

GoldenGate Stream Analytics

Release Notes for GoldenGate Stream Analytics



F18430-21
May 2024



GoldenGate Stream Analytics Release Notes for GoldenGate Stream Analytics,

F18430-21

Copyright © 2018, 2023, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	iv
Documentation Accessibility	iv
Conventions	iv
Related Documents	iv

Preface

This document lists the known issues and limitations in Oracle Stream Analytics.

Topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Conventions](#)
- [Related Documents](#)

Audience

This document is intended for all users of Stream AnalyticsOracle Stream Analytics.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Accessible Access to Oracle Support

Oracle customers who have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Related Documents

Documentation for Oracle Stream Analytics is available on [Oracle Help Center](#).

Also see the following documents for reference:

- *Understanding Oracle Stream Analytics*
- *Quick Installer for Oracle Stream Analytics*
- *Developing Custom Jars and Custom Stages in Oracle Stream Analytics*
- *Spark Extensibility for CQL in Oracle Stream Analytics*
- *Using Oracle Stream Analytics*

1.1 GoldenGate Stream Analytics Release Notes

This document contains release information such as new features, known issues, and corrected problems for GoldenGate Stream Analytics 19.1.

Oracle recommends you to review this document before installing or working with the product.

1.1.1 New Features and Improvements

This section provides a list of new features added in a release.

1.1.1.1 Release 19.1.0.0.9 - May 2024

This section lists the new features and upgrades in GGSA 19.1.0.0.9.

1.1.1.1.1 Spark and Kafka Versions Upgrade

From GGSA 19.1.0.0.9:

- Spark library version upgraded to 3.5.0
- Kafka library version upgraded to 3.7.0

For more information on downloading the libraries, and installing GGSA, see [Installing GoldenGate Stream Analytics](#).

1.1.1.1.2 Select Connection while Importing a Pipeline

GGSA supports selecting a new connection or using the existing connection while importing a pipeline and its resources. See [Importing a Pipeline](#).

1.1.1.1.3 Support for SASL Security and Protocol mechanisms for Kafka Connections

SASL security and protocol mechanism support added for Kafka connections.

1.1.1.2 Release 19.1.0.0.8 - August 2023

This section lists the new features and upgrades in GGSA 19.1.0.0.8.

1.1.1.2.1 Spark and Kafka Versions Upgrade

From 19.1.0.0.8, the Spark and Kafka library versions are upgraded to 3.4.0.

For more information on downloading the libraries, and installing GGSA, see [Installing GoldenGate Stream Analytics](#).

1.1.1.3 Release 19.1.0.0.7 - March 2022

This section lists the new features added in the GGSA 19.1.0.0.7.

1.1.1.3.1 Public REST APIs

The following Public REST APIs are added for pipeline management and user management:

Pipeline Management

- Export
- Import
- List
- Publish
- Unpublish

User Management

- Create
- List
- Retrieve
- Update

See REST API for GoldenGate Stream Analytics.

1.1.1.3.2 Custom Stage Pass-through Input

Custom Stage Pass-through Input

A custom stage pass-through input is supported to pass through all the attributes from previous stages in a pipeline, in addition to the ones that are defined in the `outputspec` of the custom stage class. You can enable this feature by adding an optional parameter `canPassThroughInput=true`.

If you do not set this parameter, it is defaulted to false, and only the attributes defined in the `outputspec` will be output in the custom stage. To implement `PassThroughInput`, it is also required to implement `EventSpecAccessor` and set `outputspec` from `setoutputspec`, and get `indexesToCopy` to set the values for attributes from previous stages.

1.1.1.3.3 ONNX Predictive Model

GGSA supports the ONNX predictive model, to enhance the product machine learning capability.

1.1.1.4 Release 19.1.0.0.6_MLR33682853 - December 2021

This release addresses a security vulnerability:

CVE-2021-44228 and CVE-2021-45046

Oracle recommends that you upgrade to this patch of GoldenGate Stream Analytics to include this security fix.

1.1.1.5 Release 19.1.0.0.6_MLR33251626 - September 2021

This section lists the new features added in the GGSA 19.1.0.0.6_MLR33251626.

1.1.1.5.1 Segment Detector Pattern

Use the Segment Detector Pattern to detect contiguous events (range/segment) having unchanged values for the selected attributes, over a specified period of time.

1.1.1.5.2 Pivot Pattern

Use the Pivot pattern to pivot all or selected attributes of an incoming event, into new columns, based on the selected pivot value and key provided.

1.1.1.5.3 Custom CQL Stage

You can add a Custom CQL stage for pipelines, to write a CQL query to perform any *Select* operation on the data of the previous stage. Currently, creating certain business rules are not supported by the existing pipeline stages or patterns stages. You can use the CQL query as part of the Custom CQL stage, for such requirements.

1.1.1.6 Release 19.1.0.0.6 - April 2021

This section lists the new features added in the GGSA 19.1.0.0.6 bundle patch release.

1.1.1.6.1 Kafka MTLS Support

GGSA supports a 2-way authentication between the user and the Kafka broker. Users can enable the MTLS authentication by providing the Bootstrap URL, SSL truststore and keystore information in **System Settings**. Users can connect to MTLS enabled Kafka set-up while creating Kafka connections.

1.1.1.6.2 GoldenGate Full Records

The GoldenGate Extract process captures either full data records or transactional data changes, depending on the configuration parameters. To minimize the overhead or performance impact on the transactional database, GGSA users configure the Extract to capture only the transactional changes. This also helps to reduce payload size needs to transfer over the network, thus increasing the performance and security. But a few customers also need the unchanged columns (full data records), making them available to the processes that require up-to-date data feed, or to replicate this data to various big data targets for analysis.

