captures node-level and cluster-level errors and event messages; each node's log is unique to that node and resides in the SAP HANA trace directory. The timezone for this file is the local timezone.</p> <p>Location</p> <pre data-bbox="922 1167 1211 1227">\$SAP_RETRIEVAL_PATH/ trace/DB_<DB-NAME></pre>
streamingserver~<workspace name>.<project name>.<project instance number>~<machine hostname>.<5 digit cluster node's rpc port>.<3 digit log serial number>.trc	<p>Project log. Captures errors and events in a running project. Resides in the SAP HANA trace directory. The timezone for this file is the local timezone.</p> <p>Location</p> <pre data-bbox="922 1406 1211 1467">\$SAP_RETRIEVAL_PATH/ trace/DB_<DB-NAME></pre>
streamingserver~<workspace name>.<project name>.<project instance number>~<machine hostname>.<5 digit cluster node's rpc port>.<3 digit log serial number>.out	<p>Standard streams log. Captures output written to stdout and stderr, including licensing information and messages from third party applications. Resides in the SAP HANA trace directory. The timezone for this file is the local timezone.</p> <p>Location</p> <pre data-bbox="922 1675 1211 1736">\$SAP_RETRIEVAL_PATH/ trace/DB_<DB-NAME></pre>
workspace, cluster (also workspace, streaming)	<p>A named scope (similar to a directory) to which you deploy projects and their supporting files, adapters, and data services. Cluster workspaces provide the option to manage the permissions of related objects together. Every cluster has at least one workspace, and a workspace can contain projects, supporting files, adapters, and data services from multiple nodes.</p>

7 High Availability

SAP HANA streaming analytics supports high availability features that promote failure recovery and data redundancy.

SAP HANA index server With multiple SAP HANA index servers configured in your SAP HANA landscape, you can configure streaming analytics to support index server failover. Update the `jdbc-url` property in the `cluster.cfg` file to include additional index servers separated by semicolons:

```
jdbc-url=jdbc:sap://indexserver1:3<instance-number>15;indexserver2:3<instance-number>15
```

Set this feature in the cluster bootstrap file located at `/hana/shared/<SID>/streaming/cluster/<db-name>/config/cluster.cfg`.

Multinode Clustering A multinode cluster of two or more nodes can protect against node failure. When a server in a node fails, the projects running on the failed node restart on another node if their affinities are configured.

To add multiple nodes, repeat the procedure for adding a host.

Warm standby (Recommended) We recommend using a multinode cluster and enabling warm standby for your projects. In warm standby, the cluster starts a single copy of the project and monitors its health. If the cluster manager detects that a project is no longer running, it restarts the project on the same or a different node based on availability. All connections are redirected to the new project instance.

If the project uses windows attached to log stores, it will recover any data from the log stores once it restarts. Note that for data recovery to work in a multinode cluster, the log store needs to reside on disk space that is shared by all nodes within the cluster.

Set this feature in the project configuration.

Hot Standby When you deploy a project in hot standby mode, two instances of the same project run in the cluster, preferably on separate machines. One version of the project is designated the primary instance, and the other is designated the secondary instance. All connections from outside the cluster (adapters, clients, studio) are directed to the primary project server. If the primary instance fails, a hot failover occurs and all connections are automatically directed to the secondary instance.

Data between primary and secondary instances is continuously synchronized. The primary instance receives each message first. To maintain redundancy, the secondary instance must acknowledge receipt of the message before the primary instance begins processing.

The main advantage of hot standby is that regardless of the size of the project and amount of data, the secondary project takes over from the primary in a matter of seconds. However, there are also several limitations (see below) you should consider

before choosing between warm or hot standby. For most use cases, we recommend enabling warm standby instead.

- Because the primary and secondary projects run in parallel, they may receive input data in a different order. Therefore, the two projects may produce different results, especially in the case of complex projects where data can flow down different paths and join later. If you need the primary and secondary projects to produce identical data, do not use this mode.
- You can't use guaranteed delivery (GD) because the sequence numbers that correspond to a logical event may not be the same between the primary and secondary projects due to the issue described above. It can't be guaranteed that the GD subscriber has received all events when the secondary project takes over.
- You can't use internal (or external) managed input or output adapters because the same adapters run in both the primary and secondary projects. For output adapters, this results in duplicate rows in the target and missing rows for input adapters. In this mode you can only use adapters that start and stop outside of the project, that is unmanaged or cluster managed adapters.

Set this feature in the project configuration.

Zero Data Loss

Zero data loss features—guaranteed delivery, consistent recovery, and auto checkpoint—protect against data loss in the event of a server crash or loss of connection.

Guaranteed Delivery

Guaranteed delivery (GD) uses log stores to ensure that a GD subscriber registered with a GD stream or window receives all the data processed by that stream or window even if the client is not connected when the data is produced. GD is supported on streams and windows and each GD stream or window requires a GD log store.

Set this feature through Windows properties in studio; adapter cnxml files; client applications (via SDKs); and binding parameters in CCR files.

Consistent Recovery

Recovery consistency depends on guidelines for log stores:

- When consistent recovery is enabled, the server uses coordinated checkpoints to save data in log stores.
- When any log store fails to complete a checkpoint, all the log stores for that project roll back to their state as of the previous successful checkpoint. This rule ensures that even if a server or connection fails, all log stores in a project are consistent with one another. However, any input data that has not been checkpointed is not recovered upon restart.

Set this feature in the project configuration.

Auto Checkpoint

Auto checkpoint lets you control how often log store checkpoints occur across all input streams and windows in the project. More frequent checkpoints mean that less data is lost if the server crashes. At the maximum checkpoint frequency of every input transaction (value of 1), all input data is protected except the data from the last transaction, which might not be

checkpointed before a crash. Here, input transaction refers to one record, one transaction, or one envelope published to the server and currently in flight. This applies to a single publisher; if there are multiple publishers, then the last transaction for every publisher that is in flight may be lost. When you set checkpoint frequency, you make a trade-off: with frequent checkpoints you can reduce the amount of data at risk, but performance and latency may suffer as a result. The alternative is to increase performance but risk a larger amount of data loss by setting infrequent checkpoints.

Set this feature in the project configuration.

Related Information

[Editing Deployment Options in the Project Configuration File](#)
[Zero Data Loss](#)

8 Set Up System Replication

System replication enables recovery from a data center outage by switching to a secondary site. You can use the `hdbsutil` tool or SAP HANA studio to set up system replication between two SAP HANA systems. Since SAP HANA streaming analytics is part of the SAP HANA system landscape, you do not need to separately set up system replication for it. However, there are some additional tasks you need to perform on the streaming analytics side so that replication is supported successfully.

Prerequisites

- Verify that your SAP HANA systems meet the prerequisites listed throughout the *Setting Up System Replication* section in the *SAP HANA Administration Guide*.
- Specify the same worker group values for the hosts on the secondary site that you wish to map to the hosts on the primary site.
- Do not provision any streaming analytics services until after you enable replication.
- Specify the same streaming analytics node names for both the primary and secondary systems. See [Edit the Streaming Node Name \[page 81\]](#) for detailed instructions. To easily identify which host a streaming analytics node runs on, specify the same value for the node names and the worker group.
- Ensure both the primary and secondary systems have the same number of streaming analytics hosts. For example, if you have a host in the primary system that you don't have in the secondary that host won't be available on the secondary system in case of failover.
- The `.ini` file parameters should be the same between the primary and secondary systems. Set the `global.ini/[inifile_checker]/replicate = true` parameter on the primary system to enable automatic replication of the `.ini` file parameters between the systems. The default interval value is 3600 seconds, but if you want to test takeover on your secondary system, consider temporarily setting this to a lower value such as 60 seconds. See *Monitoring INI File Parameter Changes* in the *SAP HANA Administration Guide* for complete details on setting this parameter.
- If you change files that are not stored in the SAP HANA database after you set up system replication, ensure that both the primary and secondary systems have the updated content. Those files are not replicated automatically, therefore, they require manual maintenance. Some of the files include:
 - Cluster log configuration files (`cluster_log4j2.xml`).
 - Adapter configuration files.
 - Any content aside from passwords in the cluster bootstrap file (`cluster.cfg`). Passwords are automatically moved over.
 - Any streaming analytics keys that are not in these default locations:
 - `$(STREAMING_HOME)/../cluster/<db-name>/keys`
 - `$(STREAMING_HOME)/../cluster/<db-name>/adapters/adapter.key`
 - `$(STREAMING_HOME)/wsp/keystore_rsa.jks`
 - `$(STREAMING_HOME)/wsp/wsp.key`Keys within these default locations are automatically moved over during registration only.

- Any DSN entries in the `.odbc.ini` file, which are referenced by data services used by projects.
- Manually update both the primary and secondary systems with any passwords, keys, or node names that you change after configuring system replication.

Context

The following steps outline the general process of setting up system replication between two SAP HANA systems, failing over to a secondary system, and failing back to a primary system. Steps that are unique to streaming analytics are marked as such.

Procedure

1. Set up system replication on primary and secondary systems:
 - a. Start the primary system.
 - b. Create an initial data backup on the primary system.
 - c. Enable system replication on the primary system (`sr_enable`).
 - d. Register the secondary system with the primary system (`sr_register`).
 - e. (Streaming analytics-specific) Make any input files used by your projects available on both the primary and secondary host file systems.
 - f. (Streaming analytics-specific) Ensure your data services point to a valid data source. Otherwise, manually update them after failover and before restarting your projects.
 - g. Start the secondary system.
 - h. On the primary system:
 1. Create the tenant database that will be used to host the streaming server, back it up, and verify that SAP HANA shows its replication status as normal.
 2. As an admin on the system database, add the streaming service to the tenant database. For example, `ALTER DATABASE DB1 ADD 'streamingserver';`
 3. As an admin on the tenant database, initialize the streaming service. For example, `ALTER SYSTEM INITIALIZE SERVICE 'streamingserver' WITH CREDENTIAL TYPE 'PASSWORD' USING 'ABCabc123';`

Note

It is expected for the streaming service to be in a stopped state on the secondary SAP HANA system while replication is in progress. The streaming service will be started normally once takeover is performed.

2. During failover, the secondary system takes over from primary system:
 - a. Secondary system in data center B takes over from primary in data center A (`sr_takeover`).
 - b. Stop primary system in data center A.
 - c. (Streaming analytics-specific) If your data services do not point to a valid data source, manually update them.
 - d. (Streaming analytics-specific) Start your projects on the secondary system.

The projects do not restart automatically or start from where they left off before failover. You have to manually restart all projects.

- e. When the primary system is available again, register it with the secondary system (`sr_register`).

The roles are switched, the original primary is registered as a secondary system. The original secondary is the production system.

- f. Start the system in data center A.

3. Failback to the original primary system:

- a. Send a takeover command from the system in data center A (`sr_takeover`).

- b. Stop the system in data center B.

- c. Register the system in data center B as secondary again (`sr_register`).

- d. Start the system in data center B.

In this section:

[Configure the Primary and Secondary Systems \[page 41\]](#)

Use `hdbnsutil` to configure the primary and secondary SAP HANA systems that have SAP HANA streaming analytics installed.

[Modify the Assigned Worker Group \[page 42\]](#)

Update the worker group value assigned during installation, using SAP HANA studio or SQL.

[Enabling the Activation of Content in the SAP HANA Repository \[page 43\]](#)

If you enabled the activation plugin on your primary site but failover occurred and now your secondary site has taken over as your new primary, you'll need to make a few updates to the plugin. If you did not enable the activation plugin on your original primary site, skip this topic and only perform the steps outlined in [Activating Content in the SAP HANA Repository \[page 160\]](#) for your new primary site.

Related Information

[SAP HANA Administration Guide](#)

8.1 Configure the Primary and Secondary Systems

Use `hdbnsutil` to configure the primary and secondary SAP HANA systems that have SAP HANA streaming analytics installed.

Context

There are no unique configuration steps to perform on the streaming analytics side.

Follow the instructions in the [SAP HANA System Replication](#) section of the *SAP HANA Administration Guide*.

8.2 Modify the Assigned Worker Group

Update the worker group value assigned during installation, using SAP HANA studio or SQL.

Modify Using SAP HANA Studio

Prerequisites

You have the INIFILE ADMIN system privilege.

Procedure

1. In the SAP HANA Studio Administration console, double-click the system in the Systems panel, and then choose the **Configuration** tab.
A list of all configuration files appears.
2. Expand ► **nameserver.ini** ► **landscape** ▾.
3. Right-click `workergroups_<hostname>`, and choose **Change...**
4. Enter the new value.

Modify Using SQL Console

Prerequisites

You have the INIFILE ADMIN system privilege.

Procedure

1. In a SQL console, connect to the system to be modified.
2. Do one of:
 - Run the built-in procedure `SYS.UPDATE_LANDSCAPE_CONFIGURATION('SET WORKERGROUPS', '<hostname>', '<workergroup_name>')`
 - Execute the ALTER SYSTEM ALTER CONFIGURATION statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('nameserver.ini', 'SYSTEM')
SET ('landscape', 'workergroups_<hostname>') = '<workergroup_name>'
```

where:

- `<hostname>` is the streaming host machine name.
- `<workergroup_name>` is the worker group name.

8.3 Enabling the Activation of Content in the SAP HANA Repository

If you enabled the activation plugin on your primary site but failover occurred and now your secondary site has taken over as your new primary, you'll need to make a few updates to the plugin. If you did not enable the activation plugin on your original primary site, skip this topic and only perform the steps outlined in [Activating Content in the SAP HANA Repository \[page 160\]](#) for your new primary site.

Prerequisites

You performed the steps outlined in [Activating Content in the SAP HANA Repository \[page 160\]](#) on your original primary site.

Context

Do the following once takeover is complete on the new primary site.

Procedure

1. From your SAP HANA machine as the `<sid>adm` user, run these commands:

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<lowercase tenant database name>/config/cluster.cfg --deactivate-job  
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<lowercase tenant database name>/config/cluster.cfg --delete-truststore  
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<lowercase tenant database name>/config/cluster.cfg --install
```

2. From the same machine as the same user, run this command for each streaming analytics host used by the tenant on the new primary site:

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<lowercase tenant database name>/config/cluster.cfg --add-host --hdb-host  
<streaming host on new primary site>
```

9 Deploying an Adapter to a Cluster

You can start, stop, and remove an adapter from a cluster using the SAP HANA cockpit, streaming analytics plugin for SAP HANA studio, or the `streamingclusteradmin` command line utility.

Context

To manage an adapter using a cluster, add the adapter to the cluster using the `streamingclusteradmin` command line utility. You can only do this for adapters developed using the adapter toolkit (either preconfigured or custom), so ensure that you have one of these ready to deploy. For a complete list of adapters developed using the adapter toolkit, see *Preconfigured Adapters Included with the Adapter Toolkit* in the *SAP HANA Streaming Analytics: Building Custom Adapters* guide.

Procedure

1. Add the adapter to the cluster by executing:

For default (single) tenant systems,

```
streamingclusteradmin --uri=esps://<host>:3<instance-number>26 --
username=<username> --password=<password> --add_adapter --workspace-
name=<workspace-name> --adapter-name=<adapter-name> --adapter-
type=toolkit_adapter --arc=/config/adapter_config.xml
```

For multiple tenant systems, the external port can be a different value than 3xx26. The external port is always the internal port number plus one. For example, in single tenant systems, the internal port is 3xx16 (where xx is the instance number) and the external port is 3xx26. In multiple tenant systems, if the internal port is 3xx40 then the external port is 3xx41, if 3xx49 then 3xx50, and so on.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='<streamingserver>' AND HOST='<host name>'` statement and look at the PORT column in the `M_SERVICES` view.

Note

If you omit the password parameter when you call the `streamingclusteradmin` tool, streaming analytics prompts you for the password and hides it as you type, which improves security.

2. (Optional) Start the adapter by executing:

For default (single) tenant systems,

```
streamingclusteradmin --uri=esps://<host>:3<instance-number>26 --
username=<username> --password=<password> --start_adapter --workspace-
name=<workspace-name> --adapter-name=<adapter-name>
```

If you've got a multiple tenant system, see the step above for important information regarding external port numbers.

In this section:

[Application Deployment Parameters \[page 45\]](#)

Application deployment parameters specify whether warm standby or affinity are enabled on the application being deployed to the cluster. Manually create an application deployment (XML) configuration file to set these options. Each application to be deployed to the cluster, such as an adapter, requires its own configuration file.

[Sample Application Deployment Configuration File \[page 47\]](#)

Example content of an application deployment configuration file.

[Adapter Warm Standby \[page 48\]](#)

An adapter fails when it does not run properly or stops running properly. If warm standby is enabled, the cluster restarts the adapter if it is killed or stopped abnormally.

[Adapter Affinity \[page 48\]](#)

Adapters in cluster-managed mode support controller affinities, which lets you set which node the adapter can run on. Each adapter allows a single controller affinity.

[Running Custom Adapters on a Remote Server \[page 48\]](#)

Set up the components required to run a custom toolkit adapter on a remote server.

9.1 Application Deployment Parameters

Application deployment parameters specify whether warm standby or affinity are enabled on the application being deployed to the cluster. Manually create an application deployment (XML) configuration file to set these options. Each application to be deployed to the cluster, such as an adapter, requires its own configuration file.

The streaming analytics installation directory does not include a sample application deployment configuration file. See *Sample Application Deployment Configuration File* for sample file content.

These options are only valid for adapters built using the adapter toolkit. For a full list of these adapters, see *Preconfigured Adapters Included with the Adapter Toolkit* in the *SAP HANA Streaming Analytics: Building Custom Adapters* guide.

XML Element	Description
ApplicationDeployment	(Required) Root element containing all of the following elements.
Instances	(Required) Section containing the Instance element and its sub-elements.

XML Element	Description
Instance	<p>Type: <code>string</code></p> <p>(Required) Specifies the application instance. Contains two attributes:</p> <ul style="list-style-type: none"> <code>name</code> - the application name <code>type</code> - the application instance type <p>The only valid value for the <code>type</code> attribute is <code>toolkit_adapter</code>.</p> <p>This element contains the <code>Failover</code> and <code>Affinities</code> elements.</p>
Failover	<p>Type: <code>boolean</code></p> <p>(Required) Specifies whether you want to enable warm standby for the adapter. If disabled, adapter restarts are not permitted. If enabled, the <code>FailureInterval</code> and <code>FailurePerInterval</code> elements are available and adapter restarts are permitted.</p> <p>Contains the <code>FailureInterval</code> and <code>FailurePerInterval</code> elements.</p>
FailureInterval	<p>Type: <code>integer</code></p> <p>(Required) Specifies the time, in seconds, of an interval.</p> <p>If warm standby is enabled, a value is required. If warm standby is enabled and this value is empty, the adapter cannot be successfully added to the cluster.</p>
FailurePerInterval	<p>Type: <code>integer</code></p> <p>(Required) Specifies the number of restarts the application can attempt within a given interval. This count resets to zero if you restart the application manually, or if failures are dropped from the list because they are older than the size of the interval.</p>
Affinities	<p>(Optional) Section containing the <code>Affinity</code> element.</p>

XML Element	Description
Affinity	<p>Type: string</p> <p>(Optional) Specifies a controller affinity, which determines the node the application can run on. The only valid value is controller.</p> <p>This element contains four attributes: charge, strength, type, and value. If an affinity is specified, values for these attributes must be specified.</p> <p>Charge can be positive or negative. If positive, the application runs on the node. If negative, the application does not run on the node.</p> <p>Strength can be strong or weak. Strong requires the application to run only on a specific node. If you have strong, positive affinity set for a node, and that node fails, the failover process tries to restart the application on that node. If the node has not recovered, the application restart fails, and you must restart manually.</p> <p>Type specifies the affinity type. Controller is the only valid value.</p> <p>Value specifies the value of the affinity. For example, if the affinity type is controller, its value should be the node name for which the affinity is set.</p>

9.2 Sample Application Deployment Configuration File

Example content of an application deployment configuration file.

```
<?xml version="1.0" encoding="utf-8"?>
<ApplicationDeployment xmlns="http://www.sybase.com/esp/application/2014/01/">

  <Instances>
    <Instance name="adapter1" type="toolkit_adapter">
      <Failover enable="true">
        <FailureInterval>1</FailureInterval>
        <FailuresPerInterval>1</FailuresPerInterval>
      </Failover>
      <Affinities>
        <Affinity charge="negative" strength="strong"
type="controller" value="node2"/>
      </Affinities>
    </Instance>
  </Instances>
</ApplicationDeployment>
```

9.3 Adapter Warm Standby

An adapter fails when it does not run properly or stops running properly. If warm standby is enabled, the cluster restarts the adapter if it is killed or stopped abnormally.

Failover typically results in a restart, though a strong positive affinity to a node that is not available can prevent an adapter from restarting. Restarts can be limited based on failure intervals and restarts per interval.

Enable warm standby in the application deployment configuration file. See [Application Deployment Parameters \[page 45\]](#) for detailed information on configuring failover for adapters managed by a cluster.

9.4 Adapter Affinity

Adapters in cluster-managed mode support controller affinities, which lets you set which node the adapter can run on. Each adapter allows a single controller affinity.

Set adapter affinity in the application deployment configuration file. See [Application Deployment Parameters \[page 45\]](#) for detailed information on configuring adapters managed by a cluster.

9.5 Running Custom Adapters on a Remote Server

Set up the components required to run a custom toolkit adapter on a remote server.

Context

To run custom adapters developed using the adapter toolkit on a remote server, first set up adapter toolkit directories on the target machine.

Procedure

1. On the remote server, create a new folder to house your adapter toolkit directory.
2. Set the STREAMING_HOME environment variable to point to the new folder.
 - a. Go to Control Panel, then select **System > Advanced System Settings > Environment Variables**.
 - b. Set the STREAMING_HOME variable to the newly created adapter directory.
3. From an existing installation of streaming analytics, copy the following directories into the new \$STREAMING_HOME location:

- `$STREAMING_HOME/adapters/framework`
- `$STREAMING_HOME/lib/jre`

Ensure that the original directory structure under `STREAMING_HOME` is preserved on the remote server.

i Note

If you prefer to use your own copy of the JRE, use version 6 or higher. Modify the scripts under `$STREAMING_HOME/adapters/framework/bin` to point to the correct JRE location.

4. (Optional) If you are running only custom adapters on the remote server, and none of the preconfigured adapters included with streaming analytics, remove:
 - `$STREAMING_HOME/adapters/framework/examples`
 - `$STREAMING_HOME/adapters/framework/instances`

10 Deploying a Project to a Cluster

To run a project, deploy it to a streaming cluster. You can deploy projects through SAP HANA studio, or using the `streamingclusteradmin` utility.

In this section:

[Deploying a Project to a Cluster from Studio \[page 50\]](#)

The streaming analytics plugin for SAP HANA studio provides a graphic interface for connecting to a server.

[Deploying a Project to a Cluster Using `streamingclusteradmin` \[page 51\]](#)

You can set project options using the project configuration view in SAP HANA studio, or by editing the project's `.ccr` file. Then, use the `streamingclusteradmin` command line utility to deploy projects to a cluster.

[Editing Cluster Parameters in Project Configuration \[page 52\]](#)

Configure the cluster that your project binds to for receiving and providing data. The cluster can then be used when configuring bindings.

[Accessing Project Contents using SAP HANA Smart Data Access \[page 53\]](#)

Use the Smart Data Access feature in SAP HANA to connect to a project for streaming analytics and query the contents of its windows, using the content as a data source for SAP HANA.

10.1 Deploying a Project to a Cluster from Studio

The streaming analytics plugin for SAP HANA studio provides a graphic interface for connecting to a server.

Prerequisites

If you do not have a server connection already defined, create a new connection by selecting **New Server URL** in the server view in the SAP HANA Streaming Run-Test perspective.

Context

View cluster setup details for streaming analytics plugin for SAP HANA studio in: `$(STREAMING_HOME)/studio/clustercfg/studio.xml`. Do not modify this file.

Procedure

1. If you are not already in the SAP HANA Streaming Run-Test perspective, click the SAP HANA Streaming Run-Test tab at the top of the window, or select **Window > Open Perspective > Other > [X]**.
2. If you do not see the server view window, select **Window > Show View > Other > Server View [X]**.
3. Right-click on the server you want to connect to and select **Connect Server**.
4. To connect all of the listed servers, select the **Reconnect All** icon from the top-right corner of the server view window.

Unselecting **Filter Metadata Streams** causes all metadata streams to appear in the server view.

5. Right-click on the target workspace and select **Load Project(s) into Workspace**.

10.2 Deploying a Project to a Cluster Using streamingclusteradmin

You can set project options using the project configuration view in SAP HANA studio, or by editing the project's `.ccr` file. Then, use the `streamingclusteradmin` command line utility to deploy projects to a cluster.

Procedure

1. Start the cluster, or ensure that the cluster is running.
2. Add the project to the node by executing:

For default (single) tenant systems,