To enable streaming of full data record (value of all fields), GGSA provides the **Generate Full Records** option while creating a GoldenGate stream. Enable this option to stream all the records, irrespective of the database transactional changes made to a single row, a subset or all the columns of a row.

1.1.1.6.3 Big Data Connections and Targets

GGSA supports the following new Big Data connections and targets:

- ElasticSearch
- HBase
- MongoDB

1.1.1.6.4 Continuous Merge Pattern

Use the Continuous Merge pattern to merge two or more continuous events into a single event, based on the key attributes.

1.1.1.6.5 Oracle Machine Learning Service Pattern (Preview Feature)

Use the Oracle Machine Learning (OML) Service pattern to use OML models to apply scoring on the ingested events.

1.1.1.6.6 Internal Kafka Topics

The internal Kafka topics and Group ID's used by GGSA are standardized to the following naming conventions:

Kafka Topics

Topic	Resource	Operations
sx_backend_notification_<UUID>	Topic	CREATE,DELETE,DESCRIBE,READ,WRITE
sx_messages_<UUID>	Topic	CREATE,DELETE,DESCRIBE,READ,WRITE
sx_<application_name>_<stage_name>_public	Topic	CREATE,DELETE,DESCRIBE,READ,WRITE
sx_<application_name>_<stage_name>_draft	Topic	CREATE,DELETE,DESCRIBE,READ,WRITE
sx_<application_name>_public_<offset_number>_<stage_name>_offset	Topic	CREATE,DELETE,DESCRIBE,READ,WRITE

Group IDs

Group ID	Resource	Operations
sx_<UUID>_receiver	Group	DESCRIBE,READ
sx_<UUID>	Group	DESCRIBE,READ
sx_<application_name>_public_<offset_number>_<stage_name>	Group	DESCRIBE,READ

1.1.1.6.7 Zookeeper Dependency Removed

Zookeeper configuration is no longer a prerequisite to use GGSA.

1.1.1.7 Release 19.1.0.0.5 - September 2020

This section lists the new features added in the GGSA 19.1.0.0.5 bundle patch release.

1.1.1.7.1 Embedded Ignite Cluster

GGSA implements an Embedded Ignite Cache. Ignite can be used as a target, or a reference for any pipeline, to persist events. In case of GoldenGate, the pipeline persists incoming events containing full record, and also updates the cache with modified events, making available the latest modified records.

You can start or stop the Ignite Cluster from the **Cache Cluster** tab on the **System Settings** page.

1.1.1.7.2 NFS Target

GGSA supports a Network File System (NFS) target.

1.1.1.7.3 Ignite Connection

GGSA supports both embedded Ignite cache connection and external Ignite cache connection.

You can enable the internal cache from the **Cache Cluster** tab on the **System Settings** page. Enabling internal cache creates and deploys ignite cluster on all spark worker nodes.

You can create a connection to external Ignite by selecting Ignite Cache type, while creating a new connection from the **Catalog** tab. Provide Ignite server IP address, Port numbers for unsecured cluster, and security credentials for protected cluster.

1.1.1.7.4 Ignite Target

GGSA now supports and uses Ignite as a target to write Json data.

GGSA supports two types of Ignite cluster:

- External cluster managed outside GGSA
- Embedded Ignite Cluster, started as Spark application and managed by GGSA

With the addition of Ignite target, GGSA can write data to both the internal and external ignite cache clusters.

1.1.1.7.5 Ignite Reference

Ignite reference enriches the stream data with the Ignite cache cluster reference data, at query stage. Ignite reference has data in key-value pair; key is String type and value is a Json object. GGSA supports only single-value equality join in query stage to reference.

1.1.1.7.6 GoldenGate Big Data Connections and Targets

GGSA supports new GoldenGate Big Data connections and targets:

- OCI Object Store
- AWS S3
- HDFS

- Microsoft Azure Datalake Gen-2
- Hive

1.1.1.7.7 Split Pattern

Use the Split pattern to transform a single event from a stage into multiple events, by splitting the value field. For example, you can flatten an array of json element or json object from the source, to individually process it and also to push it to some targets.

1.1.1.7.8 Cache Monitoring

You can monitor, reset, and reload the caches in both the internal (embedded) and external clusters, using the **Monitor** option on the GGSA application's homepage.

1.1.1.7.9 Advance Kafka Settings

The **Advance Kafka Settings** that is added to the **System Settings** page, enables you to provide Bootstrap URL, SSL truststore, and the SASL information. The bootstrap URL is given preference over the zookeeper, to connect to Kafka, and to even deploy the pipelines.

1.1.1.7.10 Protection Policy

A **Protection Policy** feature is added under **Hadoop Authentication** of the **Environment** tab, on the **System Settings** page. Protection policy defines how data is transferred between hadoop services and clients (GGSA). Policy types are authentication, integrity and privacy.

1.1.1.8 Release 19.1.0.0.3 - April 2020

This section lists the new features added in the OSA 19.1.0.0.3 bundle patch release.

1.1.1.8.1 Configure POST and PUT methods in REST Target

REST target can send requests to REST end-point, using HTTP POST and PUT methods. Default setting is POST.

1.1.1.8.2 Kafka Offset Changes

Infer shape from Kafka topic using Earlier/ Latest offset

Detects the shape based on the earliest or the latest offset of the kafka topic. The default option is earliest. Use this option if currently no data is flowing in the topic.

Allow User to reset the offset during pipeline publish

You have two new options on the Pipeline Settings screen:

- **Input Topics Offset** — select the Kafka topic offset value from this drop-down list. Values listed are **earliest** and **latest**. The default value is **latest**. When you publish the pipeline for the first time, the input stream is read based on the offset value you have selected in this drop-down list. On a subsequent publish, the value you have selected here is not considered, and the input stream is read from where it was last left off.
- **Reset Offset** — select this checkbox to read the input stream based on the offset value selected in the **Input Topics Offset** drop-down list.

Allow User to build dashboard from Beginning or Latest offset

While creating a dashboard, you can build dashboard from the beginning or the latest offset. Available options are **latest** and **earliest**. If you select the latest option, the dashboard shows only the new records that have been received after opening the dashboard. If you select the earliest option, then the records are read from the beginning of the OSA pipeline topics.

1.1.1.8.3 Define Key Field while creating a Kafka Target

You can define key fields while defining Kafka Target. Data will be partitioned based on the selected keys.

- You can select multiple fields as keys
- Key selection is not mandatory
- Key selection is allowed only when data format is JSON

1.1.1.8.4 ToJson Pattern

ToJson Pattern transforms event(s) coming from a stage in the pipeline into a JSON text.

Use this pattern to transform multiple events into a single JSON document, and send it to a downstream system through OSA pipeline targets. For example, you can configure a Database target after the toJson pattern stage, to write the json payload (of a single or multiple events), into a database table.

1.1.1.8.5 Condition Expression Builder Enhancements

The Condition Expression Builder enhancements enable you to:

- Create blocks without adding condition expression, which you can add at any later stage
- Link the blocks using AND/ OR
- Define complex conditions. Example:

```
IF (
  ((op_type == 'I') AND (after_SHHOLD is not null))
  OR (
    ((before_SHHOLD is null) OR (before_SHHOLD = " "))
    AND (
      after_SHHOLD is not null
    )
  )
)
```

1.1.1.8.6 Delete Intermediate Stage(s) from a Pipeline

You can now delete any intermediate Stages in an unpublished pipeline.

Note:

- You cannot delete a Source stage.
- When you delete any intermediate stage, all its child stages will also be deleted.

1.1.1.8.7 OCI Notification Service

OSA is integrated with OCI notification service, to broadcast messages to distributed components using publish subscriber pattern. This service delivers secure, highly reliable, low latency, and durable messages to applications hosted on OCI, and externally. To integrate with OCI Notification service, you have to define an OCI Notification service topic. You have options to create an **ONS** type connection and a **Notification** type target to enable this feature.

**Note:**

Currently, only JSON data format type is supported for the data to be published.

1.1.1.9 Release 19.1.0.0.2 - February 2020

This section lists the new features added in the OSA 19.1.0.0.2 bundle patch release.

1.1.1.9.1 Integration with Oracle Goldengate

The Integration with Oracle Goldengate feature enables you to connect to a Goldengate microservices instance, and generate a change data stream (GG Change Data) from an extract process. Configuring GG Change Data creates a new Kafka topic which you can then use to create a Stream.

1.1.1.9.2 Coherence Target

OSA now supports Coherence Target, to output data to a Coherence server. The output OSA events (Tuple) are pushed to external coherence cluster using the cache name and the coherence server details that the user provides.

1.1.1.9.3 Wallet Support for DB Target/Reference & Metadata

Oracle Stream Analytics User Interface allows a user to create and test a ATP/ADW database connection.

The user can use this database connection to create a database reference or a database target.

1.1.1.9.4 Support MySQL DB for Reference

This feature allows a user to create MySQL database tables as a reference in OSA. It supports a JDBC based connection to a MySQL database.

1.1.1.9.5 Integration with OCI Streaming Service & Secure Kafka

Oracle Stream Analytics now supports ingesting from OCI Streaming service using the Kafka compatibility feature of OSS.

1.1.1.9.6 Oracle Advanced Queue Source

OSA can read messages from Oracle Advanced Queue. This option is available as general JMS connector. User can create an Advanced Queue connection and use it to create a JMS type stream.

1.1.1.9.7 Support for secure REST Target

OSA now supports REST target type, which is SSL enabled and requires Basic authentication.

In addition to SSL enabled if the end point requires authentication, then you can pass it as a custom header field. Currently, OSA supports only Basic authentication. Example Custom Header - `Authorization Basic admin:admin`.

1.1.1.9.8 Kerberos Authentication

Hadoop now supports Kerberos Authentication for OSA running on Yarn-based Spark cluster. When using Kerberos authentication, the user is authenticated by obtaining a Kerberos ticket from the Kerberos server.

1.1.1.9.9 SSL And Authentication Enabled REST Endpoint

Creating a target of type "REST" which is SSL enabled and requires Basic authentication: If the REST End point is only SSL enabled, then you could connect in one of the following ways:

- Upload Truststore file, enter truststore password. Truststore password is optional.
- If you do not have the Truststore file and password, you can connect to the REST endpoint by clicking **Trust password**

Note:

Using the Trust password option, will connect using untrusted certificates and is an insecure connection.

In addition to SSL enabled if the end point requires authentication, then you can pass it as a custom header field. Currently, OSA supports only Basic authentication. The custom header would be, for example, "Authorization Basic admin:admin".

1.1.1.9.10 New Data Patterns

OSA supports 4 new Data Patterns:

- **Current And Previous Pattern** - automatically correlates the current and previous events
- **Delay Pattern** - delays delivering the event to downstream node in the pipeline for specified number of seconds
- **Row Window Snapshot** - dumps entire window contents to downstream node on the arrival of a new event. Window capacity is limited by specifying the maximum number of events.
- **Time Window Snapshot** - dumps entire window contents to downstream node on the arrival of a new event. Event in the window expires after specified time range.

1.1.1.9.11 Backpressure in OSA

Backpressure is Spark's way to ensure stability in the streaming application, which means the streaming application receives data only as fast as it can process.

When the Backpressure feature is enabled, a signal is passed from the downstream components towards the upstream components within Spark, based on the present batch processing, and scheduling delay statistics.

The Backpressure is currently enabled by default for all published applications, regardless of whether HA is enabled. Also, this setting is enabled for both Yarn as well as the Standalone cluster manager.