```
streamingclusteradmin --uri=esps://<host>:3<instance-number>26 --  
username=<username> --password=<password>  
> add project <workspace-name>/<project-name> <project-name>.ccx <project-  
name>.ccr
```

For multiple tenant systems, the external port can be a different value than 3xx26. The external port is always the internal port number plus one. For example, in single tenant systems, the internal port is 3xx16 (where xx is the instance number) and the external port is 3xx26. In multiple tenant systems, if the internal port is 3xx40 then the external port is 3xx41, if 3xx49 then 3xx50, and so on.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='<streamingserver>' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view.

i Note

If you omit the password parameter when you call the `streamingclusteradmin` tool, streaming analytics prompts you for the password and hides it as you type, which improves security.

- (Optional) Start the deployed project by executing:

For default (single) tenant systems,

```
streamingclusteradmin --uri=esps://<host>:3<instance-number>26 --
username=<username> --password=<password>--start_project
--workspace-name=<workspace-name> --project-name=<project-name>
```

If you've got a multiple tenant system, see the step above for important information regarding external port numbers.

10.3 Editing Cluster Parameters in Project Configuration

Configure the cluster that your project binds to for receiving and providing data. The cluster can then be used when configuring bindings.

Procedure

- In the CCR Project Configuration editor window, select the **Clusters** tab.
- Select a cluster to edit using one of the following methods:
 - Click the name of an existing cluster in the **All Clusters** pane to edit that cluster's information.
 - Click **Discover** to search for available clusters.
 - Click **Add** to add a new cluster.

The editor displays the **Cluster Details** pane.

- Enter the requested information in the **Cluster Details** pane.

Field	Description
Cluster URL	Property name as it appears in the ccr file: <code>name</code> Enter the hostname and port number of the cluster. For example, <code>esps://<host>:<port></code> .
Type	Property name as it appears in the ccr file: <code>type</code> The cluster connection type. Toggle between local (no server information necessary) and remote (server information must be known). Local refers to the cluster where the project is running. Remote refers to a different cluster from the one the project is running under.
Username	Property name as it appears in the ccr file: <code>Username</code> A username to log in to the cluster.

Field	Description
Password	Property name as it appears in the ccr file: <code>password</code> A password to log in to the cluster.

4. (Optional) Click **Encrypt** after entering the username or password.
 - a. Enter the required information in the **Content Encryption** pane, including **Cluster URL**, comprised of your hostname and port number (`esps://<host>:<port>`), and credential fields.
 - b. Click **Encrypt**.
The editor redisplay the **Cluster Details** pane with the field you encrypted (either the username or password) displaying randomized encryption characters.

i Note

To reset the encryption, click **Encrypt** beside the appropriate field, and click **Reset** when the **Already Encrypted** popup appears.

10.4 Accessing Project Contents using SAP HANA Smart Data Access

Use the Smart Data Access feature in SAP HANA to connect to a project for streaming analytics and query the contents of its windows, using the content as a data source for SAP HANA.

Procedure

1. On the SAP HANA machine, create a remote source over the streaming analytics window.
2. On the SAP HANA machine, create a virtual table over the remote source.
See the *SAP HANA Administration Guide* for complete information on SAP HANA Smart Data Access.

Related Information

[SAP HANA Administration Guide](#)

11 Monitoring Streaming Analytics

You can use various tools to monitor the status of streaming analytics and optimize performance.

The tools available are:

- SAP HANA cockpit.
- SAP HANA studio:
 - The system monitor provides an overview of all your SAP HANA systems at a glance, including system availability, and current resource usage information. From the system monitor, you can drill down into each individual system in the administration editor.
 - The administration editor provides detailed information about resource usage, current alerts, system performance, and system configuration, as well as tools for analyzing and troubleshooting issues in your system.
 - Access to SAP HANA cockpit.

For more information about SAP HANA studio, see the *SAP HANA Administration Guide*.

- `streamingmonitor` and `streamingsubscribe` utilities. See [Monitoring a Project with Utilities \[page 55\]](#).

In this section:

[Monitoring a Project with Utilities \[page 55\]](#)

Use the command-line interface to monitor project performance.

[Streaming Analytics Statistics in SAP HANA Studio \[page 57\]](#)

The statistics service (or statistics server) collects and evaluates information about status, performance, and resource consumption from all components belonging to the system, and issues alerts when necessary.

[Editing the XS Job for SAP HANA Streaming Analytics \[page 60\]](#)

An Extended Application Services (XS) job called `sap.hana.streaming.repoplugin::plugin.xsjob` is created as part of the SAP HANA streaming analytics installation.

[Monitoring with Metadata Streams \[page 61\]](#)

Metadata streams are automatically created by SAP HANA streaming analytics. You can query and subscribe to these streams to obtain important health and performance information about your running projects and adapters.

Related Information

[SAP HANA Administration Guide](#)

11.1 Monitoring a Project with Utilities

Use the command-line interface to monitor project performance.

The `streamingmonitor` tool reads performance data from a running instance of the project server and displays it on standard output. Monitoring data is only available if the time-granularity option is enabled. Set the option in one of three ways:

- Directly in the project configuration (CCR) file.
- Using the server view of the SAP HANA Streaming Run-Test perspective in studio.
- Using the `time_granularity` command of the `streamingprojectclient` utility. See `streamingprojectclient` in the *SAP HANA Streaming Analytics: Utilities Guide* for more information.

The time-granularity option specifies, in seconds, how often the set of performance records — one per stream and one per gateway connection — is obtained from the running streaming analytics. By default, time-granularity is set to 5. To increase performance, disable the time-granularity project option when monitoring is not required. For more information on configuring projects with the streaming analytics plugin for SAP HANA studio, see *Project Configurations* in the *SAP HANA Streaming Analytics: Developer Guide*.

i Note

The `streamingprojectclient` stream contains basic information about the connected clients, but performance-related fields are populated only with the monitoring option.

For example, to monitor a project running in a cluster that has a node on the host "myhost.sybase.com" with RPC port 31415, use the following command:

```
streamingmonitor -p myhost.sybase.com:31415/<workspace-name>/<project-name>
```

You can also use `streamingsubscribe` to monitor the status of project streams through heartbeat messages. To monitor streaming analytics heartbeat messages, use the following command:

```
streamingsubscribe -p <host>:11180/default/prj1 -H
```

For more information on the `streamingmonitor` and `streamingsubscribe` tools, see the *SAP HANA Streaming Analytics: Utilities Guide*.

In this section:

[Monitoring Project Status \[page 56\]](#)

Use the `streamingclusteradmin --get_projects` command to monitor the status of your projects.

11.1.1 Monitoring Project Status

Use the `streamingclusteradmin --get_projects` command to monitor the status of your projects.

This command lists all available projects, independent of a workspace. The project status is provided for each project reported. To show the statuses of all currently existing projects in command line mode, enter:

```
streamingclusteradmin --uri=esps://<host>:<port> --username=<username> --password=<password> --get_projects
```

Status Types and Definitions

For each project, `--get_projects` provides two types of statuses: requested status and current status.

The requested status is the project's desired status. This status can be one of the following:

Requested Status	Description
Started-Running	The project is fully functional.
Stopped	The project is stopped entirely.

When the project is requested to start or stop, the requested status becomes started-running or stopped, respectively. The project then goes through a cycle of starting up or shutting down. The cycle is not instantaneous.

The current status is the actual status at the time of the report. This status can be one of the following:

Current Status	Description
Starting	The project has received a request to start, but the process has not started yet.
Started	The controller has received the request and is about to start the project. The requested status is started-running.
Started-Initializing	The project process has started and is in contact with the cluster, but it is not yet functional. The requested status is started-running.
Started-Running	The project is fully functional. The requested status is started-running.
Stopping	The project has received a request to stop, but the shutdown cycle is not complete yet. The request status will be stopped.
Stopped	The project is stopped. The request status is stopped.
Failed	The project achieved started-running status, but it has suddenly stopped without a stop request. The requested status is started-running.
Start-Failed	An attempt was made to start the project, but the project never achieved started-running status. The requested status is started-running.

Use the `--get_projects` command in either command line mode or interactive mode. For further instruction, consult the *Server Executables* section of the *SAP HANA Streaming Analytics: Utilities Guide*.

Investigating a Failed Status

If your project fails for any reason, review the SAP HANA trace directory, examining the project log files for reported errors. For more information on log files, consult the *File and Directory Infrastructure* and *Project Log Files* sections of the *SAP HANA Streaming Analytics: Configuration and Administration Guide*. For information on troubleshooting project errors, consult the *Project Problems* section of the *SAP HANA Streaming Analytics: Configuration and Administration Guide*.

11.2 Streaming Analytics Statistics in SAP HANA Studio

The statistics service (or statistics server) collects and evaluates information about status, performance, and resource consumption from all components belonging to the system, and issues alerts when necessary.

Certain statistics are specific to streaming analytics.

Note

To view streaming analytics statistics, you must have migrated to the statistics service. For more information, see SAP Note 1917938.

_SYS_STATISTICS.STREAMING_PROJECTS_STATISTICS

`_SYS_STATISTICS.STREAMING_PROJECTS_STATISTICS` provides detailed data for streaming analytics projects.

- Interval: 3 min
- Base Table: `M_STREAMING_PROJECTS`

Name	SQL Data Type
SNAPSHOT_ID	TIMESTAMP
SERVER_TIMESTAMP	TIMESTAMP
INDEX	NVARCHAR
SCHEMA_NAME	NVARCHAR
PROJECT_NAME	NVARCHAR
CPU_SYSTEM_UTILIZATION	REAL
CPU_USER_UTILIZATION	REAL

Name	SQL Data Type
CPU_UTILIZATION	REAL
RESIDENT_MEMORY_SIZE	BIGINT
THREAD_COUNT	INTEGER
VIRTUAL_MEMORY_SIZE	BIGINT

_SYS_STATISTICS.STREAMING_PROJECTS_CONNECTIONS_STATISTICS

_SYS_STATISTICS.STREAMING_PROJECTS_CONNECTIONS_STATISTICS provides detailed data for the connections of streaming analytics projects.

- Interval: 3 min
- Base Tables: M_STREAMING_PROJECTS, M_STREAMING_PROJECT_PUBLISHERS, M_STREAMING_PROJECT_SUBSCRIBERS

Name	SQL Data Type
SNAPSHOT_ID	TIMESTAMP
SERVER_TIMESTAMP	TIMESTAMP
INDEX	NVARCHAR
SCHEMA_NAME	NVARCHAR
PROJECT_NAME	NVARCHAR
TOTAL_PUB_COUNT	BIGINT
TOTAL_PUB_ROW_THROUGHPUT	REAL
TOTAL_PUB_TRANSACTION_THROUGHPUT	REAL
TOTAL_SUB_COUNT	BIGINT
TOTAL_SUB_ROW_THROUGHPUT	REAL
TOTAL_SUB_TRANSACTION_THROUGHPUT	REAL
TOTAL_CONN_COUNT	BIGINT

_SYS_STATISTICS.STREAMING_PROJECTS_STREAMS_STATISTICS

_SYS_STATISTICS.STREAMING_PROJECTS_STREAMS_STATISTICS provides detailed data for the streams of streaming analytics projects.

- Interval: 3 min
- Base Table: M_STREAMING_PROJECT_STREAMS

Name	SQL Data Type
SNAPSHOT_ID	TIMESTAMP
SERVER_TIMESTAMP	TIMESTAMP
INDEX	NVARCHAR
SCHEMA_NAME	NVARCHAR
PROJECT_NAME	NVARCHAR
STREAM_NAME	NVARCHAR
CPU_SYSTEM_UTILIZATION	REAL
CPU_USER_UTILIZATION	REAL
CPU_UTILIZATION	REAL
ROW_THROUGHPUT	REAL
TRANSACTION_THROUGHPUT	REAL

_SYS_STATISTICS.STREAMING_PUBLISHERS_STATISTICS

_SYS_STATISTICS.STREAMING_PUBLISHERS_STATISTICS provides detailed data for the publishers of streaming analytics projects.

- Interval: 3 min
- Base Table: M_STREAMING_PROJECT_PUBLISHERS

Name	SQL Data Type
SNAPSHOT_ID	TIMESTAMP
SERVER_TIMESTAMP	TIMESTAMP
INDEX	NVARCHAR
SCHEMA_NAME	NVARCHAR
PROJECT_NAME	NVARCHAR
CLIENT_IP	VARCHAR
CLIENT_PORT	INTEGER
CPU_SYSTEM_UTILIZATION	REAL
CPU_USER_UTILIZATION	REAL
CPU_UTILIZATION	REAL
ROW_THROUGHPUT	REAL
TRANSACTION_THROUGHPUT	REAL

_SYS_STATISTICS.STREAMING_SUBSCRIBERS_STATISTICS

_SYS_STATISTICS.STREAMING_SUBSCRIBERS_STATISTICS provides detailed data for the subscribers of streaming analytics projects.

- Interval: 3 min
- Base Table: M_STREAMING_PROJECT_SUBSCRIBERS

Name	SQL Data Type
SNAPSHOT_ID	TIMESTAMP
SERVER_TIMESTAMP	TIMESTAMP
INDEX	NVARCHAR
SCHEMA_NAME	NVARCHAR
PROJECT_NAME	NVARCHAR
CLIENT_IP	VARCHAR
CLIENT_PORT	INTEGER
CPU_SYSTEM_UTILIZATION	REAL
CPU_USER_UTILIZATION	REAL
CPU_UTILIZATION	REAL
ROW_THROUGHPUT	REAL
TRANSACTION_THROUGHPUT	REAL

11.3 Editing the XS Job for SAP HANA Streaming Analytics

An Extended Application Services (XS) job called `sap.hana.streaming.repopugin::plugin.xsjob` is created as part of the SAP HANA streaming analytics installation.

This job periodically checks the `sap.hana.streaming.repopugin::ACTIVATIONS` database table, located in the `SAP_HANA_STREAMING` schema, for new entries signaling new streaming analytics project activations. By default, the job performs this check every 30 seconds and deletes entries older than 30 days from the database table. To change these intervals, edit the XS job using the XS Job Dashboard in the SAP HANA XS Administration Tool.

To avoid potential negative impacts on system performance, change the amount of time entries are stored in the database table. Edit the `deleteActivationsAfterThisManyDays` parameter of the job schedule. To change how often the job runs, edit the value of the `XCron` keyword.

See the *Scheduling XS Jobs* section in the *SAP HANA Administration Guide* for detailed information on XS jobs.

Related Information

[SAP HANA Administration Guide](#)

11.4 Monitoring with Metadata Streams

Metadata streams are automatically created by SAP HANA streaming analytics. You can query and subscribe to these streams to obtain important health and performance information about your running projects and adapters.

Much of the information that metadata streams provide also appears in the SAP HANA cockpit. Refer to the section *Managing Streaming Analytics with SAP HANA Cockpit* in this guide.

Some metadata streams contain static information that never changes while the project is running, for example, `_ESP_Streams`. Other streams continuously update at various periods or on various events. You can subscribe, query, and view metadata streams in the same way as regular streams. For example, you can use the `streamingsubscribe` utility to subscribe to the streams `_ESP_Connectors` and `_ESP_Streams` from the project `default/prj1`, and print all stream data in XML format to standard output:

```
streamingsubscribe -c user-id:password -s _ESP_Connectors,_ESP_Streams -p <host-name>/default/prj1
```

For details on `streamingsubscribe`, see the *SAP HANA Streaming Analytics: Utilities Guide*.

Note

The schema for metadata streams can change between releases as the set of statistics the streams report expands. New columns are added to the end of the schema for existing metadata streams. Keep this in mind when coding.

Metadata streams differ from general streams in the following ways:

- Metadata streams have reserved names. No other objects can use these names.
- Metadata streams store their records in a special store called `ESPMetadataStore`. No other streams can use this store.
- Metadata streams cannot be used in CCL or serve as an input for a stream in a project. For example, the following usage is not possible:

```
INSERT INTO myStream SELECT * FROM _ESP_Connectors WHERE latency > 1
```

In this section:

[_ESP_Adapter_Statistics \[page 63\]](#)

Reports statistics unique to each adapter. Both internal and external adapters can publish statistics to this stream.

[_ESP_Clients \[page 64\]](#)

Contains information about all the currently active gateway client connections.

[_ESP_Clients_Monitor \[page 64\]](#)

Contains information about the performance of all currently active gateway client connections and a copy of data from the `_ESP_Clients` stream. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

[_ESP_Clockupdates \[page 66\]](#)

Delivers notifications of changes in the logical clock of the project.

[_ESP_Columns \[page 66\]](#)

Contains information about all columns of all streams.

[_ESP_Config \[page 67\]](#)

Contains the current project configuration (`.ccx`) information.

[_ESP_Connectors \[page 67\]](#)

Contains information about all internal adapters defined in the project.

[_ESP_GD_Sessions \[page 68\]](#)

Contains information about registered guaranteed delivery sessions that may be active or inactive. In a guaranteed delivery session, a guaranteed delivery window or stream transmits event information to a registered guaranteed delivery subscriber.

[_ESP_Keycolumns \[page 69\]](#)

Contains information about the primary key columns of all the streams. If a stream has a primary key, the columns that make up the key are listed in this stream.

[_ESP_Project_Monitor \[page 69\]](#)

Contains information on project CPU usage, memory consumption, and number of threads. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

[_ESP_Resource_Info \[page 70\]](#)

Stores thread object information, collects memory information for every thread, and displays alert messages.

[_ESP_RunUpdates \[page 72\]](#)

Delivers notifications of changes during debugging. The streaming analytics server sends notifications only when the project is in trace mode.

[_ESP_Streams \[page 72\]](#)

Contains information about all streams and windows.

[_ESP_Streams_Monitor \[page 73\]](#)

Contains information about the performance of streams, and a copy of data from the `_ESP_Streams` stream. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

[_ESP_Streams_Topology \[page 74\]](#)

Contains pairs of names for streams that are directly connected.

[_ESP_Subscriptions \[page 74\]](#)

Contains information about all currently active subscriptions. A dropped connection is considered unsubscribed from everything to which it was subscribed.

[_ESP_Subscriptions_Ext \[page 75\]](#)

Contains information about all currently active subscriptions.

11.4.1 _ESP_Adapter_Statistics

Reports statistics unique to each adapter. Both internal and external adapters can publish statistics to this stream.

The information this stream provides is especially useful for custom and unmanaged adapters. For project-managed and cluster-managed adapters, use the statistics in the SAP HANA cockpit. See *Managing Streaming Analytics with SAP HANA Cockpit > Managing Streaming Adapters* in this guide.

Column	Type	Description
adapter_name	string	A unique name of the adapter instance.
stat_name	string	The name of an adapter statistic, as defined by the adapter.
last_update	bigdatetime	The time that the statistic was last updated.
value	string	The value of the statistic (converted to a string).

Several adapters have been created using the adapter toolkit and consist of various modules (transporters, formatters, ESPPublisher, or ESPSubscriber). See *Preconfigured Adapters Included with the Adapter Toolkit* in the *SAP HANA Streaming Analytics: Building Custom Adapters* for a full list of these adapters. Note that:

- Output transporter modules that run in streaming mode do not report statistics to the `_ESP_Adapter_Statistics` metadata stream.
- Output transporter modules that run in row mode report statistics to the `_ESP_Adapter_Statistics` metadata stream.
- Input transporter modules that run in either streaming or row mode report statistics to the `_ESP_Adapter_Statistics` metadata stream.

See *Transporters Currently Available from SAP* and *Formatters Currently Available from SAP* in the *SAP HANA Streaming Analytics: Building Custom Adapters* guide to determine in which mode a module operates. See *Adapter Controller Parameters* for the adapter of your choice in the *SAP HANA Streaming Analytics: Adapters Guide* for details on tuning the frequency that the adapter reports its statistics to the `_ESP_Adapter_Statistics` metadata stream.

For information on the statistics that each adapter publishes, see the *SAP HANA Streaming Analytics: Adapters Guide*.

11.4.2 _ESP_Clients

Contains information about all the currently active gateway client connections.

Column	Type	Description
Handle	long	A unique integer ID of the connection.
user_name	string	The username to login to the connection, shown once the user is authenticated.
IP	string	The address of the client machine.
host	string	The symbolic host name of the client machine, if available. If not available, host is the IP address of the client machine.
port	integer	The TCP port number from which the connection originates.
login_time	msdate	The time the server accepts (but does not authenticate) the connection, in GMT.
conn_tag	string	The user-set symbolic connection tag name. If not set by the user, conn_tag is NULL.

11.4.3 _ESP_Clients_Monitor

Contains information about the performance of all currently active gateway client connections and a copy of data from the _ESP_Clients stream. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

Column	Type	Description
handle	long	A unique integer ID of the connection.
IP	string	The IP address of the client machine, as a string.
host	string	The symbolic host name of the client machine, if available. If not available, host is the IP address of the client machine.
port	integer	The TCP port number from which the connection originates.
login_time	bigdatetime	The time the server accepts (but does not authenticate) the connection.
cpu_pct	float	Total CPU usage for the client thread, as a percentage of a single CPU core.
last_update	bigdatetime	The time of the current update.
subscribed	integer	The status of a subscription to a stream: 1 if subscribed; 0 if not subscribed, indicating a publisher.
sub_trans_per_sec	float	The client's performance, in transactions per second, received by the client since the last update.

Column	Type	Description
sub_rows_per_sec	float	The client's performance, in data rows per second, received by the client since the last update.
sub_inc_trans	long	The number of transactions, envelopes, or messages received by the client since the last update.
sub_inc_rows	long	The number of data rows received by the client since the last update.
sub_total_trans	long	The total number of transactions, envelopes, or messages received by the client. Transactions, envelopes, and messages still in the queue are not counted.
sub_total_rows	long	The total number of data rows received by the client. Rows still in the queue are not counted.
sub_dropped_rows	long	The total number of data rows dropped in the gateway because they were not read quickly enough by the client. For lossy subscriptions.
sub_accum_size	integer	The current number of rows collected in the accumulator to be sent in the next pulse. For pulsed subscriptions.
sub_accum_ops	integer	The current number of rows processed by the accumulator since the last pulse. For pulsed subscriptions.
sub_queue	integer	The number of rows queued for transmission to the client.
sub_queue_fill_pct	float	The current sub_queue, as a percentage, relative to the queue size limit. If sub_queue_fill_pct reaches 100 percent, any future attempts to post data to this client are blocked, propagating the flow control back to the source of the post.
sub_work_queue	integer	The number of rows for transmission to the client that are being transferred from the proper queue to the socket buffer. The rows can be regrouped by envelopes.
pub_trans_per_sec	float	The client's performance, in transactions per second, sent by the client since the last update. Envelopes and any service messages count as transactions.
pub_rows_per_sec	float	The client's performance, in data rows per second, sent by the client since the last update.
pub_inc_trans	long	The number of transactions, envelopes, or messages sent by the client since the last update.
pub_inc_rows	long	The number of data rows sent by the client since the last update.
pub_total_trans	long	The total number of transactions, envelopes, or messages sent by the client. Does not include the transactions, envelopes, or messages in the queue.
pub_total_rows	long	The total number of data rows sent by the client. Does not include the rows in the queue.
pub_stream_id	long	The numeric ID of the stream to which the client is trying to currently publish data. Typically, pub_stream_id is -1, meaning that data is not currently being published to the stream. Since the server publishes data in short bursts, the monitor will not report any value other than -1.

Column	Type	Description
node_cpu_pct	float	Total CPU usage for the client, as a percentage of all CPU cores on the machine. Total CPU usage equals system CPU usage plus user CPU usage.
node_cpu_pct_system	float	System CPU usage for the client, as a percentage of all CPU cores on the machine.
node_cpu_pct_user	float	User CPU usage for the client, as a percentage of all CPU cores on the machine.
cpu_time	interval	Total CPU time since the creation of the client, in microseconds. Total CPU time equals system CPU time plus user CPU time.
cpu_time_system	interval	Total system CPU time, in microseconds, since the creation of the client thread.
cpu_time_user	interval	Total user CPU time, in microseconds, since the creation of the client thread.
time_since_start	interval	Duration of lapsed real time since the creation of the client thread.

11.4.4 _ESP_Clockupdates

Delivers notifications of changes in the logical clock of the project.

Column	Type	Description
Key	string	The type of the update, which is always <code>CLOCK</code> .
Rate	float	The rate of the logical clock relative to the real time.
Time	float	The current time in seconds since the UNIX epoch.
Real	integer	The real time flag, 1 if the logical clock matches the system time, and 0 if the times do not match.
stop_depth	integer	The number of times the clock resume command must be called to resume the flow of time. When the clock is running, <code>stop_depth</code> is 0.
max_sleep	integer	The time, in real milliseconds, which guarantees all sleepers discover changes in the physical clock rate or time.

11.4.5 _ESP_Columns

Contains information about all columns of all streams.

Column	Type	Description
username	string	Hard-coded as <code>user</code> .
relname	string	The name of the stream that contains columns described by this row.

Column	Type	Description
attname	string	The name of the column described by this row.
attypid	integer	The internal PostgreSQL value representing the type of this column. Valid values: <ul style="list-style-type: none"> integer - 23 long - 20 money - 701 float - 701 date - 1114 timestamp - 1114 string - 1043
attnum	integer	The position of this column in the schema, starting from 0.

11.4.6 _ESP_Config

Contains the current project configuration (.ccx) information.

Column	Type	Description
key	string	The format of the information in the .ccx file. This value is always xml.
value	string	The text of the current .ccx file.

11.4.7 _ESP_Connectors

Contains information about all internal adapters defined in the project.

Column	Type	Description
name	string	The name of the adapter, as defined in the project.
stream	string	The name of the stream on which the adapter is defined.
type	string	The adapter type defined in the ATTACH ADAPTER statement.
input	integer	The function of the adapter, specifying whether it is input or output. The values are 1 for InConnection, and 0 for OutConnection.
ingroup	string	The StartUp group where this connector belongs.

Column	Type	Description
state	string	The state of the adapter: <ul style="list-style-type: none"> • READY – ready to be started. • INITIAL – performing start-up and initialization. • CONTINUOUS – continuously receiving real-time data. • IDLE – currently not receiving data but attempting to reconnect to the data source or link. • DONE – no remaining input or output data; the adapter is about to exit. • DEAD – the adapter thread exited. The adapter remains in this state until explicitly requested to restart.
total_rows	long	The total number of data records recognized in the input data.
good_rows	long	The number of data records successfully processed.
bad_rows	long	The number of data records that experienced errors. The fields total_rows, good_rows, and bad_rows are updated once in a few seconds to reduce the overhead.
last_error_time	seconddate	The time that the error occurred in YYYY-MM-DD hh:mm:ss format.
last_error_msg	string	The complete text of the error message as written to the log.
latency	interval	The latency introduced by the adapter. For an input adapter, this is the amount of time it takes the adapter to receive data from its source and publish the data to the stream. For an output adapter, this is the amount of time it takes for the adapter to receive a message from the stream and publish the data to its destination. The SAP HANA Output adapter and the Database Input and Output adapters report latency in milliseconds. Other internal adapters report latency in seconds. For adapters that do not report latency information, the column value is NULL.

11.4.8 _ESP_GD_Sessions

Contains information about registered guaranteed delivery sessions that may be active or inactive. In a guaranteed delivery session, a guaranteed delivery window or stream transmits event information to a registered guaranteed delivery subscriber.

The _ESP_GD_Sessions metadata stream tracks all the guaranteed delivery (GD) sessions for a project. You can monitor the streams or windows being subscribed to in a given guaranteed delivery session, the client handle associated with a session and the last sequence number committed for a given GD session, stream or window combination.

When a project contains at least one stream or window that supports guaranteed delivery, streaming analytics stores _ESP_GD_Sessions in a metadata log store so the stream can be recovered after a restart. If a project contains no GD streams, _ESP_GD_Sessions is stored in a memory store, but it is not used.

In some situations, there is a delay in updating this stream as new subscriptions are added or existing windows or streams are dropped.

Column	Type	Description
gd_key	string	The automatically generated key for the guaranteed delivery session. The gd_key is unique for a given gd_name/user_name combination.
stream_name	string	The name of the stream or window this guaranteed delivery session subscribes to.
user_name	string	The user associated with this guaranteed delivery session.
gd_name	string	The name of this guaranteed delivery session.
sequence_no	long	The sequence number of the last event committed from stream_name in this guaranteed delivery session. A value of 0 indicates that no commits have been issued.
client_handle	long	The active client handle associated with this guaranteed delivery session. A value of -1 indicates that there are no active clients for this guaranteed delivery session. For active connectors, the value is 0.
last_update	bigdatetime	The last time this stream entry was updated.

The key for this stream consists of the `gd_key` plus the `stream_name`—that is, there is one stream entry per `gd_key/stream_name` pair.

11.4.9 _ESP_Keycolumns

Contains information about the primary key columns of all the streams. If a stream has a primary key, the columns that make up the key are listed in this stream.

Column	Type	Description
table	string	The name of the stream owning the column described by this row.
field	string	The name of the column described by this row.
type	integer	The internal PostgreSQL value representing the type of this column. The possible values are: <ul style="list-style-type: none"> integer - 23 long - 20 money - 701 float - 701 date - 1114 timestamp - 1114 string - 1043

11.4.10 _ESP_Project_Monitor

Contains information on project CPU usage, memory consumption, and number of threads. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The

frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

Column	Type	Description
<code>project_name</code>	string	Currently hard-coded to the word <code>project</code> .
<code>node_cpu_pct</code>	float	Total CPU usage, as a percentage, by the project since the last update. Total CPU usage equals the system CPU usage plus the user CPU usage. Valid values are in the range from 0.0 to 100.00%. On multicore machines, the percentage is relative to the total number of available cores. A value of 100% indicates a usage of 100% of all cores on the machine.
<code>node_cpu_pct_system</code>	float	The system (Kernel on Windows) CPU usage, as a percentage, by the project since the last update. Valid values are in the range from 0.0 to 100.00%. On multicore machines, the percentage is relative to the total number of available cores. A value of 100% indicates a usage of 100% of all cores on the machine.
<code>node_cpu_pct_user</code>	float	The user CPU usage, as a percentage, by the project since the last update. Valid values are in the range from 0.0 to 100.00%. On multicore machines, the percentage is relative to the total number of available cores. A value of 100% indicates a usage of 100% of all cores on the machine.
<code>cpu_time</code>	interval	Total CPU time for the project, in microseconds. Total CPU time is equal to the system CPU time plus the user CPU time.
<code>cpu_time_system</code>	interval	Total system CPU time for the project, in microseconds.
<code>cpu_time_user</code>	interval	Total user CPU time for the project, in microseconds.
<code>time_since_start</code>	interval	Duration of lapsed real time since the project was started, in microseconds.
<code>startmem_usage_vm</code>	long	Total amount of virtual memory, in bytes, used by the project at the time of the update.
<code>mem_usage_rss</code>	long	Total amount of system memory (RSS), in bytes, used by the project at the time of the update.
<code>num_threads</code>	integer	Total number of threads used by the project at the time of the update.
<code>last_update</code>	bigint	Time of the current update.

11.4.11 `_ESP_Resource_Info`

Stores thread object information, collects memory information for every thread, and displays alert messages.

`_ESP_Resource_Info` stores thread information, which includes the thread name, type, ID, CPU usage, memory usage, input queue size, and so on. When memory tracking is enabled, either as an option in the project configuration file or in `streamingproject`, or as a command in `streamingprojectclient`, the resource information metadata stream collects memory information and displays memory alert messages. When disabled, any previous memory query information is lost.

`_ESP_Resource_Info` contains the following columns:

Column	Type	Description
<code>name</code>	<code>string</code>	The name of the thread object, which includes engine, stream, adapter, gateway client, and so on.
<code>type</code>	<code>integer</code>	The type of the object, which includes: <ul style="list-style-type: none"> • 0 - SYSTEM_THREAD • 1 - CLUSTER_THREAD • 2 - PLATFORM_THREAD • 3 - COMMANDSERVER_THREAD • 4 - SQLSERVER_THREAD • 5 - GATEWAY_THREAD • 6 - GATEWAY_SAFEIO_THREAD • 7 - GATEWAY_PUBLISHER_THREAD • 8 - GATEWAY_SUBSCRIBER • 9 - STREAM_THREAD • 10 - STREAM_BASE_THREAD • 11 - STREAM_EXPORT_THREAD • 12 - CONNECTION_THREAD • 13 - CONNECTION_READER_THREAD • 14 - CONNECTION_WRITER_THREAD
<code>thread_id</code>	<code>integer</code>	The object thread ID.
<code>cpu_usage</code>	<code>double</code>	The thread CPU usage. This is two times the thread CPU usage percentage value, which is also shown in this metadata stream.
<code>memory_usage</code>	<code>integer</code>	The thread allocated memory. If the memory collection features are not supported or disabled, the memory usage is 0.
<code>input_queue_size</code>	<code>integer</code>	The current input queue size in the thread object. If there is no input queue, the size is 0.
<code>alert_messages</code>	<code>string</code>	The alert messages. When thread object issues based on thread type, CPU usage, memory usage, input queue size, and so on, are detected, the memory tracking feature sends memory alert messages to <code>_ESP_Resource_Info</code> .
<code>current_time</code>	<code>integer</code>	The current timestamp.

11.4.12 _ESP_RunUpdates

Delivers notifications of changes during debugging. The streaming analytics server sends notifications only when the project is in trace mode.

Column	Type	Description
key	string	The type of the update. See the following table.
value	integer	A number associated with the update, determined by the key column.
stream	string	The name of the stream if the update notifies an event related to an individual stream; otherwise, stream is NULL.
info	string	Additional information associated with the update. Its format depends on the type of the update.

The following is a table of the types of updates that the server sends as debugging commands. The Value and Stream columns in this table correspond to the value and stream rows under Column in the _ESP_RunUpdates stream.

Key	Value	Stream	Description
TRACE	0 or 1	None	Enabled (1) or disabled (0).
RUN	0 or 1	None	Streaming analytics paused (0) or running (1).
STEP	<count>	None	The project was single-stepped, either manually or automatically. The value contains the number of the steps made. No details are provided about the streams that were stepped.
BREAK	<bp-id>	<stream-name>	The breakpoint with ID <bp-id> was triggered on stream <stream-name>. These updates may come either before or after the corresponding update "RUN 0".
NO BREAK	<bp-id>	<stream-name>	A breakpoint with ID <bp-id> on the stream <stream-name> had its leftToTrigger count decreased, but has not triggered yet.
EXCEPTION	None	<stream-name>	An exception occurred on stream <stream-name>. These updates may come either before or after the corresponding update "RUN 0".
REQUESTEXIT	None	None	A request to shut down the project received.
EXIT	None	None	All the user streams have exited and the project is about to shut down.

11.4.13 _ESP_Streams

Contains information about all streams and windows.

Column	Type	Description
user_name	string	Hardcoded as user.
stream_name	string	The name of the stream described by this row.

Column	Type	Description
handle	long	The stream's numeric ID.
type	string	The type of the stream: <code>stream</code> , <code>window</code> , or <code>metadata</code> .
visibility	string	The visibility of the stream: <code>input</code> , <code>output</code> , <code>local</code> , or <code>intermediate</code> .
target	string	The name of the target stream for streams with the <code>intermediate</code> visibility value. For streams with all other visibility values, <code>target</code> is the same as the <code>stream_name</code> column.
gd_support	integer	Indicates whether guaranteed delivery is enabled for the stream or window: <ul style="list-style-type: none"> 0 – does not support guaranteed delivery 1 – supports guaranteed delivery with no checkpoint messages 2 – supports guaranteed delivery with checkpoint messages Guaranteed delivery with checkpoints is available only when the server is running in consistent recovery mode or the auto checkpoint option is enabled.

11.4.14 _ESP_Streams_Monitor

Contains information about the performance of streams, and a copy of data from the `_ESP_Streams` stream. Monitoring data is available only if the time-granularity option in the project configuration (CCR) file is set to greater than 0. The frequency of updates corresponds to the value of the time-granularity option. For example, if set to 1, an update is published every second, if set to 30, an update is published every 30 seconds, and if set to 0, reporting is disabled.

Column	Type	Description
stream_name	string	The name of the stream.
target	string	The name of the target element.
cpu_pct	float	Total CPU usage for the stream thread, as a percentage of a single CPU core.
trans_per_sec	float	The stream's performance, in transactions per second, since the last update.
rows_per_sec	float	The stream's performance, in rows per second, since the last update.
inc_trans	long	The number of transactions processed by the server since the last update.
inc_rows	long	The number of rows processed by the server since the last update.
queue	integer	The current input queue size.
store_rows	long	The current number of records in the stream's store.
last_update	bigint	The time of the current update.
sequence	long	The sequence number of the current update.

Column	Type	Description
posting_to_client	long	The numeric ID of the client connection to which the stream is trying to currently publish data. Typically, <code>posting_to_client</code> is -1, meaning the stream is not trying to currently publish data.
node_cpu_pct	float	Total CPU usage for the stream, as a percentage of all CPU cores on the machine. Total CPU usage equals system CPU usage plus user CPU usage.
node_cpu_pct_system	float	System CPU usage for the stream, as a percentage of all CPU cores on the machine.
node_cpu_pct_user	float	User CPU usage for the stream, as a percentage of all CPU cores on the machine.
cpu_time	interval	Total CPU time since the creation of the stream, in microseconds. Total CPU time equals system CPU time plus user CPU time.
cpu_time_system	interval	Total system CPU time, in microseconds, since the creation of the stream.
cpu_time_user	interval	Total user CPU time, in microseconds, since the creation of the stream.
time_since_start	interval	Duration of lapsed real time since the project was started, in microseconds.

11.4.15 _ESP_Streams_Topology

Contains pairs of names for streams that are directly connected.

Column	Type	Description
src_stream	string	The name of the source stream.
dst_stream	string	The name of the destination stream.

11.4.16 _ESP_Subscriptions

Contains information about all currently active subscriptions. A dropped connection is considered unsubscribed from everything to which it was subscribed.

Column	Type	Description
stream_handle	long	The handle the server assigns to the subscribed stream.
conn_handle	long	The handle the server assigns to the subscribed connection.

11.4.17 _ESP_Subscriptions_Ext

Contains information about all currently active subscriptions.

In some situations, there is a delay in updating this stream as new subscriptions are added or existing streams are dropped.

Column	Type	Description
stream_handle	long	The handle of the stream the server subscribes to.
conn_handle	long	The handle of the connection the server subscribes to.
stream_name	string	The name of the stream.
stream_user	string	The user name of the owner of the stream.
subscriber_user	string	The logon name of the user account that owns the subscription.
ip	string	The IP address of the client machine.
host	string	The symbolic host name of the client machine, if available. If not available, its value is the IP address of the client machine.
port	integer	The TCP port number from which the connection owning this subscription originates.
login_time	bigdatetime	The time the server accepts the connection owning the subscription.

12 Opening SAP HANA Cockpit

Access SAP HANA cockpit from a Web browser.

Prerequisites

- Your cockpit user ID is assigned to a resource group that contains the SAP HANA system you want to access.
- Your Web browser supports the SAPUI5 library `sap.m`. For more information about SAPUI5 browser support, see *SAP Note 1716423*.

Context

For more information, see *SAP HANA Administration Guide > Getting Started > SAP HANA Administration Tools > SAP HANA Cockpit*.

Procedure

1. In a browser, enter the URL for the SAP HANA cockpit and sign in. The URL takes this form:

```
https://<cockpit-host>:port
```

The port number is configured during cockpit installation.

2. Select **Resource Directory** and choose an SAP HANA system.
The System Overview for the SAP HANA system opens.

Related Information

[SAP HANA Administration Guide](#)

13 Managing Streaming Analytics with SAP HANA Cockpit

SAP HANA cockpit allows you to manage the nodes, workspaces, adapters, and projects involved in streaming analytics, to monitor streaming alerts and memory usage, to manage permissions, and to configure cluster settings.

The various streaming pages, and the tasks associated with each page, require a user to have specific permissions. Refer to the topic for each specific task to find out which privileges are necessary, and see [Managing Streaming Permissions \[page 118\]](#) for more information.

For information about the SAP HANA cockpit in general, and for details on performing SAP HANA cockpit tasks not specifically related to streaming analytics, see the *SAP HANA Administration Guide*.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

In this section:

[Managing Cluster Settings \[page 78\]](#)

The streaming cluster is the group of nodes, workspaces, projects, and adapters associated with the SAP HANA system. Streaming analytics has many preconfigured cluster settings for applications, services, and nodes, which require no additional configuration. However, there are certain settings which you can modify using the SAP HANA cockpit.

[Managing Streaming Nodes \[page 91\]](#)

Nodes run the projects you create with streaming analytics. Opening the Manage Streaming Nodes page displays all the nodes in the cluster.

[Managing Streaming Workspaces \[page 94\]](#)

Projects are deployed in workspaces on a streaming cluster. A workspace provides a namespace for the project and also allows you to control permissions at the workspace level.

[Managing Streaming Projects \[page 97\]](#)

A project defines one or more event streams and the business logic applied to incoming event data.

[Managing Streams \[page 102\]](#)

A stream processes incoming events without retaining and storing data, and produces output events according to an applied continuous query. A window receives data, but can also retain and store data.

[Managing Streaming Adapters \[page 103\]](#)

Input and output adapters enable the streaming project to receive and send data from dynamic and static external sources and destinations.

[Managing Streaming Publishers \[page 108\]](#)

Publishers are the streaming connections which transmit streaming data.

[Managing Streaming Subscribers \[page 108\]](#)

Subscribers are the streaming connections which consume streaming data.

[Monitoring Alerts \[page 109\]](#)

Administrators actively monitor the status of the system and its services, the consumption of system resources, and are also alerted of critical situations, such as a disk becoming full, CPU usage reaching a critical level, or a server stopping unexpectedly.

[Managing Streaming Permissions \[page 118\]](#)

Use the Streaming Permissions page to create, review, and remove roles, and to grant and revoke permissions and roles.

[Managing Streaming Data Services \[page 129\]](#)

The streaming analytics server accesses external databases through data service definitions, which are managed and stored in separate databases for each cluster. Use the Streaming Data Services page to manage data services.

Related Information

[SAP HANA Administration Guide](#)

[SAP HANA Streaming Analytics: Security Guide](#)

13.1 Managing Cluster Settings

The streaming cluster is the group of nodes, workspaces, projects, and adapters associated with the SAP HANA system. Streaming analytics has many preconfigured cluster settings for applications, services, and nodes, which require no additional configuration. However, there are certain settings which you can modify using the SAP HANA cockpit.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Clustering provides scalability and protects against node failure by allowing projects configured for warm standby to restart if the node stops unexpectedly. There is one node per streaming host. Multiple nodes are grouped in a cluster.

Through the Configure Streaming Cluster page, you can:

- Set a key store encrypted password.
- Adjust heartbeat settings.
- Edit node names.
- Manage settings for the Streaming Web Service and the Web Services Provider. (You can start or stop these providers on the Manage Streaming Nodes page, as described in [Managing Streaming Nodes \[page 91\]](#))

For more information about the Streaming Web Service and Web Services Provider, see the respective sections of the *SAP HANA Streaming Analytics: Adapters Guide*.

All configuration changes made on the Configure Streaming Cluster page are stored in the SAP HANA database.

In this section:

[Store an Encrypted Password for a Key Store \[page 79\]](#)

If you have encrypted a password with the `streamingencrypt` utility, you can store the encrypted password by using the SAP HANA cockpit.

[Adjust Heartbeat Settings \[page 80\]](#)

If you have determined that a performance issue is resulting from frequent heartbeat checks, adjust the heartbeat interval and the heartbeat timeout settings.

[Edit the Streaming Node Name \[page 81\]](#)

Specify a node name for a streaming analytics host using the `node_name` property.

[Edit Application Restart Settings \[page 84\]](#)

Configure settings for restarting applications on cluster startup.

[Setting Disallowed Users \[page 85\]](#)

Specify a list of users to prevent from authenticating with the streaming analytics cluster. These users are unable to connect to streaming analytics using command-line utilities, the Streaming Web Service, the Web Services Provider, and the SDKs.

[Configuring the Streaming Web Service \[page 85\]](#)

Configure the Streaming Web Service by using the SAP HANA cockpit.

[Configuring the Web Services Provider \[page 87\]](#)

Configure the Web Services Provider by using SAP HANA cockpit.

13.1.1 Store an Encrypted Password for a Key Store

If you have encrypted a password with the `streamingencrypt` utility, you can store the encrypted password by using the SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.
- You have changed and encrypted a password, or re-encrypted an encrypted password using the `streamingencrypt` utility located in the `%STREAMING_HOME/bin` folder, and are prepared to paste this encrypted password in the following procedure.

i Note

For more information, see *Encrypting Passwords for Cluster Configuration* in the *SAP HANA Streaming Analytics: Security Guide*.

Context

The streaming analytics cluster key store is a repository of security certificates for accessing streaming analytics.

During installation, streaming analytics encrypts passwords for the cluster cache, key store, and key elements. With the `streamingencrypt` utility, you can manually encrypt passwords, for example to re-encrypt a password or property using a new key file.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Under the **General** tab, in the **Key Store Encrypted Password** field, paste the encrypted password you copied from the `streamingencrypt` utility.
3. Enable or disable the **Hide Password** option by clicking the check-box.
4. Click **Save**.
5. Restart all running nodes for any changes to take effect.

13.1.2 Adjust Heartbeat Settings

If you have determined that a performance issue is resulting from frequent heartbeat checks, adjust the heartbeat interval and the heartbeat timeout settings.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Context

Every deployed streaming project broadcasts a heartbeat to a node in the cluster, so the cluster can ensure that projects are running. If the node detects too many missed heartbeats from a project, it assumes project failure and issues a stop command. This generates an error message in the log. However, it may be that the project has not failed, but instead the cluster has had difficulty responding to the request for a heartbeat check.

If you have received multiple error messages over a period of time, the existing heartbeat interval may be too short. If this is the case, adjust the heartbeat interval to fix performance issues.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Under the **General** tab, edit the heartbeat properties:

Attribute	Description
Application Heartbeat Interval	The time, in milliseconds, between heartbeat checks of a project.
Controller Heartbeat Timeout	The time, in milliseconds, for the cluster to wait for a response to a heartbeat check before timing out. The value must be at least 4500.
Controller Heartbeat Timeout Retry	The number of times to retry a heartbeat check in case of a timeout. The value must be 5 or lower.

3. Click **Save**.
4. Restart all running nodes for any changes to take effect.

13.1.3 Edit the Streaming Node Name

Specify a node name for a streaming analytics host using the **node_name** property.

Context

The system provides a default node name based on the streaming analytics hostname. When configuring system replication, specify the same node name for the primary and secondary systems. For more information about configuring system replication, see *Set Up System Replication* in the *SAP HANA Streaming Analytics: Configuration and Administration Guide*.

⚠ Caution

When configuring system replication, set the node names on all hosts before provisioning and initializing the streaming analytics services.

See [Managing Streaming Nodes \[page 91\]](#) for more information on streaming nodes.

⚠ Caution

All user and role permissions that you granted at the level of the node are affected if you change the node's name. These permissions do not migrate to the new node name, so make sure to edit user and role permissions accordingly to match the new node name.

Tenant System

Follow these steps if you have a tenant database system (previously referred to as a multiple-container system). By default, a node gets named based on the streaming analytics hostname. However, if you want to copy or move your tenant database to a different host in the future, specify a unique name for the node using the **node_name** property. If you don't do this before initializing the streaming analytics service, you will have to uninitialize, name the node, and then initialize again. Use SAP HANA studio or SQL, since you can't use SAP HANA cockpit to do this on a system with multiple tenant databases.

Rename Node Using SAP HANA Studio

Prerequisites

You are logged on as a user with the INIFILE ADMIN system privilege.

Procedure

1. In the SAP HANA studio Administration console, double-click **SYSTEMDB** in the Systems panel, and then choose the Configuration tab. A list of all configuration files appears.
2. Expand **streamingserver.ini** > **startup**.
3. Double-click the **node_name** property and specify a name.

i Note

You cannot use commas (,) or hyphens (-) in the node name.

4. Save the change.
5. Restart all running nodes for any changes to take effect.

Rename Node Using SQL

Prerequisites

You are logged on as a user with the INIFILE ADMIN system privilege.

Procedure

1. In a SQL console, connect to the SYSTEMDB.
2. Execute the following ALTER SYSTEM ALTER CONFIGURATION statement.