1.1.1.9.12 IN Filter Supports DB Column Lookup

You can use the IN filter to refer to a column in a database table. When you change the database column values at runtime, the pipeline picks up the latest values from the DB column, without republishing the pipeline.

1.1.1.10 Release 19.1.0.0.1 - November 2019

This section lists the new features added in the OSA 19.1.0.0.1 bundle patch release.

1.1.1.10.1 Custom Tile Layer

Oracle Stream Analytics now allows you to customize the zoom levels for a tile layer, in specific cases where you need a detailed, higher zoom view of confined spaces such as, restaurants, airports, etc. Once you have modified the global parameters to customize the tile layer, the map is updated to use the custom tile layer. These customizations, will then be applied to all geofences.

1.1.1.10.2 Database as a Target

You can now use Oracle Database as a target for a pipeline.

1.1.1.10.3 Jetty Properties

A new `jetty.properties` file is available at `OSA-19.1.0.0.0/osa-base/etc/jetty.properties`, where you can modify certain security properties:

- **jetty.session.timeout**
You can set the timeout for OSA web session.
- **host.headers.whitelist**
You can restrict the x-forwarded-host header values to the values defined with this property.
- **xforwarded.host.headers.whitelist**
You can restrict the x-forwarded-host header values to the values defined with this property.
- **response.headers.list**
A comma separated list of response headers, which will be sent along with response for every request.

1.1.1.10.4 Spark Authentication

User authentication is now enabled for the Spark standalone server master console. When the Spark master console is accessed, you will be prompted to enter your username and password. The console page is loaded only if your login credentials are correct. You can change the username and password from the OSA System Preferences Menu. The username and password fields are left blank, by default.

1.1.1.10.5 Application Monitoring

The status of a published application is now displayed on the catalog page. Refresh the catalog page on the UI, to view the current status of any published application. The status is displayed as either 'Killed' or 'Running'.

1.1.1.10.6 Log Level for Published Pipelines

When publishing a pipeline, you can now set the log level from the UI.

1.1.1.11 Release 19.1.0.0.0 - August 2019

This section lists the new features added in the OSA 19.1.0.0.0 release.

1.1.1.11.1 Advanced Pipeline

A pipeline defines the pipeline logic and is a sequence of data processing stages. A stage can be one of the following types – Query, Pattern, Rule, Query Group, Custom, and Scoring. A pipeline starts with a stream stage and can have one or more children of any type such as Query, Rule Pattern etc.

Oracle Stream analytics 19.1 introduces advanced pipeline features like splitting a pipeline into multiple branches at any stage as per your application logic. You can end each branch with one or more targets. It provides enhanced editing capabilities like editing intermediate stages and the changes due to that are propagated to downstream stages. You can now edit the sources even if those are already being used in any pipelines. With a pipeline, you can switch the streams in the source and will be notified of the corrective actions you will need to accommodate the change in source. You can now also use output of an existing stage as one of the source for another stage in the pipeline.

1.1.1.11.2 Supported Window Types

Window support in Oracle Stream Analytics in 19.1 is enhanced by the addition of following new types:

- Current Year
- Current Month
- Current Day
- Current Hour
- Current Period
- Constant Value
- Variable Range
- Partition Window Types

1.1.1.11.3 Import/Export

The export and import feature lets you migrate your pipeline and its contents between Oracle Stream Analytics systems (such as development and production) in a matter of few clicks.

Oracle Stream Analytics 19.1 now supports the operation on all types of artifacts including Visualizations, Cubes, Dashboards, Scoring Stage and Custom Stage. When you import and export, any custom jar that you use in the pipeline is also taken care during the operation. It allows you to review the export bundle, edit names and description of artifacts and manage the tags during the export. During import, it allows you to inspect the import bundle, select items to import and manage tags at the time of import.

1.1.1.11.4 LDAP Integration

The Lightweight Directory Access Protocol (LDAP) is an open source application accepted across various industries. This application protocol is used for obtaining and maintaining distributed directory information services over a network using an Internet Protocol (IP).

Oracle Stream Analytics makes use of the LDAP support for Jetty OSA. With this feature, you can use the directory information services for authentication/user management. The authentication/user management can be through either internal LDAP or external LDAP.

1.1.1.11.5 Coherence POJO

Oracle Stream Analytics 19.1 enables the support for referencing coherence cache where cache value is stored as POJO. The current support involves cache value as a `Map<String, Object>`. Oracle Stream Analytics application can enrich stream data with POJO data.

1.1.1.11.6 IN Operator Support in Filters

The IN Operator is supported in filters at various stages of the application in Oracle Stream Analytics. This option is available as an operator in the drop-list where you apply any filter.

1.1.1.11.7 Multihost Coherence

Oracle Stream Analytics now supports configuring multiple hosts while creating the coherence connection. In the coherence connection application, you can now specify multiple server URLs.

Oracle Stream Analytics supports POJO and Map for Coherence cache.

1.1.1.11.8 Webtier HA

Oracle Stream Analytics 19.1 supports High Availability for its Webtier. Based on the expected load on the Oracle Stream Analytics 19.1 Webtier, you can now configure multiple hosts running the webtier and use NGNIX as a load balancer to configure the High Availability of the OSA Webtier.

1.1.1.11.9 Google Tile Layer

Oracle Stream Analytics has added the support for Google tile layer which displays Google maps in Spatial patterns and visualizations. You can configure the Google tile layer under User Preferences of the Map tab.

**Note:**

In order to use Google maps tile layer, the usage of the maps must meet the terms of service defined by Google (<http://code.google.com/apis/maps/faq.html#tos>).

1.1.2 Known Issues and Workarounds

This section lists the known issues and workarounds for a Oracle Stream Analytics release.

1.1.2.1 Release 19.1.0.0.8 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.8.

1.1.2.1.1 Edit DB Target fails to detect the Updated/ Added columns in a Table

Issue

Changes to the columns in a DB table are not reflected in the DB Target Edit dialog box.

Workaround

If you have multiple tables created in the DB, then in the DB Target Edit dialog box, switch to another table, and then switch back to the table in use. This will update the fields.

If you have only one table, create a dummy table in the DB and then toggle between them, to refresh the changes.

1.1.2.2 Release 19.1.0.0.7 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.7.

1.1.2.2.1 SASL Property in Kafka Connection is not Migrated during Upgrade

Issue

After upgrading GGSA environment from 19.1.0.0.5.1 to 19.1.0.0.6.1, the SASL property in Kafka connections is not migrated correctly and is always set to `false`. When manually setting the property to `true`, any user or password information in the connection is lost.

Workaround

Manually reconfigure the connection credentials.

1.1.2.2.2 Error in Pipeline using a DB Reference with a Hyphen in Display Name

Issue

Using a DB Reference with a hyphen in its display name throws an error in pipeline.

Workaround

Do not use unsupported special characters in the Oracle DB Reference display name.

1.1.2.2.3 Create Session Fails with Error while using Rest API End Points

Issue

While using any REST API client which saves cookies, a first time REST API call to create user is successful, but the subsequent request to create user fails.

Workaround

Remove the saved cookie from the client (Postman), while making the second API call.

1.1.2.3 Release 19.1.0.0.6_MLR33750853 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.6.3

1.1.2.3.1 Unable to Deploy more than 17 Pipelines

Issue

Unable to deploy more than 17 pipelines.

Workaround

The default Spark setup is set to a maximum of 16 pipelines. You can change this limit:

1. Open an SSH terminal on the GGSA instance
2. run `mysql -uroot -p`
3. Enter the password mentioned in `/home/opc/README.txt`
4. `mysql> use OSADB`
5. `mysql> insert into osa_system_property (mkey, value) values('spark.port.maxRetries', '40');`
6. Restart GGSA by running `sudo systemctl restart osa`

Also ensure that the environment has sufficient core and memory resources in Spark for the required number of pipelines.

1.1.2.4 Release 19.1.0.0.6_MLR33251626 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.6_MLR33251626

1.1.2.4.1 Export of a pipeline having a Custom CQL stage fails

Issue

When a pipeline with a custom CQL stage is exported from the Catalog, it fails.

Workaround

None.

1.1.2.4.2 Custom CQL Stage fails if Previous Stage Display Name contains Spaces

Issue

A custom CQL stage fails if the display name of the previous stage contains spaces.

Workaround

None.

1.1.2.4.3 Timestamp and Interval type are changed to Double and Text type in a Custom CQL Stage

Issue

Timestamp and Interval type fields in a CQL query of a custom CQL stage are changed to Double and Text type respectively.

Workaround

Add a query stage, and change the data types back to their original Timestamp and Interval data types.

1.1.2.5 Release 19.1.0.0.6_MLR32849357 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.6_MLR32849357.

1.1.2.5.1 Ignite Caching and Full Records disabled in the Marketplace Instance

Issue

The following features are disabled in the GGSA marketplace instance:

- Ignite caching
- Generate Full Records (while creating a GG Stream)

Workaround

None.

1.1.2.5.2 Enable Caching for DB Reference disabled on Marketplace

Issue

The **Enable Caching** option, while creating a DB reference, is not supported in the GGSA marketplace instance.

Workaround

None.

1.1.2.6 Release 19.1.0.0.6 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.6.

1.1.2.6.1 Issue with Multiple HDFS Targets in a Pipeline

Issue

GGSA does not support multiple HDFS targets (includes Azure Datalake , which is also an HDFS target type) in the same pipeline, if they are using different HDFS clusters.

Workaround

Create a separate pipeline for HDFS target type.

1.1.2.6.2 MONGODB Test Connection fails for SSL Authentication

Issue

When MongoDB connection is created with **Client URI** connection type and with **SSL** authentication mechanism, GGSA Jetty server does not update the certificate and password. The result of **Test Connection** is based on first certificate and password. Subsequent data is not updated. However, the pipeline deployment works fine.

Workaround

For a valid **Test Connection** results, restart the Jetty server, each time you test a MongoDB connection.

1.1.2.6.3 Kafka Topic Name cannot contain a Hyphen

Issue

GGSA allows special characters such as -, _, and . in a pipeline name and a stage name. But Kafka does not allow a hyphen in topic names.

Workaround

When a topic is created in Kafka, the hyphen (-) in the topic name is replaced with an underscore (_).

1.1.2.6.4 Pipeline Fails on Publishing with ElasticSearch target and SSL Authentication

Issue

An elastic search connection with SSL type authentication does not mandate username and password, along with trust store credentials. But GGSA uses GoldenGate Big Data (GGBD) API to communicate with Elastic Search. GGBD has implementation bug which mandates username and password along with trust store credentials, making it mandatory to provide both inputs.

Workaround

This is a known issue. When GGDB fixes this issue, user can choose to provide only the truststore credentials, or the username and password for SSL authentication type.

1.1.2.6.5 GoldenGate Change Data fails to start with a Single Character Target Trail

Issue

GoldenGate Change data fails to start if created with a single character target trail.

Workaround

Create a GoldenGate Change data with a two-character target trail name.

1.1.2.7 Release 19.1.0.0.5 - Known Issues

This section lists the known issues and workarounds for GoldenGate Stream Analytics release 19.1.0.0.5.

1.1.2.7.1 Updating Partitions in a Hive Target kills a Pipeline

Issue

When you update the partitions in an existing Hive Target, the pipeline gets killed and throws an error.

Workaround

When you update the partitions, ensure you update the other parameters such as tablename, HDFS Path, HDFS Schema File Path, in the **Target Details** wizard.