```
ALTER SYSTEM ALTER CONFIGURATION ('streamingserver.ini', 'host',  
'<streaming_host_name>')  
SET ('startup', 'node_name') = '<new_node_name>' WITH RECONFIGURE
```

3. Restart all running nodes for any changes to take effect.

Single-Container System

If you have a 2.0 SP 00 or earlier version, single-container SAP HANA system with streaming analytics installed, rename the node using SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your user has the CATALOG READ system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Under the **Nodes** tab, select a node name to edit, then enter a new name.

i Note

You cannot use commas (,) or hyphens (-) in the node name.

3. Click **Save**.
4. Restart all running nodes for any changes to take effect.

13.1.4 Edit Application Restart Settings

Configure settings for restarting applications on cluster startup.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have either:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Under the **Application Type** tab, edit the properties for your applications:

Attribute	Description
Auto Restart on Cluster Startup	On cluster startup, restarts applications of this type that were started or running.
Auto Restart on Cluster Startup Timeout	The maximum time, in milliseconds, for the cluster to attempt to restart the application on the same node before attempting a failover if warm standby is enabled.

3. Click **Save**.
4. Restart all running nodes for any changes to take effect.

13.1.5 Setting Disallowed Users

Specify a list of users to prevent from authenticating with the streaming analytics cluster. These users are unable to connect to streaming analytics using command-line utilities, the Streaming Web Service, the Web Services Provider, and the SDKs.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Select the **Security** tab.
3. To add a user, click the **Add** button and enter the username on the new line.
4. To remove a user, click the **Delete** button next to the username.
5. Click **Save**.
6. Restart all running nodes for any changes to take effect.

13.1.6 Configuring the Streaming Web Service

Configure the Streaming Web Service by using the SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.
2. Select the **Streaming Web Service** tab.
3. Select the node you wish to edit.
4. Set properties for the Streaming Web Service:

Property	Description
Web Service Enabled	Check to enable the Streaming Web Service. Default is checked.
Allow Origin	Contains a URL pointing to the default host, which stops the resource from being shared with other requesters from other origins. This restricts the response header of Access-Control-Allow-Origin for REST APIs. Specifying a value of * allows all origins. Only one URL can be specified, a list is not permitted.
Allocation Size (kilobytes)	Specify the memory allocation size, in kilobytes, for reading data. Default value is 64.
Auto Start Enabled	Check this option to start the Streaming Web Service automatically with the cluster. Default is unchecked.
Cipher File	Specify the location of the Streaming Web Service key file. Default value is <code>\$(STREAMING_CLUSTER_KEYSTORE)</code> .
Connection Backlog	Specify the maximum allowable number of connections in the TCP listen queue. Default value is 1024.
Echo Mode Enabled	Check to run the Streaming Web Service in echo mode. In echo mode, the Streaming Web Service confirms REST requests that it receives but does not process them. Use this property for troubleshooting. Default is unchecked.
Failed Authentications Threshold (requests/second)	The number of authentication requests per second that a client can make before the Streaming Web Service begins to throttle requests. When throttling is active, the Streaming Web Service adds a delay to failed authentication attempts. If you set this value to 1, the Streaming Web Service adds the delay to every failed attempt. The default value is 10.
Failed Authentications Delay (milliseconds)	The amount of time, in milliseconds, that the Streaming Web Service delays responses to failed authentication attempts when throttling is active. The default value is 2000.
Host	Specify the hostname for the server to listen on, or leave blank to bind on any interface server. No default value.
Monitor Interval (seconds)	Specify the interval, in seconds, between heartbeat checks from the cluster to the Streaming Web Service. Default value is 120.
Monitor Timeout (seconds)	Specify the time, in seconds, for the cluster to wait for a response to a heartbeat check before timing out. Default value is 60.
Number of Workers	Specify the number of worker processes to start. Default value is 1.
Port	Specify the port number for the server to listen on. Default value is 9093 for default (single) tenant databases. For multiple tenant databases, default values are 3xx45, 3xx54, 3xx63, 3xx72, 3xx81, and 3xx90 where xx is the SAP HANA instance number.
SSL Enabled	Check this option if the Streaming Web Service uses SSL encryption. Default is unchecked.

Property	Description
SSL Key Files Location	If you have SSL enabled, specify the location of the SSL key file and certificate. The Streaming Web Service requires unencrypted SSL key files. The certificate file must be named <code>server.crt</code> and the private key file must be named <code>server.key</code> . Default value is <code>\${CLUSTER_ROOT}/cluster/\${STREAMING_CLUSTER_NAME}/keys</code> .
Statistics Timer Interval (seconds)	Specify the period, in seconds, to output worker statistics, such as the number of requests per second, number of live requests, and average turnaround time. Default value is 0.
Subscription Buffer Queue Size	Specify the maximum number of events from a given worker process (from all Web-Socket clients and all of their subscribed streams) that can be buffered at any given time. The default value is 4096.
TLS	If you have SSL enabled, specify the TLS versions allowed for client connections in a semicolon-separated list. The default value is <code>TLSv1.2</code> .
Token Validity (seconds)	Specify the duration, in seconds, for which an authorization token remains valid. Default value is 300.
Trace Enabled	Check this option to enable detailed traces. Use this property for troubleshooting. Default is unchecked.
Trace HTTP Enabled	Check this option to trace HTTP content. Use this property for troubleshooting. Default is unchecked.

5. Click **Save**.
6. Restart all running nodes for any changes to take effect.

13.1.7 Configuring the Web Services Provider

Configure the Web Services Provider by using SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the `CATALOG READ` system privilege.
- You have one of the following:
 - The `admin system system config` permission.
 - The `read system system config` and `write system system config` permissions.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Configure Streaming Cluster**.

2. Select the **Web Services Provider** tab.
3. Select the node you wish to edit.
4. Set properties for the Web Services Provider.

The general properties section contains properties for REST, SOAP, and WebSocket connections:

Property	Description
Web Services Provider Enabled	Enables the Web Services Provider. Default is checked.
Allow Origin	Contains a URL pointing to the default host, which stops the resource from being shared with other requesters from other origins. This restricts the response header of Access-Control-Allow-Origin for REST APIs. Specifying a value of * allows all origins. Only one URL can be specified, a list is not permitted.
Auto Start Enabled	If enabled, starts the Web Services Provider automatically with the cluster. Default is unchecked.
Cache Size	Specifies the number of published rows the Web Services Provider caches internally before sending them on to streaming analytics in a batch. The default value is 10.
Data TimeZone	The timezone that the Web Services Provider uses for data timestamps. The default timezone is UTC.
Default Workspace	The default workspace associated with the cluster.
Max. Form Content Size	The maximum allowable size, in bytes, for a REST request. This value should be a positive integer. If you specify a negative number, the Jetty default value of 200,000 bytes is used instead. Any invalid values result in an integer parsing exception.
Monitor Interval (seconds)	The interval, in seconds, between heartbeat checks from the cluster to the Web Services Provider.
Monitor Timeout (seconds)	The time, in seconds, for the cluster to wait for a response to a heartbeat check before timing out.
Protocol	If SSL is enabled, set the protocol to HTTPS. If SSL is disabled, set the protocol to HTTP.
Project Loading	The maximum number of socket connections (either publish or subscribe) to streaming analytics to create for each project instance.
REST Port	The port number that the Web Services Provider uses for REST connections. The default value for a default (single) tenant database is 9091. The default values for multiple tenant databases are 3xx43, 3xx52, 3xx61, 3xx70, 3xx79, and 3xx88, where xx is the SAP HANA instance number.

Property	Description
SOAP Port	The port number that the Web Services Provider uses for SOAP connections. The default value for a default (single) tenant database is 9090. The default values for multiple tenant databases are 3xx42, 3xx51, 3xx60, 3xx69, 3xx78, and 3xx87, where xx is the SAP HANA instance number.
TLS	Sets one or more TLS versions for SOAP connections. Accepted values are TLSv1.0, TLSv1.1, and TLSv1.2. To specify multiple TLS versions, use semicolons to separate each element. For example: TLSv1;TLSv1.1;TLSv1.2. If not set, the value defaults to TLSv1.2.
Trace	Specifies the logging level for the Web Services Provider. The logging levels in increasing order of severity are: TRACE, DEBUG, INFO, and WARNING. The default value is INFO.

The WebSocket section contains properties for WebSocket connections:

Property	Description
Websocket Enabled	Enables WebSocket connections. Default is checked.
Websocket Port	The port number that the Web Services Provider uses for WebSocket connections. The default value for a default (single) tenant database is 9092. The default values for multiple tenant databases are 3xx44, 3xx53, 3xx62, 3xx71, 3xx80, and 3xx89, where xx is the SAP HANA instance number.
Publish Socket Loading	The maximum number of remote connections that can publish data over a single connection to streaming analytics. If the limit is exceeded, then a new publisher connection is created.
WebSocket Secured	Specifies whether to use HTTP or HTTPS as the protocol. If set to true, the protocol is HTTPS.
Subscribe Socket Loading	The maximum number of remote connections that can subscribe to data over a single connection from streaming analytics. If the limit is exceeded, then a new subscriber connection is created.
WebSocket TLS	Sets one or more TLS versions for WebSocket connections. Accepted values are TLSv1.0, TLSv1.1, and TLSv1.2. To specify multiple TLS versions, use semicolons to separate each element. For example: TLSv1;TLSv1.1;TLSv1.2. If not set, the value defaults to TLSv1.2.

The Thread Pool section sets properties for Jetty, used for REST, SOAP, and WebSocket connections:

Property	Description
Thread Pool Idle Timeout	The amount of time, in milliseconds, that a thread needs to idle before becoming eligible for stopping. The default value is 30000.

Property	Description
Max. Number of Threads	The maximum number of threads that Jetty may spawn to service incoming connections. The default value is 500.
Min. Number of Threads	The minimum number of threads you want Jetty to spawn. The default value is 20.

The Publisher section sets optional publishing properties for WebSocket connections:

Property	Description
Publish Block Mode	(Required if Publish Options Enabled is checked) When publishing is buffered, use this property to specify the publishing block mode. Valid values are <code>none</code> , <code>explicit</code> , and <code>auto</code> . Explicit lets you control the batches by using start transaction and end block calls, and auto ignores these calls and batches rows internally. The default mode is <code>auto</code> (autoblocking).
Publish Block Size	(Required if Publish Options Enabled is checked) The provider only uses this value if the publish block mode is <code>explicit</code> . Specifies the number of rows inside a message block. Recommended values are 64 to 1024.
Publish Buffer Size	(Required if Publish Options Enabled is checked) Specify a positive number to enable buffered publishing. When publishing is buffered, the data is first written to an internal queue. This is picked up by a publishing thread and then written to the streaming analytics project. This is the default operating mode.
Publish Options Enabled	Check to enable these advanced publisher options for WebSocket connections.
Publish Use Transactions	(Required if Publish Options Enabled is checked) If set to true, the provider uses transaction blocks instead of envelopes. The default value is false.

The Subscriber section sets optional subscribing properties for WebSocket connections:

Property	Description
Subscribe Base Drain Timeout	(Required if Subscribe Options Enabled is checked) Specify how long, in milliseconds, streaming analytics waits for a subscriber to read all base data. If there are multiple client connections sharing a subscription to streaming analytics, all connections are dropped if the Web service client is slow.
Subscribe Droppable	(Required if Subscribe Options Enabled is checked) Specify a positive number to enable buffered publishing. When publishing is buffered, the data is first written to an internal queue. This is picked up by a publishing thread and then written to the streaming analytics project. This is the default operating mode.
Subscribe Options Enabled	Check to enable these advanced publisher options for WebSocket connections.

Property	Description
Subscribe Queue Size	(Required if Subscribe Options Enabled is checked) Specify a positive number to enable buffered publishing. When publishing is buffered, the data is first written to an internal queue. This is picked up by a publishing thread and then written to the streaming analytics project. This is the default operating mode.

If you are using HTTPS, fill out the properties in the Security section. If you are using HTTP, this section is ignored. You can set security options for REST, SOAP, and WebSocket connections:

Property	Description
Cipher File	Path to a cipher file for encrypting keystore passwords.
Keystore File	Path to a keystore file containing a private key, typically a Java Keystore (JKS).
Keystore Password Encrypted	Indicates whether the keystore password is encrypted.
Keystore Password	The password for the keystore file, which can be encrypted. If encrypted, provide a symmetric cipher to decrypt the password in the Cipher File field.
Keystore Type	The type of keystore being used.

The Throttling Control section contains properties to add a delay to failed authentication attempts:

Property	Description
Failed Authentications Threshold (requests/second)	The number of authentication requests per second that a client can make before the Web Services Provider begins to throttle requests. When throttling is active, the Web Services Provider adds a delay to failed authentication attempts. If you set this value to 1, the Web Services Provider adds the delay to every failed attempt. The default value is 10.
Failed Authentications Delay (milliseconds)	The amount of time, in milliseconds, that the Web Services Provider delays responses to failed authentication attempts when throttling is active. The default value is 2000.

5. Click **Save**.
6. Restart all running nodes for any changes to take effect.

13.2 Managing Streaming Nodes

Nodes run the projects you create with streaming analytics. Opening the Manage Streaming Nodes page displays all the nodes in the cluster.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest

features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

When you select a node from the list, you can tab between the details, list of associated projects, and list of project-managed and cluster-managed adapters associated with the node.

With the node details, there are additional tabs for the Streaming Web Service (SWS) and Web Services Provider (WSP), where you can start or stop these providers. You can configure the providers on the Configure Streaming Cluster page, as described in [Managing Cluster Settings \[page 78\]](#).

For complete information on the Streaming Web Service and Web Services Provider, see the respective sections of the *SAP HANA Streaming Analytics: Adapters Guide*.

Tab	Information
Information	<p>Node - the name of the streaming analytics node.</p> <p>Host - the host on which the streaming analytics server is currently running.</p> <p>Port - the port on which the streaming analytics server is currently running.</p> <p>Status - the current status of the node.</p>
Projects	<p>Displays all projects associated with the node. Selecting one project launches the Streaming Projects page. For projects running in hot standby mode, the instance number and role are also displayed.</p> <p>Workspace - the name of logical group to which the project belongs.</p> <p>Project - the name of the project.</p> <p>Status - the current status of the project.</p>
Adapters	<p>Displays all adapters associated with the node. Selecting one adapter launches the Streaming Adapters page.</p> <p>ClusterManaged - indicates whether the adapter is managed by the cluster (YES) or by the project (NO).</p> <p>Workspace - the name of logical group to which the adapter belongs.</p> <p>Project - the name of the project the adapter is associated with. This field is blank for cluster-managed adapters.</p> <p>Adapter - the name of the adapter.</p> <p>Status - the current status of the adapter.</p>
SWS	Status - the current status of the Streaming Web Service.
WSP	Status - the current status of the Web Services Provider.

If you are configuring for system replication, you can edit node names on the Configure Streaming Cluster page. See [Edit the Streaming Node Name \[page 81\]](#).

You cannot add or remove a node using SAP HANA cockpit. Instead, see the [Installing SAP HANA Streaming Analytics](#) section in the *SAP HANA Streaming Analytics: Installation and Update Guide* guide for instructions on adding or removing a streaming analytics host.

In this section:

[Start or Stop the Streaming Web Service \[page 93\]](#)

Start and stop the Streaming Web Service through the SAP HANA cockpit.

[Start or Stop the Web Services Provider \[page 94\]](#)

Start and stop the Web Services Provider through SAP HANA cockpit.

13.2.1 Start or Stop the Streaming Web Service

Start and stop the Streaming Web Service through the SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `execute service service sws` permission.
 - The `start service service sws` and `stop service service sws` permissions.

Context

You can also start the Web service using the `--start_sws` command of the `streamingclusteradmin` utility, or stop it using the `--stop_sws` command.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Nodes**.
2. Select your node in the left navigation pane.
3. Select the **SWS** tab.
4. Select **Start** or **Stop**.
The Streaming Web Service starts or stops.

13.2.2 Start or Stop the Web Services Provider

Start and stop the Web Services Provider through SAP HANA cockpit.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have one of the following:
 - The `execute service service wsp` permission.
 - The `start service service wsp` and `stop service service wsp` permissions.

Context

You can also start the provider using the `--start_wsp` command of the `streamingclusteradmin` utility, or stop it using the `--stop_wsp` command.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Nodes**.
2. Select your node in the left navigation pane.
3. Select the **WSP** tab.
4. In the bottom right corner of the page, select **Start** or **Stop**.
The Web Services Provider starts or stops.

13.3 Managing Streaming Workspaces

Projects are deployed in workspaces on a streaming cluster. A workspace provides a namespace for the project and also allows you to control permissions at the workspace level.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

A streaming cluster must have at least one workspace. A workspace can contain projects, supporting files, adapters, and data services from multiple nodes.

The Manage Streaming Workspaces page displays all workspaces in the cluster. Select a workspace from the list to tab between the list of projects within the workspace, and the list of project-managed and cluster-managed adapters within the workspace.

Tab	Information
Projects	<p>Displays all projects associated with the workspace. Selecting one project launches the Streaming Projects page. For projects running in hot standby mode, the instance number and role are also displayed.</p> <p>Node - the name of node that the project is associated with.</p> <p>Project - the name of the project.</p> <p>Status - the current status of the project.</p>
Adapters	<p>Displays all adapters associated with the workspace. Selecting one adapter launches the Streaming Adapters page.</p> <p>Node - the name of node that the adapter is associated with.</p> <p>ClusterManaged - indicates whether the adapter is managed by the cluster (YES) or by the project (NO).</p> <p>Project - the name of the project the adapter is associated with. This field is blank for cluster-managed adapters.</p> <p>Adapter - the name of the adapter.</p> <p>Status - the current status of the adapter.</p>

In this section:

[Add a Streaming Workspace \[page 95\]](#)

Create a new workspace to manage the permissions of related objects.

[Remove a Streaming Workspace \[page 96\]](#)

Remove a workspace from the cluster.

13.3.1 Add a Streaming Workspace

Create a new workspace to manage the permissions of related objects.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `add workspace workspace <workspace-name>`

Context

Creating a workspace allows you to use it when subsequently creating projects or adapters. You can also create a new workspace while you are creating a new project or adapter.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Workspaces**.
2. Select the plus icon.
3. Enter a unique name for the workspace.
4. Click **Add**.
The workspace is added to the workspaces list.

13.3.2 Remove a Streaming Workspace

Remove a workspace from the cluster.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `remove workspace workspace <workspace-name>`

Context

i Note

You cannot delete a workspace that contains projects.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Workspaces**.
2. Highlight the workspace in the workspaces list.
3. Select **Delete**.

4. Confirm the deletion.
The workspace is removed from the workspaces list.

13.4 Managing Streaming Projects

A project defines one or more event streams and the business logic applied to incoming event data.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

It may also include adapters to establish connections to event sources, as well as destinations (including HANA tables). At its most basic, a project consists of streams and windows, and adapters.

- Adapters connect a stream or window to a data source or destination.
- A stream processes incoming events without retaining and storing data, and produces output events according to an applied continuous query.
- A window receives data, but can also retain and store data. Incoming events can add rows, update existing rows, or delete rows in the window's table.

The Manage Streaming Projects page displays all the projects in the cluster. The project list also displays the workspace associated with the project, and the project status (whether the project is running or stopped).

Projects can be run in hot standby mode. In this mode, two project instances are created—primary (0) and secondary—that run simultaneously in a cluster. To start or stop a project running in hot standby mode, start or stop each instance. For more information, see [High Availability \[page 36\]](#).

Select any project from the list to tab between system, network, streams, and adapters information. When you select a running project, the tabs are populated with information:

Tab	Statistics
System	<p>CPU History - the percentage of total, user and system CPU usage over time for the project.</p> <p>Memory Usage History - the total memory usage over time for the project.</p> <p>Thread Usage History - the number of threads used over time for the project.</p>
Network	<p>Rows Transferred History - the rows received and rows sent per second over time for the project.</p> <p>Connections History - displays total number of publisher and subscriber connections over time for the project.</p>

Tab	Statistics
Streams	<p>Displays all streams and windows associated with the project. Selecting one stream or window launches the Streaming Streams page.</p> <p>Stream - the name of the stream.</p> <p>Type - the type of stream. Possible values are stream or window.</p> <p>Visibility - the visibility of the stream. Valid values are input, output, local, or intermediate.</p> <p>Queue Depth - the number of rows waiting to be processed.</p> <p>Rows In Store - the current number of records in the stream's store.</p>
Adapters	<p>Displays all adapters associated with the project. Selecting one adapter launches the Streaming Adapters page.</p> <p>StreamName - the stream in which the adapter operates.</p> <p>Adapter - the name of the adapter.</p> <p>Type - the unique adapter ID assigned to each adapter.</p> <p>IsInput - indicates whether the adapter is an input or output adapter.</p> <p>Status - the current status of the adapter.</p>

The data on the graphs begin as of the time you create the project.

When you select a project running in hot standby mode, the instance details display above the project details. You can tab between instances to see statistics, and start or stop the instance.

Element	Description
Node	The node running this instance.
Active Active Role	Indicates whether this instance is primary or secondary instance.
Requested Status	The expected status of the instance. Valid values are stopped or started-running .
Current Status	The actual status of the instance. Valid values are start , starting , started , started-initializing , started-running , stopping , stopped , stopped-failed , or stopped-failed-start .

In this section:

[Add a Streaming Project \[page 99\]](#)

Use SAP HANA cockpit to create a streaming analytics project.

[Remove a Streaming Project \[page 100\]](#)

Use SAP HANA cockpit to remove a project from the cluster.

[Start or Stop a Streaming Project \[page 101\]](#)

Use SAP HANA cockpit to start or stop a streaming project.




13.4.1 Add a Streaming Project

Use SAP HANA cockpit to create a streaming analytics project.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `add application application|workspace <application-name>|<workspace-name>`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Projects**.
2. Click .
3. Enter or select the name of the existing or new workspace within which to create the project.
4. Enter a unique name for the project.
5. Click  to browse to and specify the project configuration CCX file. This file contains the properties and commands used to start and stop the project, as well as other parameters to configure the project from within SAP HANA studio.
6. (Optional) Click  to browse to and specify the project runtime properties CCR file. This file specifies overrides to the project default properties.
7. Click **Add**.
The project is added to the projects list.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.4.2 Remove a Streaming Project

Use SAP HANA cockpit to remove a project from the cluster.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `remove application application|workspace <application-name>|<workspace-name>`

Context

Deleting the project also deletes any adapters that are managed by the project.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Projects**.
2. Highlight the project in the projects list.
3. Select **Delete**.
4. (Optional) To delete persistency files, check **Remove project and adapter log stores, state and metadata**.
5. Confirm deletion.
The project is removed from the projects list.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.4.3 Start or Stop a Streaming Project

Use SAP HANA cockpit to start or stop a streaming project.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `start|stop application application|workspace <application-name>|<workspace-name>`

Context

Adapters that are not managed by the cluster are started and stopped as the project starts or stops. Starting or stopping a project adds or removes these adapters from the list on the Streaming Adapters page.

i Note

To start or stop a project configured for hot standby, select **Start** or **Stop** for each project instance.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Projects**.
2. Highlight the project in the projects list.
3. Select **Start** or **Stop**.
4. Confirm your selection.
The project starts or stops.
5. (Optional) To force stop a frozen project, select the project from the projects list and click **Force Stop**.

If your project is hanging and you attempt to stop it, the project status remains at "stopping" and the project and adapter processes keep running. Force the project to stop, which immediately shuts down all project related processes.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.5 Managing Streams

A stream processes incoming events without retaining and storing data, and produces output events according to an applied continuous query. A window receives data, but can also retain and store data.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Streams subscribe to incoming events and process the event data to publish output events. Since they are stateless, they cannot retain data—and they use little memory because they do not store events. Input streams are the point at which data enters the project from external sources through adapters. A project may have any number of input streams. Output streams take their input from other streams or windows in a project and they apply a continuous query to produce their output.

By contrast, a window consists of a table where incoming events add rows, update existing rows, or delete rows. Windows are required for performing aggregate operations not available on streams.

On the Manage Streaming Projects page, selecting the streams tab when you have highlighted a running project displays all streams and windows associated with the project. Selecting one stream or window launches the Streaming Streams page. Stream Details displays additional information about the stream or window:

Tab	Statistics
Information	<p>Node - the name of node that the stream is associated with.</p> <p>Workspace - the name of the workspace that the stream is associated with.</p> <p>Project - the name of the project that the stream is associated with.</p> <p>Visibility - the visibility of the stream. Valid values are input, output, local, or intermediate.</p> <p>Queue Depth - the number of rows waiting to be processed.</p> <p>Rows In Store - the current number of records in the stream's store.</p>
System	<p>CPU Usage History - the percentages of total, user, and system CPU usage over time for the stream.</p>
Network	<p>Rows/Transactions Throughput History - the rows per second and transactions per second over time for the stream.</p>

13.6 Managing Streaming Adapters

Input and output adapters enable the streaming project to receive and send data from dynamic and static external sources and destinations.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Input adapters connect to an external datasource, and translate incoming messages from the external sources into a format that is accepted by the streaming project. Output adapters translate rows processed by the streaming project into message formats compatible with external destinations, and send those messages downstream. External sources or destinations can include data feeds, sensor devices, messaging systems, RFID reader, e-mail servers, or relational databases.

SAP HANA streaming analytics includes an extensive set of input and output adapters that you can use to subscribe to and publish data. It also provides an adapter toolkit, SDKs, and an internal adapter API that you can use to write a custom adapter. Adapters that use the adapter toolkit can be managed through the cluster, instead of through a specific project.

Open the Manage Streaming Adapters page to display all the adapters managed by the cluster, and all adapters on running projects. Select one adapter from the list to tab between adapter detail information, common statistics, and custom statistics.

Tab	Statistics
Information	<p>Adapter - the name of the adapter.</p> <p>Type - the unique adapter ID assigned to each adapter.</p> <p>Status - valid values are running or stopped.</p> <p>Workspace - the name of logical group to which the adapter belongs.</p> <p>Project - the name of the project the adapter is associated with. For projects running in hot standby mode, the instance number and role are also displayed.</p> <p>StreamName - the stream in which the adapter operates.</p> <p>IsInput - indicates whether the adapter is an input or output adapter.</p> <p>ClusterManaged - indicates whether the adapter is managed by the cluster (YES) or by the project (NO).</p>

Tab	Statistics
Common Statistics	<p>Displays the key performance indicators (KPIs) for the adapter.</p> <p>AdapterTotalRows - the total number of data records recognized in the data.</p> <p>AdapterGoodRows - the number of data records successfully processed.</p> <p>AdapterBadRows - the number of data records that experienced errors.</p> <p>AdapterLastErrDate - the date of the error occurred. The format is <code>dow mon dd hh:mm:ss zzz yyyy</code>, where:</p> <ul style="list-style-type: none"> • <code>dow</code> is the day of the week • <code>mon</code> is the month (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec) • <code>dd</code> is the day of the month (01 through 31) • <code>hh</code> is the hour of the day (00 through 23) • <code>mm</code> is the minute within the hour (00 through 59) • <code>ss</code> is the second within the minute (00 through 61) • <code>zzz</code> is the time zone (and may reflect daylight saving time). Standard time zone abbreviations include those recognized by the method <code>parse</code>. If time zone information is not available, then <code>zzz</code> is empty meaning it consists of no characters at all. • <code>yyyy</code> is the year <p>AdapterLastErrMsg - the complete text of the error message as written to the log.</p> <p>AdapterLatency - the time taken for data to be processed. For an input adapter, this is the amount of time it takes the adapter to receive data from its source and publish the data to the stream. For an output adapter, this is the amount of time it takes for the adapter to receive a message from the stream and publish the data to its destination. The common statistics are the key performance indicators (KPIs) for the adapter. The custom statistics are displayed only if the selected adapter is configured to publish them.</p>
Custom Statistics	<p>Displays any custom statistics defined for the adapter. Note that not all adapters have custom statistics, or are configured to publish them. For more information, see <i>SAP HANA Streaming Analytics: Adapters Guide</i>.</p>

In this section:

[Add an Adapter \[page 105\]](#)

Use the SAP HANA cockpit to create a cluster-managed adapter.

[Remove an Adapter \[page 106\]](#)

Use the SAP HANA cockpit to remove a cluster-managed adapter from the cluster.

[Start or Stop an Adapter \[page 107\]](#)

Use SAP HANA cockpit to start or stop a cluster-managed or a project-managed adapter.

13.6.1 Add an Adapter

Use the SAP HANA cockpit to create a cluster-managed adapter.

Prerequisites




- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `add application application|workspace <application-name>|<workspace-name>`

Context

This method can add an adapter that is not managed by a project.

If you want to add an adapter that will start and stop with a project, attach the adapter to a project using the ATTACH ADAPTER statement and edit its properties using studio.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Adapters**.
2. Click .
3. Do one of the following:
 - Select an existing workspace.
 - Enter the name of a new workspace within which to create the adapter.
4. Enter a unique name for the adapter.
5. Click  to browse to and specify the adapter configuration XML file. This file contains the properties and commands used to start and stop the adapter, as well as other information that allows you to configure the adapter from SAP HANA studio. The adapter configuration file also constructs the name for your custom adapter DLL file. The library name is referenced when you load your adapter.
6. (Optional) Click  on the Adapter Deployment File line to browse to and specify the adapter deployment file. This file specifies whether warm standby and affinity are enabled when the adapter is deployed to the cluster.
7. Click **Add**.
The adapter is added to the adapters list.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.6.2 Remove an Adapter

Use the SAP HANA cockpit to remove a cluster-managed adapter from the cluster.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `remove application application|workspace <application-name>|<workspace-name>`

Context

This method can only remove adapters that are not managed by projects. You can't use the cockpit to delete adapters that start and stop with a project, as that requires editing the project configuration file, which SAP HANA cockpit cannot access.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Adapters**.
2. Highlight the adapter in the adapters list.

i Note

You can only delete cluster-managed adapters through SAP HANA cockpit.

3. Select **Delete**.
4. (Optional) To delete persistency files, check **Remove adapter state and metadata**.
5. Confirm deletion.
The adapter is removed from the adapters list.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.6.3 Start or Stop an Adapter

Use SAP HANA cockpit to start or stop a cluster-managed or a project-managed adapter.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `start|stop application application|workspace <application-name>|<workspace-name>`

Context

Project-managed adapters (adapters that are not managed by the cluster) are started and stopped as the project starts or stops. Starting or stopping a project adds or removes these project-managed adapters from the list on the Streaming Adapters page.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Adapters**.
2. Highlight the adapter in the adapters list.
3. Select **Start** or **Stop**.
4. If required, the password input dialog displays. Enter your login ID and password.
5. Confirm your selection.
The adapter starts or stops.

Related Information

[Managing Streaming Permissions \[page 118\]](#)

13.7 Managing Streaming Publishers

Publishers are the streaming connections which transmit streaming data.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

A publisher can be a streaming project or an external data source (such as a financial market feed). Data from external sources is typically delivered through an adapter.

Opening the Manage Streaming Publishers page displays all the publishers in the cluster. Select a publisher from the list to view the CPU usage and row history (throughput) associated with that streaming connection.

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph.
Rows Transferred History	Line graph displaying rows sent per second.

13.8 Managing Streaming Subscribers

Subscribers are the streaming connections which consume streaming data.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

A subscriber consumes data produced by a publisher. A subscriber can be a streaming project, an external adapter, or a third-party application.

Opening the Manage Streaming Subscribers page displays all the subscribers in the cluster. Select a subscriber from the list to view the CPU usage and row history (throughput) associated with that streaming connection.

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph.
Rows Transferred History	Line graph displaying rows sent per second.

13.9 Monitoring Alerts

Administrators actively monitor the status of the system and its services, the consumption of system resources, and are also alerted of critical situations, such as a disk becoming full, CPU usage reaching a critical level, or a server stopping unexpectedly.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Alerts appear on the system overview of the SAP HANA cockpit. To view alert details, select an alert category, or select **Show all** to see all alerts.

The internal monitoring infrastructure of the SAP HANA database continuously collects and evaluates information about status, performance, and resource usage from all its components. In addition, it performs regular checks on the data in system tables and views, and when configurable threshold values are exceeded, issues alerts. In this way, you are proactively warned of potential problems.

The priority of the alert indicates the severity of the problem and depends on the nature of the check and configured threshold values. For example, if 90% of available disk space is used, a low priority alert is issued; if 98% is used, a high priority alert is issued. For more information about the technical implementation of monitoring and alerting features in SAP HANA, see *The Statistics Service* in the *SAP HANA Administration Guide*.

In this section:

[Alert Details \[page 109\]](#)

When you select an alert in the Alerts page, detailed information about the alert is displayed on the right.

[Streaming Analytics Alerts in SAP HANA Cockpit \[page 110\]](#)

Monitor alerts that are specific to streaming analytics.

13.9.1 Alert Details

When you select an alert in the Alerts page, detailed information about the alert is displayed on the right.

Detail	Description
Category	The category of the check that issued the alert. Checks are grouped into categories, for example, memory usage, transaction management.

Detail	Description
Next Scheduled Run	The next time the related check is scheduled to run. If the check has been switched (check status Switched Off), or it failed the last time it ran (check status Failed), this field is empty because the check is no longer scheduled.
Interval	The frequency with which the related check runs. If the check has been switched (check status Switched Off), or it failed the last time it ran (check status Failed), this field is empty because the check is no longer scheduled.
Check	Name and description of the related check.
Proposed Solution	Possible ways of resolving the problem identified in the alert.
Past Occurrences of Alert	Configurable graphical display indicating how often the alert occurred in the past.

13.9.2 Streaming Analytics Alerts in SAP HANA Cockpit

Monitor alerts that are specific to streaming analytics.

Project Alerts

Detail	Description
Alert Name	Streaming analytics License Expiry
ID	93
Short Description	Identifies a streaming analytics license close to expiry.
Alert Text	Your streaming analytics license will expire in <code><days_before_license_expiration></code> days. Once your license expires, you can no longer start streaming analytics projects.
Proposed Solution	Obtain a valid license and install it. For the exact expiration date, see the monitoring view M_LICENSES , or navigate to Manage Licenses .
Interval	1 day
Threshold	Low Priority < 30 days Medium Priority < 14 days High Priority < 7 days
Source of information	SYS.M_LICENSES

Detail	Description
Alert Name	Inactive Streaming Applications

Detail	Description
ID	600
Short Description	Identifies inactive streaming analytics applications such as projects, or adapters managed by the cluster rather than the project.
Alert Text	The status of 'streaming project'/'streaming cluster managed adapter' (<schema-name>/<application-name>) has been inactive for xxxx seconds.
Proposed Solution	Investigate why the streaming analytics application is inactive: <ul style="list-style-type: none"> Check the application's trace files. See <i>Project Log Files</i> and <i>Adapter Logging Configuration</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i> for more information. For more information on troubleshooting project errors, <i>Project Problems</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i>.
Interval	60 seconds
Threshold	Warning3 = Duration time at inactive state > 60 seconds
Source of information	SYS.M_STREAMING_APPLICATIONS system view

Detail	Description
Alert Name	Streaming Project Physical Memory Usage
ID	602
Short Description	Determines the percentage of total physical memory available on the host that is used for the streaming project.
Alert Text	The streaming project <schema-name>/<project-name> on <host> is running out of physical memory. Total available physical memory: xxxx MB, total used physical memory: xxxx MB.
Proposed Solution	Investigate the memory usage of the streaming project: <ul style="list-style-type: none"> Generate memory reports to see which components of the project are consuming the most memory. See <i>Monitor Project Memory Usage</i> in the <i>SAP HANA Streaming Analytics: Developer Guide</i> for information.
Interval	60 seconds
Threshold	Warning1 = PROJECT_USED_PHYSICAL_MEMORY_PERC> 80 Warning 2 = PROJECT_USED_PHYSICAL_MEMORY_PERC> 85 Warning3 = PROJECT_USED_PHYSICAL_MEMORY_PERC> 90
Source of information	SYS.M_STREAMING_PROJECTS system view

Detail	Description
Alert Name	Streaming Project CPU Usage
ID	603
Short Description	Determines the percentage CPU usage for a streaming project on the host and therefore whether or not CPU resources are running out.

Detail	Description
Alert Text	The streaming project schema_name/project_name on <host> runs out of CPU resources! CPU consumption user mode <cpu-user-utilization>%, system mode <cpu-system-utilization>%
Proposed Solution	Investigate the CPU usage. See <i>CPU Usage</i> in the <i>SAP HANA Streaming Analytics: Developer Guide</i> for information.
Interval	60 seconds
Threshold	Warning1 = PROJECT_CPU_UTILIZATION_PERC > 80 Warning2 = PROJECT_CPU_UTILIZATION_PERC > 90 Warning3 = PROJECT_CPU_UTILIZATION_PERC > 95
Source of information	SYS.M_STREAMING_PROJECTS system view

Streaming Stream Alerts

Detail	Description
Alert Name	Large Queue of Stream of Streaming Project
ID	612
Short Description	Identifies which stream or window in the streaming project has a large queue.
Alert Text	The stream <stream-name> of streaming project <schema-name>/<project-name> has large queue.
Proposed Solution	Investigate why the stream or window has a large queue: <ul style="list-style-type: none"> See Streaming Data Flow for more information on what can lead to a queue backup, and an example. Use watch variables in your stream to help debug issues. See Watch Variables for usage instructions.
Interval	60 seconds
Threshold	Warning1 = STREAM_QUEUE_SIZE > 1000
Source of information	SYS.M_STREAMING_PROJECT_STREAMS system view

Detail	Description
Alert Name	Large Store of Stream of Streaming Project
ID	613
Short Description	Identifies which stream or window in the streaming project has a large store.
Alert Text	The stream <stream-name> of streaming project <schema-name>/<project-name> has large store.

Detail	Description
Proposed Solution	Investigate why the stream or window has a large store: <ul style="list-style-type: none"> Check the retention policy of the window, and consider modifying it to retain less information at a time. See Retention Policies for more information .
Interval	60 seconds
Threshold	Warning1 = STREAM_STORE_SIZE > 10000 rows
Source of information	SYS.M_STREAMING_PROJECT_STREAMS system view

Adapter Alerts

Detail	Description
Alert Name	Inactive Streaming Applications
ID	600
Short Description	Identifies inactive streaming analytics applications such as projects, or adapters managed by the cluster rather than the project.
Alert Text	The status of 'streaming project'/'streaming cluster managed adapter' (<schema-name>/<application-name>) has been inactive for xxxx seconds.
Proposed Solution	Investigate why the streaming analytics application is inactive: <ul style="list-style-type: none"> Check the application's trace files. See <i>Project Log Files</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i>, and <i>Adapter Logging Configuration</i> in the <i>SAP HANA Streaming Analytics: Building Custom Adapters</i> guide for more information. For more information on troubleshooting project errors, <i>Project Problems</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i>.
Interval	60 seconds
Threshold	Warning3 = Duration time at inactive state > 60 seconds
Source of information	SYS.M_STREAMING_APPLICATIONS system view

Detail	Description
Alert Name	Inactive Streaming Project Managed Adapters
ID	601
Short Description	Identifies inactive streaming analytics adapters that are managed by projects rather than by the cluster.
Alert Text	The status of project managed adapter (<schema-name>/<project-name>/<adapter-name>) has been inactive for xxxx seconds.

Detail	Description
Proposed Solution	Investigate why the adapter is inactive: <ul style="list-style-type: none"> Check the log files of the adapter. See Adapter Logging Configuration for more information.
Interval	60 seconds
Threshold	Warning3 = Duration time at inactive state > 60 seconds
Source of information	SYS.M_STREAMING_PROJECT_ADAPTERS system view

Detail	Description
Alert Name	Bad Rows of Project Managed Adapter
ID	610
Short Description	Identifies a project managed adapter that has many rows with errors.
Alert Text	The project managed adapter <code><adapter-name></code> of streaming project <code><schema-name>/<project-name></code> has many rows with errors.
Proposed Solution	Investigate why the adapter has so many rows with errors: <ul style="list-style-type: none"> Ensure your adapter is able to convert rows passed to it into the correct format. View your server log file to find the errors logged for each bad row. See <i>File and Directory Infrastructure</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i> for more information on file locations. To learn more about adapter row types, see the corresponding section for your adapter type in the <i>SAP HANA Streaming Analytics: Adapters Guide</i>.
Interval	60 seconds
Threshold	Warning1 = ADAPTER_BAD_ROW_COUNT > 100
Source of information	SYS.M_STREAMING_PROJECT_ADAPTERS system view

Detail	Description
Alert Name	High Latency of Project Managed Adapter
ID	611
Short Description	Identifies a project managed adapter that has high latency.
Alert Text	The project managed adapter <code><adapter-name></code> of streaming project <code><schema-name>/<project-name></code> has high latency.
Proposed Solution	Investigate why the adapter has high latency: <ul style="list-style-type: none"> For information on maximizing performance and balancing throughput with latency, see Appendix: Performance and Tuning Tips. Consider changing the properties of the adapter to improve latency. See the corresponding section for your adapter type in the <i>SAP HANA Streaming Analytics: Adapters Guide</i> to learn more about the properties.
Interval	60 seconds
Threshold	Warning1 = ADAPTER_LATENCY > 1000000 microseconds

Detail	Description
Source of information	SYS.M_STREAMING_PROJECT_ADAPTERS system view

Publisher Alerts

Detail	Description
Alert Name	Number of Publishers of Streaming Project
ID	604
Short Description	Identifies a high number of publishers of the streaming project to ensure that they will not break the project.
Alert Text	The number of publishers in the streaming project <code><schema-name>/<project-name></code> exceeds the limit.
Proposed Solution	Investigate whether these publishers were created intentionally: <ul style="list-style-type: none"> Open the Cockpit Manage Streaming Publishers page to view all publishers in the cluster. See Managing Streaming Publishers [page 108] for more information.
Interval	60 seconds
Threshold	Warning1 = PROJECT_PUBLISHER_COUNT > 100000 Warning2 = PROJECT_PUBLISHER_COUNT > 1000000 Warning3 = PROJECT_PUBLISHER_COUNT > 10000000
Source of information	SYS.M_STREAMING_PROJECT_PUBLISHERS system view

Detail	Description
Alert Name	Row Throughput of Publisher of Streaming Project
ID	608
Short Description	Identifies a publisher in the streaming project that has low throughput, measured in rows per second.
Alert Text	The publisher <code><client-host>:<client-port></code> of streaming project <code><schema-name>/<project-name></code> has low row throughput.
Proposed Solution	Investigate why the publisher works slowly: <ul style="list-style-type: none"> Check which publisher is working slowly. See Managing Streaming Publishers [page 108] for more information. Monitor the streaming project being published to to determine if any elements in it are experiencing a bottleneck. See <i>Monitoring Streaming Analytics</i> and <i>CPU Usage</i> in the <i>Appendix: Performance and Tuning Tips</i> section of the <i>SAP HANA Streaming Analytics: Developer Guide</i> for more information.
Interval	60 seconds
Threshold	Warning1 = PUBLISHER_ROW_THROUGHPUT < 1000 rows per second

Detail	Description
Source of information	SYS.M_STREAMING_PROJECT_publisher system view

Detail	Description
Alert Name	Transaction Throughput of Publisher of Streaming Project
ID	609
Short Description	Identifies a publisher in the streaming project that has low throughput, measured in transactions per second.
Alert Text	The publisher <client-host>:<client-port> of streaming project <schema-name>/<project-name> has low transaction throughput.
Proposed Solution	Investigate why the publisher works slowly: <ul style="list-style-type: none"> • Check which publisher is working slowly. See Managing Streaming Publishers [page 108] for more information. • Monitor the streaming project being published to to determine if any elements in it are experiencing a bottleneck. See Monitoring Streaming Analytics [page 54] and CPU Usage.
Interval	60 seconds
Threshold	Warning1 = PUBLISHER_TRANSACTION_THROUGHPUT < 1000 transactions per second
Source of information	SYS.M_STREAMING_PROJECT_publisher system view

Subscriber Alerts

Detail	Description
Alert Name	Number of Subscribers of Streaming Project
ID	605
Short Description	Identifies a high number of subscribers in the streaming project to ensure that they will not break the project.
Alert Text	The number of subscribers in the streaming project <schema-name>/<project-name> exceeds the limit.
Proposed Solution	Investigate whether these subscribers were created intentionally: <ul style="list-style-type: none"> • Open the Cockpit Manage Streaming Subscribers page to view all subscribers in the cluster. See Managing Streaming Subscribers [page 108] for more information.
Interval	60 seconds
Threshold	Warning1 = PROJECT_SUBSCRIBER_COUNT > 100000 Warning2 = PROJECT_SUBSCRIBER_COUNT > 1000000 Warning3 = PROJECT_SUBSCRIBER_COUNT > 10000000
Source of information	SYS.M_STREAMING_PROJECT_SUBSCRIBER system view

Detail	Description
Alert Name	Row Throughput of Subscriber of Streaming Project
ID	606
Short Description	Identifies a subscriber in the streaming project that has low throughput, measured in rows per second.
Alert Text	The subscriber <client-host>:<client-port> of streaming project <schema-name>/<project-name> has low row throughput.
Proposed Solution	Investigate why the subscriber works slowly: <ul style="list-style-type: none"> • Check which subscriber is working slowly. See Managing Streaming Subscribers [page 108] for more information. • Monitor the streaming project being subscribed to to determine if any elements in it are experiencing a bottleneck. See <i>Monitoring Streaming Analytics</i> in the <i>SAP HANA Streaming Analytics: Configuration and Administration Guide</i>, and <i>CPU Usage</i> in the <i>Appendix: Performance and Tuning Tips</i> section of the <i>SAP HANA Streaming Analytics: Developer Guide</i> for more information.
Interval	60 seconds
Threshold	Warning1 = SUBSCRIBER_ROW_THROUGHPUT < 1000 rows per second
Source of information	SYS.M_STREAMING_PROJECT_SUBSCRIBER system view

Detail	Description
Alert Name	Transaction Throughput of Subscriber of Streaming Project
ID	607
Short Description	Identifies a subscriber in the streaming project that has low throughput, measured in transactions per second.
Alert Text	The subscriber <client-host>:<client-port> of streaming project <schema-name>/<project-name> has low transaction throughput.
Proposed Solution	Investigate why the subscriber works slowly: <ul style="list-style-type: none"> • Check which subscriber is working slowly. See Managing Streaming Subscribers [page 108] for more information. • Monitor the streaming project being subscribed to to determine if any elements in it are experiencing a bottleneck. See Monitoring Streaming Analytics [page 54] and CPU Usage for more information.
Interval	60 seconds
Threshold	Warning1 = SUBSCRIBER_TRANSACTION_THROUGHPUT < 1000 transactions per second
Source of information	SYS.M_STREAMING_PROJECT_SUBSCRIBER system view

For more information on accessing project trace files (also referred to as log files), see *Project Log Files* in the *SAP HANA Streaming Analytics: Configuration and Administration Guide*.

13.10 Managing Streaming Permissions

Use the Streaming Permissions page to create, review, and remove roles, and to grant and revoke permissions and roles.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Overview of Permissions

Permissions determine the operations that a user may perform when accessing SAP HANA streaming analytics through a streaming client interface. These client interfaces include the streaming analytics plugin for SAP HANA studio, SAP HANA cockpit, command line utilities, an SDK, or an adapter.

Manage a user's access to and control over SAP HANA streaming analytics by giving them the permissions necessary to complete specific tasks. You can do that by directly granting individual permissions to a user, or grouping permissions into roles.

The syntax of a permission consists of a privilege, privilege type, resource type, and an optional resource:

- **Privilege:** the action being permitted.
- **Privilege type:** the scope or range of a particular privilege's action.
- **Resource type:** the location where the permitted action takes place.
- **Resource:** (optional) a specific instance of a resource type. This can be any case sensitive name.

If you are granting a permission that applies to all possible instances of a resource, such as all workspaces, leave the **Resource** field blank.

When you do specify a resource, make sure that the name is identical to what it is called in CCL. For example, to specify the File/Hadoop CSV Input adapter, enter `toolkit_file_csv_input` as your resource.

Privileges

Privilege	Description
all	<p>Grants all privileges to a user, either for all resources or for a specific resource.</p> <p>For example, you can give someone permission to do anything in the streaming system:</p> <ul style="list-style-type: none">• Using the cockpit: <code>all all all</code>• Using <code>streamingclusteradmin:grant perm all all on all</code>

Privilege	Description
admin	<p>Grants the user administrative privileges on the specified system resource, which can be either config or policy.</p> <p>For example, you can give someone permission to edit authorization policies:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>admin system system policy</code> • Using <code>streamingclusteradmin:grant perm admin system on system policy</code>
add	<p>Allows users to add resources of the specified privilege type to the specified resource type. Can be restricted further to a particular resource type or resource.</p> <p>For example, you can give someone permission to add applications (projects or adapters) anywhere in the system:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>add application all</code> • Using <code>streamingclusteradmin:grant perm add application on all</code>
remove	<p>Allows users to remove the specified privilege type from a specified resource type. Can be restricted by resource name, as in the case where you need to remove a specific project from a workspace.</p> <p>For example, you can give someone permission to remove applications (projects or adapters) from anywhere in the system:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>remove application all</code> • Using <code>streamingclusteradmin:grant perm remove application on all</code>
start	<p>Allows users to start the specified privilege type. Can be restricted further by resource type and instance name.</p> <p>For example, you can give someone permission to start an application in the workspace default:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>start application workspace default</code> • Using <code>streamingclusteradmin:grant perm start application on workspace default</code>
stop	<p>Allows users to stop the specified privilege type on a specified resource type. Can be restricted by resource name.</p> <p>For example, you can give someone permission to stop an application in the workspace default:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>stop application workspace default</code> • Using <code>streamingclusteradmin:grant perm stop application on workspace default</code>

Privilege	Description
control	<p>Grants users control over the runtime for managed adapters within a project.</p> <p>For example, you can give someone control over starting or stopping a project-managed adapter in the project <code>project1</code>:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>control project application default/project1</code> • Using <code>streamingclusteradmin:grant perm control project on application default/project1</code>
execute	<p>Allows users to start and stop a specific service, or all services.</p> <p>For example, you can give someone permission to start and stop the Streaming Web Service:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>execute service service sws</code> • Using <code>streamingclusteradmin:grant perm execute service on service sws</code>
read	<p>Allows users read access to resources of the specified privilege type. Can be restricted by resource type and instance, as in the case where you require read access for a stream.</p> <p>For example, you can give someone permission to read applications (projects or adapters) in the workspace <code>default</code>:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>read application workspace default</code> • Using <code>streamingclusteradmin:grant perm read application on workspace default</code>
view	<p>Allows users to view resources of the specified privilege type. Can be restricted by resource type and instance.</p> <p>For example, you can give someone permission to view applications (projects or adapters) in the workspace <code>default</code>:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>view application workspace default</code> • Using <code>streamingclusteradmin:grant perm view application on workspace default</code>
write	<p>Allows users to write to or modify elements of the specified privilege type. Can be restricted by resource type and instance.</p> <p>For example, you can give someone permission to write to applications (projects or adapters) in the workspace <code>default</code>:</p> <ul style="list-style-type: none"> • Using the cockpit: <code>write application workspace default</code> • Using <code>streamingclusteradmin:grant perm write application on workspace default</code>

Privilege Types and Resource Types

Privilege or Resource Type	Description
stream	An element in a project for streaming analytics that processes incoming events and produces output events.
adapter	Connects streaming analytics projects to external data sources and destinations, and handles any required data conversion.
project	A set of event streams, data sources, and the continuous queries that process incoming data to produce the desired information.
application	A project from the view of the host machine.
dataservice	A connection to a database for reading or writing data.
workspace	One of the subdirectories within the working directory for studio.
node	A host machine, running streaming analytics, belonging to a cluster.
cluster	A group of one or more nodes that are managed as a unit.
service	<p>A stateless application that performs predefined functions for streaming analytics.</p> <p>The streaming analytics services are:</p> <ul style="list-style-type: none"> • <code>discovery</code> • <code>compiler</code> • <code>sws</code> (Streaming Web Service) • <code>wsp</code> (Web Services Provider) • <code>palmodelservice</code> (machine learning models) <p>All services are case-sensitive.</p>
system	<p>Specifies a privilege or resource type for administering the cluster.</p> <p>The system resource names are:</p> <p>config Access to edit cluster configuration</p> <p>policy Grant permissions to any users or roles (except to themselves)</p>
all	Includes all of the privilege types listed here.

Manage Streaming Permissions Page

In the SAP HANA cockpit, clicking on the Manage Streaming Permissions link takes you to the streaming permissions page, where you can manage users, roles, and their permissions through two tabs:

Tab	Information
Users	<p>Displays all users that have been granted at least one streaming analytics permission or role.</p> <ul style="list-style-type: none">• Permissions View - lists the permissions granted to the selected user.• Roles View - lists the roles that have been assigned to the selected user.
Roles	<p>Displays all roles created for streaming analytics.</p> <ul style="list-style-type: none">• Permissions View - lists the permissions granted to the selected role.

i Note

- To view and access the Manage Streaming Permissions page in the SAP HANA cockpit, make sure that you are either logged in as the SYSTEM or SYS_STREAMING user, or use an account that has been granted the permission equivalent to `grant perm admin all on all or grant perm admin system on system policy`.
- You may be asked to re-enter your login password while you use the Streaming Permissions page, due to periodic connection check-ins between SAP HANA cockpit and the streaming analytics cluster.

Reviewing User and Role Authorizations

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the user or role you want to review.

i Note

The **Users** tab only lists users that have been granted at least one permission or role.

3. Use the **User Details** and **Role Details** views to review user and role authorizations.
 - To review an individual user's permissions, open the **Users** tab, then select a user. The permissions granted to that user appear in the **Permissions** tab of the **User Details** view.
 - To review the roles that have been assigned to that user, select the **Roles** tab in the **User Details** view.
 - To review the permissions granted to a specific role, open the **Roles** tab, then select a role. The permissions granted to that role appear in the **Permissions** tab of the **Role Details** view.

In this section:

[Grant a Permission to a User \[page 123\]](#)

Control a user's access to and control over streaming analytics by giving them the permissions necessary to complete specific tasks.

[Revoke a Permission from a User \[page 124\]](#)

Revoke permissions that are no longer needed by the user.

[Create a Role \[page 125\]](#)

Add a blank role to the streaming roles list. Afterward, you can grant the role permissions.

[Remove a Role \[page 126\]](#)

Remove a collection of permissions that are no longer required.

[Grant a Permission to a Role \[page 126\]](#)

Define a set of privileges required to perform a particular task. You can then grant the role or multiple roles to any streaming user.

[Revoke a Permission from a Role \[page 127\]](#)

Revoke permissions from a role whenever the requirements for the related task change.

[Grant a Role to a User \[page 128\]](#)

A role is a set of permissions. You can grant a role to user instead of individual permissions, since a role normally contains all the permissions necessary to perform a certain task.

[Revoke a Role from a User \[page 129\]](#)

Revoke a role from a user when the user no longer needs to perform a particular task.

13.10.1 Grant a Permission to a User

Control a user's access to and control over streaming analytics by giving them the permissions necessary to complete specific tasks.

Prerequisites


- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Context

For a user's granted permissions to take effect, that user must be authenticated by the SAP HANA system. For more information, see the *SAP HANA Security Guide*.

You can also give a user the authority to grant the same role to another user. Enabling this option can help streamline development by bypassing the need to contact an admin user to grant the permission.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Users** tab, then select the user you want to grant permissions to. If the user doesn't appear in the tab's list:
 - a. At the bottom of the tab, select .
 - b. In the Add User view, enter a new username, then click **Next**.


Note

The process of adding a username to the permissions tile and assigning it authorization policies doesn't create a user. You can assign permissions and roles to any username, even if the user doesn't exist in the SAP HANA system. Either before or after you grant permissions or roles, create a corresponding SAP HANA user so that they can take effect. See the *SAP HANA Security Guide*.

3. In the User Details view, select **Edit**, then .
4. Use the dropdowns to select a privilege, privilege type, and resource type, and optionally, provide a resource.

If necessary, you can abandon a new permission entry as you are creating it by selecting .

If you don't enter a resource name, then the privilege applies to all instances of the resource type. For example, you might enter the name of a workspace in the resource field to grant access only to that workspace, or leave it blank to grant access to all workspaces on the server.

5. (Optional) Enable the **Grant Option**, which allows the user to grant this permission to other users and roles.
6. Click  to save the permission, and grant it to the user.


13.10.2 Revoke a Permission from a User

Revoke permissions that are no longer needed by the user.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. In the **Users** tab, select the user you want to revoke the permission from.
3. Select **Edit**.
4. Click  on the **Permission** row you want to revoke.
5. Select **Done** to save and exit Edit mode.


13.10.3 Create a Role

Add a blank role to the streaming roles list. Afterward, you can grant the role permissions.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Roles** tab in the banner.
3. Select .
4. Enter the name of the role you would like to add, then click **Save**.

Next Steps

After you create a role, you can grant that role permissions. See [Grant a Permission to a Role \[page 126\]](#).

13.10.4 Remove a Role

Remove a collection of permissions that are no longer required.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Roles** tab in the banner.
3. From the list, select the role you want to remove.
4. Select **Delete** at the bottom of the screen, then confirm your selection.

13.10.5 Grant a Permission to a Role

Define a set of privileges required to perform a particular task. You can then grant the role or multiple roles to any streaming user.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`


Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Roles** tab, then select a role from the list.
If you can't find the role in the list, you need to create it. See [Create a Role \[page 125\]](#).

3. In the Role Details view, select **Edit**, then .
4. Use the dropdowns to select a privilege, privilege type, and resource type, and optionally, provide a resource.

If necessary, you can abandon a new permission entry as you are creating it by selecting .

If you don't enter a resource name, then the privilege applies to all instances of the resource type. For example, you might enter the name of a workspace in the resource field to grant access only to that workspace, or leave it blank to grant access to all workspaces on the server.

5. Select  to save the permission for the role.


13.10.6 Revoke a Permission from a Role

Revoke permissions from a role whenever the requirements for the related task change.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select **Edit**, then click  on the **Permission** row you want to revoke.
3. Select **Done** to save and exit Edit mode.

13.10.7 Grant a Role to a User

A role is a set of permissions. You can grant a role to user instead of individual permissions, since a role normally contains all the permissions necessary to perform a certain task.


Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Context



A user can have more than one role. You can also give a user the authority to grant the same role to another user. Enabling this option can help streamline development by bypassing the need to contact an admin user to grant the permission.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Users** tab, then select the user you want to grant the role to. If the user doesn't appear in the tab's list:
 - a. At the bottom of the tab, select .
 - b. In the Add User view, enter a new username, then click **Next**.

i Note

The process of adding a username to the permissions tile and assigning it authorization policies doesn't create a user. You can assign permissions and roles to any username, even if the user doesn't exist in the SAP HANA system. Either before or after you grant permissions or roles, create a corresponding SAP HANA user so that they can take effect. See the *SAP HANA Security Guide*.

3. In the User Details view, select the **Roles** tab.
4. Select **Edit** to enter Edit mode, then click .
5. In the new row, use the dropdown to select a role.
6. (Optional) Enable the **Grant Option**, which allows the user to grant this role to other users.
7. Select  to grant the role.

8. Click **Done** to exit Edit mode.


13.10.8 Revoke a Role from a User

Revoke a role from a user when the user no longer needs to perform a particular task.

Prerequisites

- You are logged in to SAP HANA cockpit.
- Your tenant database user has the CATALOG READ system privilege.
- You have the following permission: `admin system system policy`

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Permissions**.
2. Select the **Users** tab in the banner.
3. Select the user you want to revoke a role from.
4. Click  for the role you want to revoke, then confirm your selection.
5. Select **Done** to save and exit edit mode.

13.11 Managing Streaming Data Services

The streaming analytics server accesses external databases through data service definitions, which are managed and stored in separate databases for each cluster. Use the Streaming Data Services page to manage data services.

i Note

If any of the functionality described throughout this section is not visible in your SAP HANA cockpit environment, you may have an older version of SAP HANA cockpit installed. To access all of the latest features, upgrade to the newest SP or revision. See the *SAP HANA Cockpit Installation and Update Guide* for more information.

Create a data service before running any projects that connect to external databases. You can create a data service either for a specific workspace, or make it available to all workspaces on the server.

Adapters that require database access obtain connections from the database manager, which stores data service connection information. When you are setting up a database adapter, specify the relevant data service

name. For example, you can define a service for connecting to an SAP ASE database through JDBC, or to SQL Server through ODBC. At runtime, the adapter obtains a connection from the database manager based on the properties in the data service, and executes queries over it.

For the server to communicate with external databases, you must have a working JDBC, ODBC, or Open Client connection. A full installation of streaming analytics includes a preconfigured ODBC driver manager. Otherwise, the appropriate JDBC or OCS driver for the desired external database must be installed manually.

The SAP HANA ODBC client must be installed to connect to an SAP HANA database. Use version 1.0.73 or later of the SAP HANA ODBC client.

Here is a list of database adapters and their supported drivers and databases.

Adapter	Supported Drivers	Supported Databases
SAP Adaptive Server Enterprise (ASE) Output adapter	Open Client™	SAP ASE
Database Input and Output adapters	JDBC	<ul style="list-style-type: none"> • SAP ASE • IBM DB2 • Oracle • Kx Systems KDB+ • Microsoft SQL Server • SAP HANA®
Database Input and Output adapters	ODBC	<ul style="list-style-type: none"> • SAP ASE • SAP IQ • SQL Anywhere® • IBM DB2 • Oracle • Microsoft SQL Server • TimesTen • MySQL 5.x • PostgreSQL • SAP HANA
SAP HANA adapter	ODBC	SAP HANA
SAP IQ adapter	ODBC	SAP IQ

In this section:

[Adding a JDBC Connection to an External Database \[page 131\]](#)

Create a service definition for a JDBC connection to the database of your choice.

[Adding an ODBC Connection to an External Database \[page 132\]](#)

Create a service definition for an ODBC connection to the database of your choice.

[Adding an OCS Connection to an External Database \[page 133\]](#)

Create a service definition for an Open Client (OCS) connection to the SAP Adaptive Server Enterprise (ASE) database. OCS connections are supported only through the SAP ASE Output adapter.

[Adding a Connection to an R Data Service \[page 134\]](#)

Create a service definition for a connection to an R data service.

[Adding a Connection to an SAP HANA Database \[page 135\]](#)

Create a service definition for an ODBC connection to an SAP HANA database.

[Updating or Deleting a Data Service \[page 136\]](#)

Modify or remove an existing database service definition.

13.11.1 Adding a JDBC Connection to an External Database

Create a service definition for a JDBC connection to the database of your choice.

Prerequisites

- Obtain the third-party JDBC `.jar` files from a given vendor, and copy them to your `$(STREAMING_HOME)/libj` directory. Save the driver to `$(STREAMING_HOME)/libj` before using it.
- Set up a connection to the streaming analytics server.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
2. In the data services view, click **Add +**, and select **Add JDBC Service**.
3. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.

This service name is the value you specify to components, such as the database adapter, that access external databases.

4. Select the scope of the data service.
 - To create a data service that is available to all workspaces on the server, choose **Server-wide**.
 - To create a data service that is only available to one workspace, choose a specific workspace.
5. Select the **Driver** from the dropdown list.

i Note

If you select **Other JDBC Database Connection**, the **Host**, **Port**, and **Database/Instance** are not required.

6. Set the **User** and **Password** parameters to the credentials necessary for communicating with the external database.
7. (Optional) Add a description of the database service in the **Description** parameter.

8. Define your connection using one of two methods:
 - a. Enter values for the **Host** and **Port** parameters.
 - b. Enter values for the **Connection String** parameter.

i Note

If you provide a connection string, streaming analytics ignores the **Host** and **Port** values.

9. (Optional) Enter a value in the **Connection Pool Size** field.
10. (Optional) Provide the name of the **Database** or **Instance**.
11. Click **Add**.

13.11.2 Adding an ODBC Connection to an External Database

Create a service definition for an ODBC connection to the database of your choice.

Prerequisites

i Note

Do not use this service for connecting to an SAP HANA database. See *Adding a Connection to an SAP HANA Database* in this guide for instructions on creating an SAP HANA service.

- Install an ODBC driver.
- Configure a DSN for the ODBC connection in the `$DIR_INSTANCE/streaming/cluster/<lowercase-db-name>/config/odbc.ini` file.
- Set up a connection to the streaming analytics server.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
2. In the data services view, click **Add +**, and select **Add ODBC Service**.
3. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.This service name is the value you specify to components, such as the database adapter, that access external databases.
4. Select the scope of the data service.

- To create a data service that is available to all workspaces on the server, choose **Server-wide**.
 - To create a data service that is only available to one workspace, choose a specific workspace.
5. Set the **User** and **Password** parameters to the credentials necessary for communicating with the external database.
 6. (Optional) Add a description of the database service in the **Description** parameter.
 7. Set the **ODBC DSN** parameter to the data source name to be used by your service. You should already have this data source set up with the ODBC driver manager.
 8. Leave the **Connection to HANA Database** checkbox blank.
If you need to connect to an SAP HANA database, use the SAP HANA service instead.
 9. Click **Add**.

13.11.3 Adding an OCS Connection to an External Database

Create a service definition for an Open Client (OCS) connection to the SAP Adaptive Server Enterprise (ASE) database. OCS connections are supported only through the SAP ASE Output adapter.

Prerequisites

- Set up a connection to the streaming analytics server.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
2. In the data services view, click **Add +**, and select **Add OCS Service**.
3. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.
 This service name is the value you specify to components, such as the database adapter, that access external databases.
4. Select the scope of the data service.
 - To create a data service that is available to all workspaces on the server, choose **Server-wide**.
 - To create a data service that is only available to one workspace, choose a specific workspace.
5. Set the **User** and **Password** parameters to the credentials necessary for communicating with the external database.
6. (Optional) Add a description of the database service in the **Description** parameter.

7. Enter details for the **Host** and **Port**.
8. (Optional) Set the **TDS Packet Size** parameter for optimal performance. If not set, streaming analytics uses the default value for Open Client.
See the **CS_PACKETSIZE** connection property in the Open Client documentation for more information.
9. (Optional) Set the **Application Name** parameter to help identify Open Client database connections used by the ASE Output adapter.
See the **CS_APPNAME** connection property in the Open Client documentation for more information.
10. Click **Add**.

13.11.4 Adding a Connection to an R Data Service

Create a service definition for a connection to an R data service.

Prerequisites

- Rserve authentication is enabled.
- Set up a connection to the streaming analytics server.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
2. In the data services view, click **Add +**, and select **Add R Service**.
3. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.This service name is the value you specify to components that access Rserve.
4. Select the scope of the data service.
 - To create a data service that is available to all workspaces on the server, choose **Server-wide**.
 - To create a data service that is only available to one workspace, choose a specific workspace.
5. Set the **User** and **Password** parameters to the credentials necessary for communicating with Rserve.
6. (Optional) Add a description for the data service in the **Description** parameter.
7. Enter details for the **Host** and **Port**.
8. Click **Add**.

13.11.5 Adding a Connection to an SAP HANA Database

Create a service definition for an ODBC connection to an SAP HANA database.

Context

i Note

The streaming analytics server that you connect to must be version 1.0 SP 10 or newer. If you connect to an older version, the option to create a HANA service does not appear in the **Data Services** page in HANA cockpit.

Procedure

1. On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
2. In the data services view, click **Add +**, and select **Add HANA Service**.
3. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.

This service name is the value you specify to components, such as the database adapter, that access external databases.
4. Select the scope of the data service.
 - To create a data service that is available to all workspaces on the server, choose **Server-wide**.
 - To create a data service that is only available to one workspace, choose a specific workspace.
5. Set the **User** and **Password** parameters to the credentials necessary for communicating with the external database.
6. (Optional) Add a description of the database service in the **Description** parameter.
7. (Optional) Select **Use Default HANA Server** if you are connecting to the local SAP HANA server.

If you choose to connect to the local SAP HANA server, you do not have to provide an SAP HANA hostname or instance number.
8. (Required if connecting to a remote SAP HANA database) Specify a database type:
 - Select **Single Tenant** if the SAP HANA system is the only database at the location, or;
 - Select **Multi Tenant** if you are connecting to a system with multiple databases. If you select this option, specify the name of the database that you are connecting to in the **Database Name** field.
9. (Required if connecting to a remote SAP HANA database) Set the **HANA Hostname** parameter to the fully qualified domain name (FQDN) of the host on which the SAP HANA database is installed.
10. (Required if connecting to a remote SAP HANA database) Set the **HANA Instance Number** parameter to the instance number of the target database.

11. (Optional) Enable **Multi-Byte Character Support** if the database you are connecting to uses multi-byte characters.
12. Click **Add**.



13.11.6 Updating or Deleting a Data Service

Modify or remove an existing database service definition.

Prerequisites

Set up a connection to the streaming analytics server.

Procedure

- On the SAP HANA system overview, under **Streaming Analytics**, click **Manage Streaming Data Services**.
- To update a data service:
 - a. In the data services view, select a data service.
 - b. Revise any fields as appropriate.
 - c. Click **Save**.
- To rename a data service:
 - a. In the data services view, select a data service.
 - b. Click **Rename**.
 - c. Give the data service a name that:
 - begins with a letter,
 - is unique,
 - is case-sensitive,
 - contains a character string consisting of any letters, numbers, underscores, dots, and colons.This service name is the value you specify to components, such as the database adapter, that access external databases.
 - d. Click **OK**.
- To delete a data service:
 - a. In the data services view, select a data service.
 - b. Click **Delete**.
- To sort the data services in ascending or descending order, click **Sort** , and select the appropriate option.
- To filter the data services according to their scope, click **Filter** , and select the appropriate option.

14 Administer a Cluster

Attending to such items as logging, data backup, and data restoration ensures that your cluster is administered properly.

In this section:

[Safeguarding Your Data \[page 137\]](#)

Protect your data to improve system redundancy and prevent unauthorized access.

[Memory Usage \[page 141\]](#)

Memory usage is directly proportional to the number of records stored in a project.

[Log Files \[page 142\]](#)

SAP HANA streaming analytics produces log files for projects and the cluster. A cluster node may contain multiple projects, each with its own project log file.

[Cluster Administrative Tool \[page 148\]](#)

The cluster administrative tool is one of several options available for cluster administration. Use it to add and remove projects and workspaces, and to query, start, and stop existing projects.

14.1 Safeguarding Your Data

Protect your data to improve system redundancy and prevent unauthorized access.

Procedure

- Secure data using OS security.
- Take steps to secure files. We recommend using disk volume encryption and storing security-related configuration on a separate disk.
- Use third-party source control to manage your project source files and provide redundancy. When source files are checked out of the source control system, use studio to browse your source folder and change source files.
- Perform regular backups of project data, including log stores.
- Configure and monitor the log file size. Having the size set too low will result in files automatically overwriting themselves, resulting in a loss of data; having it set too high will eventually flood a server's hard drive with logging data. For more information on log file size, see the *SAP HANA Streaming Analytics: Developer Guide*.

Next Steps

You can find more information on security best practices and procedures in the *SAP HANA Streaming Analytics: Security Guide*.

In this section:

[Backing Up SAP HANA with SAP HANA Streaming Analytics \[page 138\]](#)

Back up your streaming analytics system.

[Recovering SAP HANA with SAP HANA Streaming Analytics \[page 140\]](#)

Recover streaming analytics after a failure.

Related Information

[SAP HANA Streaming Analytics: Security Guide](#)

[SAP HANA Streaming Analytics: Developer Guide](#)

14.1.1 Backing Up SAP HANA with SAP HANA Streaming Analytics

Back up your streaming analytics system.

Prerequisites

Before performing a backup, set the node name for all streaming analytics hosts. See the [Edit the Streaming Node Name \[page 81\]](#) topic for more information.

Back up HANA with SAP HANA Streaming Analytics Configurations

Back up your SAP HANA system and tenant database that streaming analytics is provisioned to, including streaming analytics:

1. Perform a full backup of the SAP HANA system and tenant database in Cockpit. See the [Creating Backups](#) section for more information.
2. Back up the following files and directories using the Linux tar utility:
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/keys`
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/adapters/framework/adapter.key`
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/config/cluster.cfg`

- \$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/config/cluster_log4j2.xml

Back up SAP HANA Streaming Analytics Projects and Log Stores

Back up your streaming analytics projects' log store files either online or offline. An offline backup ensures that the log stores are up-to-date at the time of your backup, and can be done with the tar utility for easy restoration. Online backup log store files require that the backup extension, .bak, which is applied during the backup process, be removed upon restoration. Before shutting down streaming analytics, verify the locations of the project files and the type of store defined for each stream.

i Note

You do not need to stop the project server during an online backup, but operation suspends while the backup files are being created, which may cause a short disruption. The length of this suspension depends on the amount of data accumulated in the log stores. Perform an online backup only when short disruptions are acceptable.

1. Stop any running projects that you'd prefer to backup offline.
2. Perform an online backup of running projects. Use the `streamingprojectclient` utility to execute the following command for each running project:

```
$STREAMING_HOME/bin/streamingprojectclient -e -p <host>:<port>/<workspace-name>/<project-name> -c <user>[:<password>] backup
```

Where `<host>:<port>` are the host name and port number of your cluster cache. For `<user>` and `<password>`, enter your user credentials.

This creates a set of backup files in the log store directories, each with the extension .bak. Only the current contents of the stores are copied over.

3. Create a backup archive that includes the `/hana/data_streaming/<SID>` directory. This captures your offline project log stores, as well as the CCR and CCX files of all deployed projects.
4. Back up any external files used by your projects.

i Note

The server does not store project CCL files, so you can only back them up from the machine they are stored on. These are required if you want to continue developing a project after a failure, but are not required to continue running the project as it was.

14.1.2 Recovering SAP HANA with SAP HANA Streaming Analytics

Recover streaming analytics after a failure.

Prerequisites

Ensure that the following components of your SAP HANA system are consistent with when you performed the backup:

- Total number of hosts
- Number of hosts with only the streaming role
- Number of hosts with multiple roles including streaming
- Node names

Recover SAP HANA with SAP HANA Streaming Analytics Configurations

1. Replace the following files and directories with your backups using the Linux tar utility:
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/keys`
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/adapters/framework/adapter.key`
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/config/cluster.cfg`
 - `$DIR_INSTANCE/streaming/cluster/<tenant-db-name>/config/cluster_log4j2.xml`
2. If you are restoring to a new environment, edit the `$DIR_INSTANCE/streaming/cluster/<db-name>/config/cluster.cfg` file to update the following properties:
 - `jdbc-url`
 - `cluster-name`
 - `cipher file`
3. Recover the SAP HANA system and tenant database that streaming analytics is provisioned to using your backup. See the [SAP HANA Recovery](#) section for more information.
4. (Optional) Update the activation plugin.

i Note

The activation plugin is only necessary if you are activating streaming analytics content in the streaming analytics repository in SAP HANA.

1. Log in as the `<sid>adm` user.
2. Stop the activation job (ignore any errors):

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<tenant-db-name>/config/cluster.cfg --deactivate-job
```