1.1.2.7.2 Custom Query in System Settings Fails unless a Synonym is created

Issue

You can use a custom query in **Oracle DB Reference and Target** under **SQL queries configuration** on the **System settings** page, to access a table to look up. But this causes a pipeline to fail.

Workaround

Create a synonym with the name. For example, to access a table `OSA2.REF_TBL` from user `OSA1`, create a synonym from user `OSA1` using

```
Create synonym REF_TBL for OSA2.REF_TBL
```

1.1.2.7.3 Jetty Server Restart required when the Protection Policy is updated

Issue

When you update the protection policy, from the **System Settings** page, the change is not applied unless you restart the jetty server.

Workaround

Restart the jetty server, if you update the hadoop protection policy on the **System Settings** page.

1.1.2.7.4 Ignite Target cannot be Edited

Issue

You cannot edit an Ignite target that is once created.

Workaround

This is not an issue. This restriction is to avoid corruption of cache data, by allowing only one target from GGSA platform to write to only one cache in the ignite server.

1.1.2.7.5 Using Keys in a Ignite Cache

Issue

Single field is allowed as key, while reading from an ignite reference. Hence the pipeline can have only one condition while using the ignite reference. In a pipeline you can use only one condition operator `equals`, with only one condition.

Workaround

None.

1.1.2.8 Release 19.1.0.0.3 - Known Issues

This section lists the known issues and workarounds for Oracle Stream Analytics release 19.1.0.0.3.

1.1.2.8.1 Duplicate Records seen in case of Pipeline having Multiple Targets

Issue

If you have multiple targets in a pipeline, and one of the targets fails, even before the failure is detected and the pipeline stopped, the data gets written into active targets. When you resolve the failed target and republish the pipeline, there will be duplicate records in the targets that were active.

Workaround

Create a different pipeline with only one target in which duplicate records are not accepted.

1.1.2.8.2 Connection to ADW/ATP fails if Password has Special Characters in it

Issue

If you try to connect to ADW and ATP, and the DB credential has a special character in it, then the test connection and the pipeline deployment will not work.

Workaround

Specify the password between double quotes. For example, "welcome?123".

1.1.2.8.3 Null Pointer Exception when a Reference Table with Interval Data Type is used

Issue

In a Reference table, columns with Year to Month Intervals with year precision greater than 2, are not supported.

Workaround

None.

1.1.2.8.4 Nullpointerexception seen in Driver during a Graceful Shutdown

Issue

If there is some exception with target or reference during graceful shutdown, undeployment of pipeline can take longer, and you will see NullPointerException in the driver logs. However, there will be no data loss.

Workaround

None.

1.1.2.8.5 Clicking on Sign In redirects to Login Page

Issue

At times, entering your OSA instance credentials and clicking **Sign In**, does not log you in, but redirects you back to the login page.

Workaround

Try one of the following options:

- Close the window and open the OSA instance in a new window.
- Clear the cache and cookies in the browser.
- Clear the JSESSIONID cookie.

To clear the cookies in a Chrome browser, click on the 3 dots at the top right corner. Select **More tools** -> **Developer tools** -> **Application Tab**. Under **Storage**, expand **Cookies**. Select your OSA URL and clear JSESSIONID cookie.

1.1.2.8.6 Test Connection for JNDI Type does not work on Jetty Server

Issue

If you create a connection of type JNDI and use the JNDI provider as Weblogic, then the **Test Connection** does not work.

Workaround

Ensure that the JNDI connection details are correct. The pipeline using the JNDI connection works as expected.

1.1.2.8.7 Republish Pipeline while Graceful Shutdown is in progress

Issue

You unpublish a pipeline and republish the same, and encounter the following exception:

```
OSA-01213 Pipeline deployment has failed. Creating container has failed.
```

Check if the pipeline with the same name is already running. At times, the graceful shutdown of pipeline takes time to complete, and it is possible that the unpublish action is still underway.

Workaround

You can wait for the pipeline to unpublish gracefully and then deploy the pipeline. You can check the status of the pipeline in the Cluster manager (YARN or Spark standalone) console.

1.1.2.8.8 Republish Pipeline with Reset Offset if Input Topic is Repartitioned

Issue

If you repartition an input topic, you cannot republish a pipeline and start from last processed offsets, because internally created offset topics should have the same number of partitions as the input topic.

Workaround

When you repartition an input topic, republish the pipeline with a reset offset. This will ensure that the earlier offset topic is deleted, and a new one with the correct partition count is created.

1.1.2.8.9 Timestamp with Timezone that has more than 6 Subsecond Precision is not supported

Issue

An input timestamp with a seconds precision greater than 6 digits, and also with a timezone, is not supported. For example, 2020-04-21 12:29:53.998877665 +5:30, would result in wrong date.

Workaround

Input a timestamp with the seconds precision less than 6 digits.

1.1.2.9 Release 19.1.0.0.2 - Known Issues

This section lists the known issues and workarounds for Oracle Stream Analytics release 19.1.0.0.2.

1.1.2.9.1 Truncate Table With Goldengate Stream

Bug Number: 30677198

Issue

Any event processing based on deleting a record, will fail for all `Truncate Table` and `Drop Table` operations. This is a known issue.

Workaround

None.

1.1.2.9.2 GG Change Data created from OSA not starting when source table has SDO_Geometry Column

Bug Number: 30709930

Issue

Creating and starting a GG Change Data fails, when the source table has column of type `SDO_GEOMETRY`.

Workaround

None.

1.1.2.9.3 OSA Support for CLOB And BLOB Columns

Bug Number: 30759814

Issue

OSA does not support a Database Reference table with a BLOB column type.

OSA supports a Database Reference table with a CLOB column type, but you cannot define any operations on the stream field which is mapped to that column.

Workaround

None.