3. Delete the trust store (ignore any errors):

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<tenant-db-name>/config/cluster.cfg --delete-truststore
```

4. Create a new trust store, import certificates, and start the activation job:

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<tenant-db-name>/config/cluster.cfg --install
```

5. For each streaming analytics host, create new HTTP destination files:

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/  
<tenant-db-name>/config/cluster.cfg --add-host --hdb-host <streaming-  
hostname>
```

Recover SAP HANA Streaming Analytics Projects and Log Stores

Recover your projects and log store files using tar:

1. Stop the cluster or project that you are restoring.
2. Extract the `/hana/data_streaming/<SID>` directory from your backup .tar files. This includes CCR and CCX files for all deployed projects. If you did an offline log store backup, this also includes log store files. If you did an online log store backup, you will also need to restore your log store .bak files and rename them to .log files.
3. Restore any external files used by your projects.

Before You Start the Streaming Web Service or Web Services Provider

1. Log onto the HANA cockpit.
2. On the SAP HANA system overview, scroll down to the **Streaming Analytics** tile and click **Configure Streaming Cluster**.
3. Select the **Streaming Web Service** tab.
4. Verify that the path for the SSL Key Files Location property is correct.
5. Select the **Web Services Provider** tab.
6. Verify that the paths for the Cipher File and Keystore File properties are correct.

14.2 Memory Usage

Memory usage is directly proportional to the number of records stored in a project.

There are no configuration settings in the project server that directly set up or control RAM usage on the machine. However, the project server does count records in the system, to ensure that only one copy of a record exists in different streams.

Each streaming analytics project launches a Java virtual machine (JVM), which runs any Java UDFs or Java internal adapters associated with the project. Memory available to the JVM is controlled by the `java-max-heap` option in the Deployment section of the project configuration (CCR) file; the default value is 256 MB.

If your project triggers Java out-of-memory errors, increase the heap size for the project's JVM. For example:

```
<Option name="java-max-heap" value="512"/>
```

Use the `memory-reserve` option to specify the amount of reserved memory, in megabytes, that you want released to shut a project for streaming analytics down gracefully in the case that all available memory is used. The default value for this option is 0. The recommended value is 10.

14.3 Log Files

SAP HANA streaming analytics produces log files for projects and the cluster. A cluster node may contain multiple projects, each with its own project log file.

Streaming analytics stores logs in flat files. The project and cluster log files reside in the SAP HANA trace directory, `$SAP_RETRIEVAL_PATH/trace/DB_<database-name>/`. The cluster log files are named `streamingserver_<host>.3<instance-number><internal-port>.<log-serial-number>.log`.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='streamingserver' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view. The external port number is always one number higher than the internal port. For example, if the internal port is 3xx40 then the external port is 3xx41.

Use the **Diagnosis Files** tab of the SAP HANA Administration editor to perform analysis on the logs. See the *SAP HANA Administration Guide* for additional details about the diagnosis files and the Administration editor.

In this section:

[Installation Log Files \[page 143\]](#)

Streaming analytics logs installation results, errors, and warnings from various components in different log files. Review these logs to help troubleshoot issues. If you require technical support, your representative may request that you send information from one or more of these logs.

[Cluster Log Configuration File \[page 144\]](#)

The configuration file for cluster logging is `cluster_log4j2.xml`, which is a `log4j` configuration file.

[Project Log Files \[page 145\]](#)

Configure project logs to capture errors in running projects. You can configure logs for single or multiple projects in a cluster.

[Logging Level \[page 146\]](#)

Logging levels range from 0 to 7, and represent a decreasing order of severity. For example, the higher the log error, the less severe the issue. The higher you set the log level, the more information you receive as errors up to and including that log level are reported. The default logging level for projects is 4.

[Locating Log and Trace Files for Troubleshooting Adapters \[page 147\]](#)

Locate log and trace files that contain valuable information for troubleshooting various adapter problems.

Related Information

[SAP HANA Administration Guide](#)
[SAP HANA Streaming Analytics: Developer Guide](#)

14.3.1 Installation Log Files

Streaming analytics logs installation results, errors, and warnings from various components in different log files. Review these logs to help troubleshoot issues. If you require technical support, your representative may request that you send information from one or more of these logs.

The `streaming_suite.log` file contains a summary of the streaming analytics installation results.

The `streaming_suite.log` file and all other files are located at (Linux) `<install-dir>/log`.

Filename	Component
<code>streaming_suite.log</code>	Streaming analytics; includes a summary of installation results
<code>conn_lang.log</code>	Open Client connectivity language modules
<code>dbilib.log</code>	Open Client DB-Library
<code>esp_framework_install.log</code>	Adapter Toolkit
<code>esp_http_install.log</code>	HTTP Output Adapter
<code>esp_logfile_input_install.log</code>	Logfile Input Adapter
<code>esp_odbc_install.log</code>	ODBC Driver
<code>esp_pde_install.log</code>	PowerDesigner Extensions for streaming analytics
<code>esp_repserver_install.log</code>	Replication Server Adapter
<code>esp_rfc_install.log</code>	SAP RFC Adapter
<code>esp_server_install.log</code>	Streaming analytics server
<code>esp_studio_install.log</code>	Streaming analytics plugin for SAP HANA studio
<code>esp_tibco_rv_install.log</code>	TIBCO Rendezvous Adapter
<code>esp_ws_install.log</code>	Web Services Adapter
<code>espcmap.log</code>	Management User Interface
<code>jre7.log</code>	SAP Java Runtime Environment
<code>lang.log</code>	Open Client language modules

Filename	Component
open_client.log	Open Client (OCS)
sysam_util.log	SySAM License Utilities

14.3.2 Cluster Log Configuration File

The configuration file for cluster logging is `cluster_log4j2.xml`, which is a `log4j` configuration file.

The location for this file is `STREAMING_HOME/./cluster/<SID>/config/cluster_log4j2.xml`, where `<SID>` is the system ID of your SAP HANA system.

Sample `cluster_log4j2.xml` file:

```
<?xml version="1.0"?>
<Configuration name="Log4j1">
  <Properties>
    <Property name=".level">INFO</Property>
    <Property name="com.sybase.esp.cluster.impl.Log4JHandler.level">FINEST</Property>
  </Properties>
  <Appenders>
    <RollingFile name="Log" fileName="${sys:com.sybase.esp.cluster.logfile}.000.log" filePattern="${sys:com.sybase.esp.cluster.logfile}.*%03i.log">
      <PatternLayout pattern="%d{MMM dd yyyy HH:mm:ss.SSS} %p - %m%n"/>
      <Policies>
        <SizeBasedTriggeringPolicy size="1MB"/>
      </Policies>
      <DefaultRolloverStrategy max="5"/>
    </RollingFile>
  </Appenders>
  <Loggers>
    <Root level="info">
      <AppenderRef ref="Log"/>
    </Root>
    <Logger name="com.sybase.esp.cluster.applications" level="info"/>
    <Logger name="com.sybase.esp" level="info"/>
  </Loggers>
</Configuration>
```

The `com.sybase.esp.cluster.logfile` property is ignored as the cluster log file is named `streamingserver_<machine hostname>.<5 digit cluster node's rpc port>.<3 digit log serial number>.log` and is located in the trace directory. This location cannot be changed.

The cluster log file is configured by default to back up its contents once the file reaches 1MB in size. The `DefaultRolloverStrategy` option specifies how many backup files to create.

You can set the root and logger levels to `error` or `info`. The `info` option produces minimum log information. Under normal circumstances, keep the root option set to the default value `info`, or the log becomes almost unreadable because of its size. You can use the `com.sybase.esp` logger to debug a node without using third-party debugging components. Do not modify the `com.sybase.esp.cluster.applications` property; the `info` value is required in this instance.

Consult `log4j 2` documentation for more information on supported properties and configuration instructions.

14.3.3 Project Log Files

Configure project logs to capture errors in running projects. You can configure logs for single or multiple projects in a cluster.

The streaming analytics project log and standard streams log files are both located in the SAP HANA trace directory. The standard streams file receives all output written to stdout and stderr, including SySAM licensing information for streaming analytics, as well as messages from third party applications that write to stdout and stderr.

Project log file	<code>streamingserver~<workspace-name>.<project-name>.<project-instance-number>~_<machine-hostname>.<cluster-node-rpc-port>.<log-serial-number>.trc</code>
Standard streams log file	<code>streamingserver~<workspace-name>.<project-name>.<project-instance-number>~_<machine-hostname>.<cluster-node-rpc-port>.<log-serial-number>.out</code>
Trace directory	<code>\$DIR_INSTANCE/<hostname>/trace</code>

The format of a message is: `[SP-<Severity>-<Message Code>] (<Time>) sp(<Process ID>) <Log Message>` where:

- Severity is a number which qualifies the importance of the issue. For example, 1 is a severe problem, 3 is an error, 4 is a warning, and 7 is an informational message.
- Message code helps identify the exact error message that is being printed. Since certain messages can be similar, the error code helps identify the exact issue that has occurred.
- Time is the timestamp in the server's local time.
- Process ID is a number assigned by the operating system to each running instance of an executable. Each instance gets a unique value.

Modify logging levels for projects in their project configuration files (`.ccr`), or using the Project Configuration editor in the SAP HANA Streaming Development perspective in studio. For more information about the `.ccr` file and the Project Configuration editor, see the *Project Configurations* section in the *SAP HANA Streaming Analytics: Developer Guide*.

To modify logging levels for a project at runtime, use `streamingprojectclient` to execute:

```
streamingprojectclient -p [<host>:]<port>/<workspace-name>/<project-name> -c <username>:<password> "loglevel <level>"
```

Log level changes made with `streamingprojectclient` do not persist— if you restart the project without also changing the logging level in the `<project-name>.ccr` file, you lose your changes to the logging level. After you change the logging level in `<project-name>.ccr`, stop and remove the project from the node, then redeploy the project to activate the new logging level.

The timezone for these log files is the local timezone.

You can modify the size and number of log files stored in a given project using the `logfile-size` and `logfile-depth` properties in the project configuration (`.ccr`) file. You can edit the `.ccr` file manually, or by using the Project Configuration editor in the SAP HANA Streaming Development perspective in studio. For

more information about the .ccr file and the Project Configuration editor, see the *Project Configurations* section in the *SAP HANA Streaming Analytics: Developer Guide*.

For information about Flex logging, see *Flex Logging* in the *SAP HANA Streaming Analytics: Developer Guide*.

Related Information

[Project Configurations](#)

[Flex Logging](#)

14.3.4 Logging Level

Logging levels range from 0 to 7, and represent a decreasing order of severity. For example, the higher the log error, the less severe the issue. The higher you set the log level, the more information you receive as errors up to and including that log level are reported. The default logging level for projects is 4.

You can set logging levels:

- In cluster configuration, logging levels apply to all projects that run on the node unless you set a different logging level in the .ccr file of a project.
- In the project configuration file, `<project-name>.ccr`.
- Using `streamingprojectclient` at runtime.

Name	Level	Description
LOG_EMERG	0	System is unusable
LOG_ALERT	1	Action must be taken immediately
LOG_CRIT	2	Critical conditions
LOG_ERR	3	Error conditions
LOG_WARNING	4	Warning conditions
LOG_NOTICE	5	Normal but significant condition
LOG_INFO	6	Informational
LOG_DEBUG	7	Debug-level messages

Adjust the log level for your streaming analytics project log file according to the type of information you want to receive in SAP HANA. The higher the trace level, the more detailed the information recorded by the trace (`streamingserver~<workspace-name>.<project-name>.<project-instance-number>~<machine-hostname>.<cluster-node-rpc-port>.<log-serial-number>.trc`).

SAP HANA Streaming Analytics Log Level	SAP HANA Trace Level
0	f (FATAL)

SAP HANA Streaming Analytics Log Level	SAP HANA Trace Level
3	e (ERROR)
5	w (WARNING)
6	i (INFO)
7	d (DEBUG)

14.3.5 Locating Log and Trace Files for Troubleshooting Adapters

Locate log and trace files that contain valuable information for troubleshooting various adapter problems.

Context

You can access log and trace files for adapters using the studio Administration Console or by navigating to `hana/shared/<SID>/HDB<instance>/<hostname>/trace/DB_<tenant-database-name>` on the streaming analytics node's file system.

Procedure

1. In the SAP HANA Administration Console, right-click on your system in the **Systems** pane and select **► Configuration and Monitoring ► Open Administration ►**.
2. Click the **Diagnosis Files** tab.
3. Enter **streamingserver** into the filter to find the `.out`, `.trc`, and `.log` files, and double-click to view.
 - The `.out` file (`streamingserver~<workspace-name>.<instance-index>~<hostname>.<port>.000.out`) reports project and adapter information and errors.
 - The `.trc` file (`streamingserver~<workspace-name>.<instance-index>~<hostname>.<port>.000.trc`) reports data flow, adapter parsing errors, and truncation warnings.
 - Each adapter `.log` file (`streamingserver~<workspace-name>.<project-or-adapter-name>.<instance-index>~<hostname>.<port>.000.log`) reports information specific to that adapter.
 - Each adapter bad data `.log` file (`streamingserver~<workspace-name>.<project-or-adapter-name>.<instance-index>.<adapter-name>_bad_data~<hostname>.<port>.log.<yyyy-mm-dd>`) reports data that can't be parsed by the adapter.

i Note

Bad data logging is not enabled by default, and is only done by the File/Hadoop CSV Input and File/Hadoop Event XML Input adapters. If necessary, enable bad data logging through the adapter's properties.

14.4 Cluster Administrative Tool

The cluster administrative tool is one of several options available for cluster administration. Use it to add and remove projects and workspaces, and to query, start, and stop existing projects.

You can perform the same tasks in streaming analytics plugin for SAP HANA studio, and in the SAP HANA cockpit.

The cluster administrative tool operates in interactive mode or command line mode. In interactive mode, connect to the cluster manager once, and execute commands until you exit. In command line mode, the utility logs you out after each command; you must enter the URI and authentication details (which vary by authentication type) to connect to the cluster manager every time you specify a command.

Interactive mode requires less typing. Command line mode is intended for scripting. To use interactive mode, the password you use to connect to the cluster manager must begin with an alphabetic character.

i Note

The parameters, excluding supported commands, are case-insensitive.

When you are logged on to the host as the `<sid>adm` user, the `streamingclusteradmin` utility is in the default path.

These interactive mode examples demonstrate the use of some of the parameters and commands:

```
streamingclusteradmin --uri=esps://cluster_server:3<instance-number><external-  
port> --username=me --password=sybase  
> get managers  
> get workspaces  
> get projects
```

For multiple tenant systems, the external port can be a different value than 3xx26. The external port is always the internal port number plus one. For example, in single tenant systems, the internal port is 3xx16 (where xx is the instance number) and the external port is 3xx26. In multiple tenant systems, if the internal port is 3xx40 then the external port is 3xx41, if 3xx49 then 3xx50, and so on.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='<streamingserver>' AND HOST='<host name>'` statement and look at the PORT column in the `M_SERVICES` view.

These command line mode examples demonstrate the use of some of the parameters and commands:

```
streamingclusteradmin --uri=esps://cluster_server:3<instance-number><external-  
port> --username=me --password=sybase --get_managers  
streamingclusteradmin --uris=esps://cluster_server:3<instance-number><external-  
port> --username=me --password=sybase --get_workspaces
```

```
streamingclusteradmin --uri=esps://cluster_server:3<instance-number><external-port> --username=me --password=sybase --get_projects
```

The following table describes the function of each `streamingclusteradmin` command:

Command	Function
Interactive mode: <code>get managers</code> Command line mode: <code>--get_managers</code>	Returns the host-name:rpc-port pairs for the managers in the cluster.
Interactive mode: <code>get controllers</code> Command line mode: <code>--get_controllers</code>	Returns the list of controllers in the cluster.
Interactive mode: <code>get workspaces</code> Command line mode: <code>--get_workspaces</code>	Returns the names of the workspaces in the cluster.
Interactive mode: <code>get projects</code> Command line mode: <code>--get_projects</code>	Returns the list of projects, with their state.
Interactive mode: <code>get project <workspace-name>/<project-name></code> Command line mode: <code>--get_projectdetail --workspace-name=<workspace-name> --project-name=<project-name></code>	Returns information about the specified project, including whether it is running, on which node it is running, and runtime details. For a project running in hot standby mode, the command returns information for each instance, and identifies the primary and secondary instances.
Interactive mode: <code>get streams <workspace-name>/<project-name></code> Command line mode: <code>--get_streams --workspace-name=<workspace-name> --project-name=<project-name></code>	Returns the streams associated with a workspace.
Interactive mode: <code>get schema <workspace-name>/<project-name> <stream-name></code> Command line mode: <code>--get_schema --workspace-name=<workspace-name> --project-name=<project-name> --stream-name=<stream-name></code>	Returns the schema of the specified stream.
Interactive mode: <code>add workspace <workspace-name> [--ignore-error]</code> Command line mode: <code>--add_workspace --workspace-name=<workspace-name> [--ignore-error]</code>	Adds a workspace. Use the optional <code>ignore-error</code> argument to add the workspace even when doing so causes a workspace error.

Command	Function
<p>Interactive mode: <code>add project <workspace-name>/<project-name> <project-name>.ccx [<project-name>.ccr]</code></p> <p>Command line mode: <code>--add_project --workspace-name=<workspace-name> --project-name=<project-name> --ccx=<project-name>.ccx [--ccr=<project-name>.ccr]</code></p>	<p>Adds a project.</p> <p><code><project-name>.ccx</code> is the compiled project file. Specify the path to the file.</p> <p><code><project-name>.ccr</code> is the project's runtime configuration file. Include the CCR file for project with affinities, or one running in hot standby mode. Specify the path to the file. <code><project-name>.ccr</code> and <code><project-name>.ccx</code> are always located in the same directory.</p>
<p>Interactive mode: <code>remove workspace <workspace-name> [--ignore-error]</code></p> <p>Command line mode: <code>--remove_workspace --workspace-name=<workspace-name> [--ignore-error]</code></p>	<p>Removes a workspace. Use the optional <code>ignore-error</code> argument to remove the workspace even when doing so causes a workspace error.</p>
<p>Interactive mode: <code>remove project <workspace-name>/<project-name></code></p> <p>Command line mode: <code>--remove_project --workspace-name=<workspace-name> --project-name=<project-name></code></p>	<p>Removes a project.</p> <p>Prerequisite: Stop the project. You cannot remove a running project.</p>
<p>Interactive mode: <code>start project <workspace-name>/<project-name> [timeout (sec)] [<instance-index>]</code></p> <p>Command line mode: <code>--start_project --workspace-name=<workspace-name> --project-name=<project-name> [--timeout=<timeout-in-seconds>] [--instance-index=<instance-index>]</code></p>	<p>Starts the project. If the project is added with a strong controller affinity and that controller is not available, start-up fails.</p> <p><code><timeout-in-seconds></code> specifies how long the call waits to verify that the project has started.</p> <p>For a project running in hot standby mode, the instance index specifies which of the two instances to start. Valid values are 0 and 1. Use <code>get_projector --get_project_detail</code> to determine whether and where the instances are running.</p>
<p>Interactive mode: <code>stop project <workspace-name>/<project-name> [timeout (sec)] [<instance-index>]</code></p> <p>Command line mode: <code>--stop_project --workspace-name=<workspace-name> --project-name=<project-name> [--timeout=<timeout-in-seconds>] [--instance-index=<instance-index>]</code></p>	<p>Stops the project.</p> <p><code><timeout-in-seconds></code> specifies how long the call waits to verify that the project has stopped.</p> <p>For a project running in hot standby mode, the instance index specifies which of the two instances to stop. Valid values are 0 and 1. Use <code>get_projector --get_project_detail</code> to determine whether and where the instances are running.</p>

Command	Function
Interactive mode: <code>encrypt <clear-text></code> Command line mode: <code>--encrypt_text --text=<clear-text></code>	Encrypts plain text data. Use this command to encrypt passwords in configuration files.
Interactive mode: <code>deploykey <new-username> <keystore> <store-pass> <key-alias> [<store-type>]</code> Command line mode: <code>--deploy_key --new-user=<new-username> --keystore=<keystore> --storepass=<store-pass> --key-alias=<key-alias> [--storetype=<store-type>]</code>	Adds a new user by deploying a new user key to the key store. When you deploy a new user key, the node to which you send the deploy command updates the key store, and the other nodes then reload that file. To test if the deploy key is working properly, log in to the cluster with the new key, but through a different node.
Interactive mode: <code>connect</code>	Connect or reconnect a project to a cluster. This command is in interactive mode only.
Interactive mode: <code>quit</code> or <code>exit</code>	Logs you out of interactive mode. To reaccess the utility, provide your user name and password.
Interactive mode: <code>help</code> Command line mode: <code>--help</code>	Displays a plain-text description of the <code>streamingclusteradmin</code> utility's commands and usage information.

15 Copy or Move a Tenant Database

You can copy or move a tenant database provisioned with the streaming analytics service to another SAP HANA system.

Context

After a copy, the new tenant database runs as a separate, independent database. After a move, the original database is deleted and the new database takes over. The steps for doing this are the same as outlined in the *SAP HANA Administration Guide* with a few additional steps specific to streaming analytics.

Procedure

1. Complete the tasks outlined in the *Preparing to Copy or Move a Tenant Database* section in the *SAP HANA Administration Guide*.
2. (Streaming analytics-specific) If you did not specify names for the streaming analytics nodes before you initialized the connection between the tenant database and the streaming analytics server:
 - Uninitialize the streaming analytics service (see [Unprovision Streaming Analytics from a Tenant Database \[page 156\]](#)).
 - Name the nodes on all the streaming analytics hosts currently being used by the source tenant database and that will be used by the target tenant database (see [Edit the Streaming Node Name \[page 81\]](#)).
 - Reinitialize the connection (see [Initializing the Connection Between a Tenant Database and a Streaming Service \[page 158\]](#)).
3. Create a tenant database in the target system as a copy of the original tenant database in the source system. See *Copy a Tenant Database to Another System* or *Move a Tenant Database to Another System* in the *SAP HANA Administration Guide* for complete instructions.
4. Enable HTTP(S) access for the target tenant database by executing this SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<target tenant
database name>')
  SET ('public_urls', 'https_url') = 'https://<virtual hosts name>:
43<inst#>' WITH RECONFIGURE;
```

See *Configure HTTP(S) Access to Tenant Databases via SAP HANA XS Classic* in the *SAP HANA Administration Guide* for additional details.

5. Finalize the copy or move. See *Copy a Tenant Database to Another System* or *Move a Tenant Database to Another System* in the *SAP HANA Administration Guide* for details on the statement you need to execute.

6. (Optional, Streaming analytics-specific) If you want to be able to activate streaming analytics projects in the SAP HANA repository on the target tenant database, execute the following SQL statements:

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'SYSTEM') SET
('repository', 'jsplugin_sap/hana/streaming/repoplugin/plugin.xsjs') = 'true'
WITH RECONFIGURE;
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'SYSTEM') SET
('scheduler', 'enabled') = 'true' WITH RECONFIGURE;
```

Related Information

[SAP HANA Administration Guide](#)

16 Provisioning the Streaming Analytics Service to a Tenant Database

Provisioning creates and assigns the streaming analytics service (`hdbstreamingserver`) to the tenant database. Depending on your SAP HANA instance configuration, this task may not be required.

Context

If your SAP HANA instance has only ever contained the original initial tenant, when you added the streaming analytics host or role, the streaming analytics service was automatically provisioned to the tenant. But if your single tenant instance ever previously contained multiple tenant databases, or your instance currently contains multiple tenants, you must manually provision the streaming analytics service to the tenant.

A tenant can run a streaming cluster consisting of multiple streaming servers that each run on a separate host. Multiple tenants can share a single streaming host for their streaming servers with each tenant running only one service on that host.

Procedure

1. In a SQL console, connect to the SYSTEMDB as a user with the DATABASE ADMIN privilege.
2. Execute the following ALTER DATABASE statement:

```
ALTER DATABASE <database_name> ADD 'streamingserver'  
    [ AT [ LOCATION<streaming_hostname>[:3<instance_number><internal-  
port>' ] ] ] '
```

The host and port information you specify is that of the streaming analytics host. If there are multiple streaming analytics hosts available and you do not specify a host or port, the SAP HANA system randomly selects from the available hosts.

The streaming analytics service runs on one of seven reserved internal ports: 3xx16 (default), 3xx40, 3xx49, 3xx58, 3xx67, 3xx76, or 3xx85, where xx is the instance number. If not specified, the default, or first available port is used. If you would like to have additional tenants on the same host, you can reserve additional ports by editing the `[multidb] reserved_instance_numbers` property in the `global.ini` file.

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='streamingserver' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view. The external port number is always one number higher than the internal port. For example, if the internal port is 3xx40 then the external port is 3xx41.

Example

The following code example adds the `streamingserver` service to the tenant `my_tenant` on `hostA`, for internal port `30316`.