1.1.2.9.4 UI allow no mapping for a DB table column with Not Null constraint

Bug Number: 30542553

Issue

While adding a Database target, the UI allows you to map to any field in the target DB table, which has a 'not null' constraint. But when you try to publish a pipeline with a 'not nullable' column mapped to a 'No Mapping' field, you will encounter an exception with error code `OSA-01213`.

This is a known issue.

Workaround

None.

1.1.2.10 Release 19.1.0.0.1 - Known Issues

This section lists the known issues and workarounds for Oracle Stream Analytics release 19.1.0.0.1.

1.1.2.10.1 Edit DB Target fails to detect the updated/added columns in a table

Bug Number: 30540965

Issue

If you add a database target with a predefined table and then edit the table manually, the target does not reflect the edited table. This is a known issue.

Workaround

If you alter the table used as target for a OSA pipeline, for the changes to reflect, you should delete the existing target in the pipeline and create the Database target again.

1.1.2.10.2 Adding DB Target Stage fails when underlying Table has a Blob column

Bug Number: 30542368

Impacted Platforms: All

If you add a database target mapped to a table having a 'BLOB' or 'Interval' data type columns, the target does not identify the columns of the two datatypes. This is a known issue.

Workaround

There is no workaround for this issue. You should use only the tables without blob and interval column types as target for their OSA pipelines.

1.1.2.10.3 Pipeline continues to process records even when target is not accessible

Bug Number: 30531844

Issue

pipeline with a Database target gets published and continues to process records even when the Database target is unavailable. This is a known issue.

Workaround

None

1.1.2.10.4 UI allow no mapping for a DB table column with Not Null constraint

Bug Number: 30542553

Issue

While adding a Database target, the UI allows you to map to any field in the target DB table, which has a 'not null' constraint. But when you try to publish a pipeline with a 'not nullable' column mapped to a 'No Mapping' field, you will encounter an exception with error code OSA-01213.

This is a known issue.

Workaround

None.

1.1.2.10.5 Unable to change DB target Property Mapping after Switching to a New Target

Bug Number: 30566018

Issue

You cannot change the target property mapping after changing the target to a newly created one.

Workaround

Delete the target stage and create new target stage pointing to the new target.

1.1.2.10.6 Timestamp seen in Live Data is different from Input Data

Bug Number: 30579122

Issue:

In a live data stream, you cannot properly parse the timestamps having more than 3 digits of precision.

Workaround

Use a timestamp format from the list of supported formats below:

11/21/2005 11:14:23.111 "MM/dd/yyyy HH:mm:ss.SSS"

11/21/2005 11:14:23.11 "MM/dd/yyyy HH:mm:ss.SS"

11/21/2005 11:14:23.1 "MM/dd/yyyy HH:mm:ss.S"

11/21/2005 11:14:23 "MM/dd/yyyy HH:mm:ss"

11/21/2005 11:14 "MM/dd/yyyy HH:mm"

11/21/2005 11:14 "MM/dd/yyyy HH"

11/21/2005 "MM/dd/yyyy"

11-21-2005 11:14:23.111 "MM-dd-yyyy HH:mm:ss.SSS"

11-21-2005 11:14:23.11 "MM-dd-yyyy HH:mm:ss.SS"

11-21-2005 11:14:23.1 "MM-dd-yyyy HH:mm:ss.S"

11-21-2005 11:14:23 "MM-dd-yyyy HH:mm:ss"

11-21-2005 11:14 "MM-dd-yyyy HH:mm"

11-21-2005 11 "MM-dd-yyyy HH"

11-21-2005 "MM-dd-yyyy"

15-DEC-01 11.14.14.111 AM"dd-MMM-yy hh.mm.ss.SSS"

15-DEC-01 11.14.14.11 "dd-MMM-yy hh.mm.ss.SS"
 15-DEC-01 11.14.14.1 "dd-MMM-yy hh.mm.ss.S"
 15-DEC-01 11.14.14 "dd-MMM-yy hh.mm.ss"
 15-DEC-01 11.14 "dd-MMM-yy hh.mm"
 15-DEC-01 11 "dd-MMM-yy hh"
 15-DEC-01 "dd-MMM-yy"
 15/DEC/01 "dd/MMM/yy"
 2013-10-5 15:16:0.756 "yyyy-MM-dd HH:mm:ss.SSS"
 2013-10-5 15.16.0.756 "yyyy-MM-dd HH.mm.ss.SSS"
 2013-10-5 15:16:0 "yyyy-MM-dd HH:mm:ss"
 2013-10-5 15.16.0 "yyyy-MM-dd HH.mm.ss"
 2013-10-5 15:16 "yyyy-MM-dd HH:mm"
 2013-10-5 15.16 "yyyy-MM-dd HH.mm"
 2013-10-5 15 "yyyy-MM-dd HH"
 2012-11-10 "yyyy-MM-dd"
 11:14:14 "HH:mm:ss"
 "yyyy-MM-dd'T'HH:mm:ss'.SSS"
 "yyyy-MM-dd'T'HH:mm:ss"

1.1.2.10.7 Log level Settings do not work with Open JDK

Bug Number: 30611896

Issue

You cannot set a Log level while publishing a pipeline, if you use Open JDK instead of Oracle JDK. The log level setting is default to INFO.

Workaround

None

1.1.2.11 Release 19.1.0.0.0 - Known Issues

This section lists the known issues and workarounds for Oracle Stream Analytics release 19.1.0.0.0.

1.1.2.11.1 Batch Delay Increases When HA is Enabled

Bug Number: 27754912

Impacted Platforms: All

While publishing a pipeline, if you enable HA you can see that the batch scheduling is delayed. This is a known issue.

Workaround

No workaround available

1.1.2.11.2 Values of Drop-down List not seen when the Live Output Table has a single row

Bug Number: 27473663

Impacted Platforms: All

Full values within the drop-down list are not seen when the live output table has a single row. This is a known issue.