```
ALTER DATABASE my_tenant ADD 'streamingserver' AT LOCATION 'hostA:30316'
```

Next Steps

- If you plan to use system replication, configuration should be done before initialization. See [\(Optional\) Configure System Replication](#).
- If not using system replication, your next task is to initialize the connection. See [Initializing the Connection Between a Tenant Database and a Streaming Service](#).

Related Information

[SAP HANA Security Guide](#)

17 Unprovision Streaming Analytics from a Tenant Database

Unassign the `hdbstreamingserver` service from a tenant database when it no longer requires streaming analytics.

Prerequisites

The connection between the tenant database and the streaming server is uninitialized.

Context

Unassigning the `hdbstreamingserver` service from a tenant database removes any streaming analytics data from the tenant, such as streaming users, and tables, and makes the streaming analytics host available for provisioning to another tenant.

Procedure

1. In a SQL console, connect to the SYSTEMDB as a user with the DATABASE ADMIN privilege.
2. Execute the following `ALTER DATABASE` statement:

```
ALTER DATABASE <database_name> REMOVE 'streamingserver'  
    AT [ LOCATION ] '<streaming_hostname>:3<instance_number><available  
internal port>'
```

To see your assigned internal port number, look at the `streamingserver` service in the landscape view in SAP HANA studio or run the `SELECT PORT FROM M_SERVICES WHERE SERVICE_NAME='streamingserver' AND HOST='<host name>'` statement and look at the `PORT` column in the `M_SERVICES` view. The external port number is always one number higher than the internal port. For example, if the internal port is 3xx40 then the external port is 3xx41.

Example

The following code example removes the service `streamingserver` from the tenant `my_tenant` on `hostA`, for instance 03.

```
ALTER DATABASE my_database REMOVE 'streamingserver' AT LOCATION 'hostA:30316'
```

18 Initializing the Connection Between a Tenant Database and a Streaming Service

Create the streaming analytics schema and database users in the tenant database and start the streaming analytics service (`hdbstreamingserver`).

Context

To allow moving or copying of the tenant database at some point in the future, specify names for all the nodes on the streaming analytics hosts currently being used by this tenant database before initializing the connection. See [Edit the Streaming Node Name \[page 81\]](#) for detailed steps.

If you skip this step now and wish to copy or move this tenant database to another SAP HANA system at a later point, you will then need to uninitialize the connection, name the nodes, and initialize the connection again. See [Copy or Move a Tenant Database \[page 152\]](#) for complete details.

Procedure

1. In a SQL console, connect to the tenant database as a user with the DATA ADMIN privilege.
2. Execute the following `ALTER SYSTEM INITIALIZE SERVICE` statement:

```
ALTER SYSTEM INITIALIZE SERVICE 'streamingserver'  
  [ AT [ LOCATION ] '<streaming_hostname>[:3<instance_number><available  
internal port>' ] ]  
  WITH CREDENTIAL TYPE 'PASSWORD' USING '<streaming_cluster_password>'
```

Example

This example initializes the `streamingserver` service on `hostA`, for internal port 30016, with the streaming cluster password = `MyStreamingClusterPassword`.

```
ALTER SYSTEM INITIALIZE SERVICE 'streamingserver' at 'hostA:30016'  
with CREDENTIAL TYPE 'PASSWORD' USING 'MyStreamingClusterPassword'
```

19 Uninitializing the Connection Between a Tenant Database and a Streaming Service

Remove the streaming analytics schema and database users from the tenant database and stop the streaming analytics service (`hdbstreamingserver`).

Procedure

1. In a SQL console, connect to the tenant database as a user with the DATABASE ADMIN privilege.
2. Execute the following `ALTER SYSTEM UNINITIALIZE SERVICE` statement:

```
ALTER SYSTEM UNINITIALIZE SERVICE 'streamingserver'  
  AT [ LOCATION ] '<hostname>:3<instance_number><available internal port>'  
  [ DROP CASCADE | RESTRICT ]
```

Specifying `CASCADE` allows the service to be uninitialized even if projects are deployed on the streaming analytics host. Uninitialization fails if projects are deployed and you either specify `RESTRICT` (default) or specify nothing.

Example

The following code example uninitialized the `streamingserver` service, even if dependent streaming projects exist.

```
ALTER SYSTEM UNINITIALIZE SERVICE 'streamingserver' drop cascade;
```

20 Activating Content in the SAP HANA Repository

To be able to activate streaming analytics content in the SAP HANA Repository, first manually enable the activation plugin.

Prerequisites

SAP HANA streaming analytics is installed and configured.

Context

Perform these steps for each tenant database you want to have the ability to enable content.

Procedure

1. Enable HTTP(S) access for your target tenant database by executing this SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'database', '<target tenant database name>')
SET ('public_urls', 'https_url') = 'https://<virtual hosts name>:43<instance number>' WITH RECONFIGURE;
```

See *Configure HTTP(S) Access to Tenant Databases via SAP HANA XS Classic* in the *SAP HANA Administration Guide* for additional details.

2. As a tenant database user with the DATABASE ADMIN privilege, run the following SQL statements:

```
CALL
SYS_REPO.GRANT_ACTIVATED_ROLE('sap.hana.xs.admin.roles::HTTPDestAdministrator', 'SYS_STREAMING_ADMIN');
CALL
SYS_REPO.GRANT_ACTIVATED_ROLE('sap.hana.xs.admin.roles::JobAdministrator', 'SYS_STREAMING_ADMIN');
CALL
SYS_REPO.GRANT_ACTIVATED_ROLE('sap.hana.xs.admin.roles::TrustStoreAdministrator', 'SYS_STREAMING_ADMIN');
```

3. From your HANA machine as the <sid>adm user, run following commands:

```
streamingclusterutil -m --task=deploy-du --db-name=<lowercase tenant DB name>
--file=HANA_STREAMING_P.tgz
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/
<lowercase tenant database name>/config/cluster.cfg --install
```


- As a tenant database user with the DATABASE ADMIN privilege, run the following SQL statements:

```
CALL
SYS_REPO.GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT('select','SAP_HANA_STREAMING', 'SYS_STREAMING_ADMIN');
CALL
SYS_REPO.GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT('update','SAP_HANA_STREAMING', 'SYS_STREAMING_ADMIN');
CALL
SYS_REPO.GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT('delete','SAP_HANA_STREAMING', 'SYS_STREAMING_ADMIN');
CALL
SYS_REPO.GRANT_SCHEMA_PRIVILEGE_ON_ACTIVATED_CONTENT('insert','SAP_HANA_STREAMING', 'SYS_STREAMING_ADMIN');
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'DATABASE') SET ('repository', 'jsplugin_sap/hana/streaming/repoplugin/plugin.xsjs') = 'true' WITH RECONFIGURE;
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'DATABASE') SET ('scheduler', 'enabled') = 'true' WITH RECONFIGURE;
```

- From your HANA machine as the <sid>adm user, run the following command for each host used by this tenant:

```
hdbstreamingserver -hdb:installer --config $DIR_INSTANCE/streaming/cluster/<lowercase tenant database name>/config/cluster.cfg --add-host --hdb-host <streaming host name>
```

Results

Now you can activate streaming analytics content in the SAP HANA Repository. For example, you can right-click on content and select **Activate** to deploy and run it in the SAP HANA Repository.

In this section:

[Disabling the Activation of Content in the SAP HANA Repository \[page 161\]](#)

If you no longer wish to activate streaming analytics content in the SAP HANA Repository, disable the activation plugin for your tenant.

20.1 Disabling the Activation of Content in the SAP HANA Repository

If you no longer wish to activate streaming analytics content in the SAP HANA Repository, disable the activation plugin for your tenant.

Prerequisites

You have performed the steps outlined in [Activating Content in the SAP HANA Repository \[page 160\]](#) for your tenant.

Context

Perform these steps for each tenant database you no longer wish to have the ability to activate content.

Procedure

1. As the tenant database user with the DATABASE ADMIN privilege, run the following SQL statement:

```
ALTER SYSTEM ALTER CONFIGURATION ('xsengine.ini', 'DATABASE') SET ('scheduler', 'enabled') = 'false' WITH RECONFIGURE;
```

2. (Optional) If you want to delete the activation plugin delivery unit, run the following command from your HANA machine as the <sid>adm user:

```
streamingclusterutil -m --task=teardown-du --db-name=<lowercase tenant database name> --file=HANA_STREAMING_P.tgz
```

21 Troubleshooting

Resolve issues with SAP HANA streaming analytics.

→ Tip

If an issue arises, review the logs. Streaming analytics writes two project-level logs and a cluster/node-level log. For the names and locations of the logs, see [Log Files \[page 142\]](#).

In this section:

[Understanding Server Messages \[page 163\]](#)

Server messages are produced to provide error and event information. Understand the pieces that make up a message to debug and monitor your system.

[Login and Connection Problems \[page 164\]](#)

Solve issues with authentication and with connecting to servers or projects.

[Project Problems \[page 168\]](#)

Resolve problems with projects.

[Data_streaming Directory Already Exists During Installation \[page 171\]](#)

Issue: If the `data_streaming` directory already exists when installing SAP HANA streaming analytics, an error message appears and the installation fails.

[Java Error When Initializing the streamingserver Service \[page 172\]](#)

Issue: When initializing the `streamingserver` service, the process fails, and returns a Java error.

[Cannot Access Streaming Analytics Pages in SAP HANA Cockpit \[page 172\]](#)

Issue: On the system overview in the SAP HANA cockpit, the links to SAP HANA Streaming Analytics pages do not appear.

[Cannot Start Up Streaming Server \[page 172\]](#)

Issue: The streaming server does not start up.

21.1 Understanding Server Messages

Server messages are produced to provide error and event information. Understand the pieces that make up a message to debug and monitor your system.

The following is the format of log messages:

```
[XXXXX]{-1}{-1/-1} YYYY-DD-MM HH:MM:SS.SSSSS n Streaming UNKNOWN(-0001):[SP-X-YYYYYY](seconds)sp(pid) error message text
```

XXXXX

Thread ID of the thread generating the message.

YYYY-DD-MM HH:MM:SS.SSSSSS	UTC timestamp for the message.
n	Type of error that resulted in the message. Values include e (error), i (info), w (warning), d (debug), or f (fatal).
Streaming	Name of the component that generated the error. This will always be <code>Streaming</code> for Streaming messages.
UNKNOWN(-0001)	File and line number causing the error. This feature is not supported in Streaming log messages.
X	Logging level of the project.
YYYYYY	The module number in which the message is logged, combined with the message number. Used mostly for SAP technical support.
pid	The process ID for the Streaming Processor process.
error message text	The specific error that occurred.

21.2 Login and Connection Problems

Solve issues with authentication and with connecting to servers or projects.

In this section:

[Error: Invalid Login Credentials \[page 165\]](#)

Issue: When you attempt to log in to SAP HANA streaming analytics, an error displays.

[Cannot Connect to the Cluster \[page 165\]](#)

Issue: When running a project, you cannot connect to the SAP HANA streaming analytics cluster.

[A Studio Project Does Not Run, Reports Login Failure \[page 166\]](#)

Issue: You are unable to run a project in a cluster and errors display.

[A Utility Fails to Connect to the Server \[page 166\]](#)

Issue: You are unable to connect to the streaming analytics server using a command-line utility such as `streamingprojectclient` or `streamingsubscribe`.

[A Utility Fails to Connect to a Project \[page 167\]](#)

Issue: You cannot connect to a project using a command-line utility such as `streamingprojectclient`.

[An External Adapter Fails to Start \[page 167\]](#)

Issue: Your attempts to start an external adapter fail.

[An Unmanaged External Adapter Fails to Connect to a Project \[page 168\]](#)

Issue: When starting an external adapter in unmanaged mode without editing its sample XML configuration file, the adapter fails to start.

[ODBC Data Source Name Not Found \[page 168\]](#)

Issue: When connecting to an ODBC data service, you receive the following error: `[[unixODBC][Driver Manager]Data source name not found, and no default driver specified]`.

[Nodes Are Missing from Cockpit and streamingclusteradmin \[page 168\]](#)

Issue: On a system with multiple streaming analytics hosts, one or more do not appear in `streamingclusteradmin` or the SAP HANA cockpit.

21.2.1 Error: Invalid Login Credentials

Issue: When you attempt to log in to SAP HANA streaming analytics, an error displays.

```
[error] security : Authentication failure:Invalid login credentials
```

Solution: Verify that the user ID and password are valid and spelled correctly.

21.2.2 Cannot Connect to the Cluster

Issue: When running a project, you cannot connect to the SAP HANA streaming analytics cluster.

Solution: Check the status of the streaming analytics server, and restart if necessary.

1. Verify that the streaming analytics host is running. See [Verify that the Streaming Service is Installed and Running \[page 27\]](#).
2. If the host is not running, start it. See [Start the Streaming Service Manually \[page 28\]](#).
3. If the host cannot be started, check the streaming analytics log files in the SAP HANA trace directory for errors.
If you see the error code 700283, there are problems with your streaming analytics license and you may need to get a permanent license key. See *Install a Permanent License* in the *SAP HANA Streaming Analytics: Installation and Update Guide* for detailed instructions.
4. Verify that the SAP HANA index server is running.

i Note

A common cause for issues with starting a streaming analytics host is the SAP HANA index server not running.

Related Information

[Verify that the Streaming Service is Installed and Running \[page 27\]](#)

[Start the Streaming Service Manually \[page 28\]](#)

[Install a Permanent License](#)

21.2.3 A Studio Project Does Not Run, Reports Login Failure

Issue: You are unable to run a project in a cluster and errors display.

```
Failed to connect to server "esps://<host>:3<instance-number><port-number>".  
Reason: "Failed to login server"
```

Studio reports these errors when it has an SSL mismatch with the streaming analytics server. For example, since SSL is enabled on the server by default and cannot be disabled, the mismatch results from the studio connection definition for that server not specifying SSL.

Solution: Correct the connection definition in studio. For details, see [Configuring a Cluster Connection](#).

21.2.4 A Utility Fails to Connect to the Server

Issue: You are unable to connect to the streaming analytics server using a command-line utility such as `streamingprojectclient` or `streamingsubscribe`.

The command might return this message:

```
Couldn't connect to server, XML-RPC Fault(-504)
```

Utilities affected by this issue include:

- `streamingprojectclient`
- `streamingcnc`
- `streamingconvert`
- `streamingkdbin`
- `streamingkdbout`
- `streamingquery`
- `streamingsubscribe`
- `streamingupload`

Solution: Verify the following:

- You are using the correct host name, port number, and login details in the command.
- The server is up and reachable. You can ping it.
- You are using `-e` correctly. Use the `-e` flag of the utility only if SSL is enabled on the server. For example:

```
> streamingprojectclient -c your_user_name:your_password  
-p <host>:51011/your_workspace/your_project -e  
streamingprojectclient> loglevel 7  
streamingprojectclient> stop  
streamingprojectclient> quit
```

For details on command syntax, see the *SAP HANA Streaming Analytics: Utilities Guide*.

Related Information

[SAP HANA Streaming Analytics: Utilities Guide](#)

21.2.5 A Utility Fails to Connect to a Project

Issue: You cannot connect to a project using a command-line utility such as `streamingprojectclient`.

For example:

```
% streamingprojectclient -c your_user_name:your_password -p <host>:51011/  
your_workspace/your_project  
ASAP_loginToCluster( your_user_name ) failed, status = -1
```

Solution: Verify the following:

- The user ID and password are valid and spelled correctly.
- The workspace name and project name are correct.
- If access control is enabled, the user has permissions that allow access to the project.
- The project is running. Use `streamingclusteradmin` to verify and start the project, if necessary.

21.2.6 An External Adapter Fails to Start

Issue: Your attempts to start an external adapter fail.

When attempting to run an external adapter an error message similar to the following appears:

```
Failed call to:https://<Streaming-hostname>:61308/RPC2 (Failed to read server's  
response: <Streaming-hostname>) java.io.IOException: Failed call to:https://  
<Streaming-hostname>:61308/RPC2 (Failed to read server's response: <Streaming-  
hostname>)
```

This error is an example of the streaming analytics server not being resolved.

Solution: Use the `ping` command to verify that the hostname of the server to which you are trying to connect can be resolved. If the hostname cannot be resolved:

1. Determine the IP address of the host on which the streaming analytics server is running. Run this command from that machine:

```
nslookup <hostname>
```

2. Add the following line to C:\Windows\System32\drivers\etc\hosts (Windows) or /etc/hosts (UNIX):

```
<ip-address-of-server-hostname> <Server-hostname>
```

21.2.7 An Unmanaged External Adapter Fails to Connect to a Project

Issue: When starting an external adapter in unmanaged mode without editing its sample XML configuration file, the adapter fails to start.

The external adapter may be unable to connect to the example workspace specified in the sample XML adapter configuration file if the URI specified in the file uses `esp` instead of `esps`. Using `esp` causes the adapter to fail to connect because SSL is enabled by default on SAP HANA streaming analytics and cannot be disabled.

Solution:

1. Open the adapter's XML configuration file in a notepad program and ensure the URI uses `esps`, not `esp`.
2. If using one of the adapter examples provided with your installation, edit the `set_example_env.bat` or `set_example_env.sh` script file to specify:

```
set ADAPTER_EXAMPLE_CLUSTER_NODE_PROTOCOL=esps
```

21.2.8 ODBC Data Source Name Not Found

Issue: When connecting to an ODBC data service, you receive the following error: `[[unixODBC] [Driver Manager]Data source name not found, and no default driver specified].`

Solution: Define the data source name (DSN) for the ODBC connection in the `$DIR_INSTANCE/streaming/cluster/<lowercase-db-name>/config/odbc.ini` file.

21.2.9 Nodes Are Missing from Cockpit and streamingclusteradmin

Issue: On a system with multiple streaming analytics hosts, one or more do not appear in `streamingclusteradmin` or the SAP HANA cockpit.

Solution: On each streaming analytics host, edit the `/etc/hosts` file. Add an entry for the host using the format `<ip-address> <fully-qualified-hostname> <hostname>`. If the `hosts` file has an entry for the host using the format `127.0.0.1 <fully-qualified-hostname> <hostname>`, remove this entry.

21.3 Project Problems

Resolve problems with projects.

In this section:

[A Utility Fails to Connect to a Project \[page 169\]](#)

Issue: You cannot connect to a project using a command-line utility such as `streamingprojectclient`.

[A Project Triggers Java Out-of-Memory Errors \[page 169\]](#)

Issue: When a project runs, you receive out-of-memory errors from the Java virtual machine.

[Published Data Lost When Project Fails \[page 170\]](#)

Issue: A subscriber receiving data through a stream does not receive all the data sent by the publisher before the streaming analytics node shuts down unexpectedly.

[User Cannot Perform Tasks Despite Being Granted Permissions \[page 170\]](#)

Issue: Despite having been granted permission to perform a task, a user receives an error message stating that they do not have the required permission for the operation.

[Error: Could Not Allocate Controller \[page 171\]](#)

Issue: When you start a streaming analytics project, the following error displays: `Could not allocate Controller`.

21.3.1 A Utility Fails to Connect to a Project

Issue: You cannot connect to a project using a command-line utility such as `streamingprojectclient`.

For example:

```
% streamingprojectclient -c your_user_name:your_password -p <host>:51011/  
your_workspace/your_project  
ASAP_loginToCluster( your_user_name ) failed, status = -1
```

Solution: Verify the following:

- The user ID and password are valid and spelled correctly.
- The workspace name and project name are correct.
- If access control is enabled, the user has permissions that allow access to the project.
- The project is running. Use `streamingclusteradmin` to verify and start the project, if necessary.

Related Information

[Cluster Administrative Tool \[page 148\]](#)

21.3.2 A Project Triggers Java Out-of-Memory Errors

Issue: When a project runs, you receive out-of-memory errors from the Java virtual machine.

Solution: Modify the project configuration (`.ccr`) file to increase the heap size for the project's Java virtual machine.

21.3.3 Published Data Lost When Project Fails

Issue: A subscriber receiving data through a stream does not receive all the data sent by the publisher before the streaming analytics node shuts down unexpectedly.

Solution: Enable guaranteed delivery on the stream, or replace it with a guaranteed delivery-enabled window.

Guaranteed delivery (GD) uses log stores to ensure that a GD subscriber registered with a GD stream or window receives all the data processed by that stream or window even if the client is not connected when the data is produced. GD is supported on streams and windows and each GD stream or window requires a GD log store.

i Note

GD-enabled windows are not recommended on a project configured for hot standby. The shared disk requirements for GD log stores are incompatible with the continuous synchronization that enables a primary instance to fail over quickly to its secondary instance.

See the *SAP HANA Streaming Analytics: Developer Guide* for information on guaranteed delivery.

Related Information

[Guaranteed Delivery](#)

21.3.4 User Cannot Perform Tasks Despite Being Granted Permissions

Issue: Despite having been granted permission to perform a task, a user receives an error message stating that they do not have the required permission for the operation.

Solution: You may have granted permission to a user that does not exist in the SAP HANA system. Verify that the username for the user to whom you are granting a permission matches their SAP HANA username:

- If you are using the SAP HANA cockpit, verify that the username under the **Users** tab in the Manage Streaming Permissions page matches the SAP HANA username.

i Note

When you add a user to the Manage Streaming Permissions page, you are not creating an SAP HANA user. Users listed under the **Users** column are only users that have been granted at least one streaming analytics permission or role.

- If you are using the `streamingclusteradmin` utility, use the `get users` and `get perms` commands to verify that the streaming analytics username matches the SAP HANA login credentials.
- If you are using the SAP HANA studio, you are not a restricted user, and you have the necessary admin permissions, you can reference the list of all current SAP HANA users by doing the following:
 1. In the **Systems** view of the **SAP HANA Administrator Console** perspective, log on to the necessary system as an admin user.

2. Expand the **Security** folder.
3. Expand the **Users** folder to see a comprehensive list of all SAP HANA users on the system. Use this list as a reference to ensure that the username you are granting permissions to on the SAP HANA cockpit Manage Streaming Permissions page matches the username you see in this list.

For information on how to set up and authenticate SAP HANA users, see the *Managing SAP HANA Users* section in the *SAP HANA Administration Guide*.

Related Information

[SAP HANA Administration Guide](#)

21.3.5 Error: Could Not Allocate Controller

Issue: When you start a streaming analytics project, the following error displays: `Could not allocate Controller.`

Solution: This error occurs when SAP HANA and streaming analytics are uninstalled and reinstalled with the same SID without first deleting the `/hana/data_streaming/<SID>` folder.

1. Uninstall streaming analytics.
2. Delete the `/hana/data_streaming/<SID>` folder.
3. Reinstall streaming analytics.

21.4 Data_streaming Directory Already Exists During Installation

Issue: If the `data_streaming` directory already exists when installing SAP HANA streaming analytics, an error message appears and the installation fails.

The existence of the directory before beginning an installation may result in permission errors, or an error message and installation failure.

Solution:

1. Delete the directory if it is empty or if it contains a single, empty subdirectory. For example, for `/hana/data_streaming/ABC`, ABC is empty.
2. If the subdirectory is not empty, ensure that the directory (for example, `/hana/data_streaming`) is owned by the root user and each subdirectory (for example, `/hana/data_streaming/ABC`, `/hana/data_streaming/DEF`) is owned by the `<sid>adm` of the system ID.

21.5 Java Error When Initializing the streamingserver Service

Issue: When initializing the `streamingserver` service, the process fails, and returns a Java error.

```
[403]: internal error: Call to streamingclusterutil failed
```

This behavior occurs when using the `/etc/hosts` file to override virtual host name resolution instead of using the DNS.

Solution:

Edit the `/etc/hosts` file on both the SAP HANA host and streaming worker host, adding the same SAP HANA server IP address and machine name with alias entry you made to the `/etc/hosts` file on the local machine.

For example:

```
10.2.4.3 machine_name-alias1.corp
10.2.4.3 machine_name-alias2.corp
```

If necessary, restart the tenant database before re-executing the initialize service statement.

21.6 Cannot Access Streaming Analytics Pages in SAP HANA Cockpit

Issue: On the system overview in the SAP HANA cockpit, the links to SAP HANA Streaming Analytics pages do not appear.

Solution: Verify that the database user has the CATALOG READ system privilege.

21.7 Cannot Start Up Streaming Server

Issue: The streaming server does not start up.

Solution 1: If you have edited the node name of an existing streaming analytics host, the node name in `streamingserver.ini` does not match the name in the cluster configuration file. To solve this, change the node name back to its default value.

To check the correct node name, run the following command as the `<sid>adm` user:

```
hdbstreamingserver -u --config $STREAMING_HOME/./cluster/<db-name>/config/cluster.cfg --show
```

The node name is in the Nodes section of the output:

```
<Cluster>
```

```
<Nodes>
  <Node enabled="true" name="<node_name">">
```

Set the node name in `streamingserver.ini` to this value and restart all running nodes for the changes to take effect.

Solution 2: This issue may also be caused by an invalid SAP HANA license. Check the streaming server log file. In the file, you see the following messages:

```
Caused by: com.sap.db.jdbc.exceptions.JDBCDriverException: SAP DBTech JDBC:
[591]: internal error: only commands for license handling are allowed in current
state
...
FATAL - CODE_700418 | Could not load config from database
FATAL - CODE_700412 | Factory of new node failed
```

If you see these messages, there is a problem with your SAP HANA license and you may need to get a permanent license key. See [Install a Permanent License](#) for detailed steps.

Related Information

[Edit the Streaming Node Name \[page 81\]](#)

[Check the Current License Key](#)



[Managing SAP HANA Licenses](#)

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

© 2022 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.