Workaround

Pull the horizontal splitter up so that more of the table is visible.

1.1.2.11.3 SDO_GEOMETRY Not Supported in Coherence Cache

Issue

Bug Number: 27480822

Impacted Platforms: All

In coherence cache, data type conversion to `SDO_GEOMETRY` is not supported. This is a limitation.

Workaround

No workaround available

1.1.2.11.4 Pattern Cannot Detect Events Correctly

Issue

Bug Number: 25132943

Impacted Platforms: All

When you have configured the batch duration as 1 second for underlying Spark Streaming, the pattern match clause cannot detect events sequence correctly happening within an interval of less than or equal to 1 second. This is a limitation.

Workaround

No workaround available

1.1.2.11.5 Default Character Limit for Text Fields

Bug Number: 27701535

Impacted Platforms: All

There is a limitation on the size of text fields in pipelines: events with text fields exceeding 1024 characters will not be processed.

1.1.2.11.6 Limited Support for Predictive Models

The import and scoring of Predictive Models might contain undocumented limitations and you should use them as is.

1.1.2.11.7 Possible loss of events in HA scenario after Executor Crash

Bug Number: 29697956

Impacted Platforms: All

Events that arrive when the executor is up and running but the CQL engine is still not bootstrapped, will result in failed tasks. By default each task is attempted 3 times by Spark and subsequent to that it reports the task as failed. This is a known issue.

Workaround

No workaround available.

1.1.2.11.8 Correlation pattern does not support number data type

Bug Number: 29801715

Impacted Platforms: All

Pattern Correlation does not support number data type. So when choosing the parameter of type Number for correlation pattern, it throws exception in the spark log file. The workaround is to use data type either of int, float, double, bigint or apply toDouble function for the number data type and use that in correlation pattern.

Workaround

The workaround is to use data type either of int, float, double, bigint or apply toDouble function for the number data type and use that in correlation pattern.

1.1.2.11.9 Rest Submission Server for Spark disabled

Oracle Stream Analytics when configured on Spark standalone cluster, uses its Rest Submission server to deploy the pipeline. Spark 2.4.0 and higher versions have Rest Submission servers disabled by default.

Ensure or enable the Spark Rest Submission server of the standalone cluster. To enable the Spark Rest Submission server, set the property `spark.master.rest.enabled` to `True`, when starting the cluster. One way is to add it in Spark default config file, `${SPARK_HOME}/conf/spark-defaults.conf`, before starting the cluster. See [Spark documentation](#) for details.

1.1.2.11.10 Loading Class Fails when using JMS Streams

Bug: 29022250

Impacted Platforms: All

When using JMS stream with durable subscription, user is possible to run into the error: `FAILED TO LOAD CLASS WEBLOGIC.JMS.COMMON.INVALIDCLIENTIDEXCEPTION`. This is because the JMS durable subscription is exclusive and an existing pipeline has created the subscription.

Workaround

Change the stream's setting to use another client ID or subscription name, or kill the unwanted pipeline that runs JMS subscriber.

1.1.2.11.11 No Output Event For Some Special Rule Configuration

Bug: **29015898**

Impacted Platforms: All

In rule stage, modifying a value by a rule, already set by another rule, leads to no events to be emitted in live output table.

Workaround

Create a separate rule stage to handle this case.

1.1.2.11.12 Disorder Of Output Events With Pipeline Contains Obr For Executor Failover

Bug: **29112584**

Impacted Platforms: All

Ordering of events across different partitions is not guaranteed. However events are ordered within same partition.

Workaround

No workaround available.

1.1.2.11.13 Query Group ->Table Stage does not work on certain Window Types

Bug: **29795496**

Impacted Platforms: All

Query group->Table stage does not work on the 'constant value', 'row without slide' and 'now' window types. Adding a group based on above window types on the query group/table stage, renders the pipeline broken, with the following UI error: Failed to start pipeline. : Uncaught exception. : Syntax Error. Invalid usage of Reserved CQL Keyword: "on".

Workaround

Create another stage from query group and configure the 'Group' and 'Constant Value' window in the new stage.

1.1.3 Corrected Problems

This section lists all the corrected problems in a Release.

1.1.3.1 Release 19.1.0.0.8 - Corrected Problems

This section lists the issues that were fixed in GoldenGate Stream Analytics release 19.1.0.0.8.

1.1.3.1.1 Logging into a GGSA Marketplace Instance using SSH Tunneling does not redirect to the Index.html page

Bug 34529208

This issue is resolved in Release 19.1.0.0.8.

1.1.3.2 Release 19.1.0.0.5 - Corrected Problems

This section lists the issues that were fixed in GoldenGate Stream Analytics release 19.1.0.0.5.

1.1.3.2.1 Pipeline Continues to Process Records even when Target is not Accessible

Fix

The pipeline terminates if the database target is not accessible.

1.1.3.2.2 Possible loss of events in HA scenario after Executor Crash

This issue is resolved in the 19.1.0.0.5 release.

1.1.3.3 Release 19.1.0.0.3 - Corrected Problems

This section lists the issues that were fixed in Oracle Stream Analytics release 19.1.0.0.3.

1.1.3.3.1 Adding DB Target Stage fails when underlying Table has an Interval Type Column

Fix

The Database target table can have Interval type columns.

1.1.3.3.2 Timestamp seen in Live Data is different from Input Data

Bug Number: 30579122

Fix

You can input timestamps with seconds of more than 3-digit precision, however, the input timestamp is truncated to millisecond precision.

SSSSSSSSSS

1.1.3.4 Release 19.1.0.0.2 - Corrected Problems

This section lists the issues that were fixed in Oracle Stream Analytics release 19.1.0.0.2.

1.1.3.4.1 Adding DB Target Stage fails when underlying Table has a Blob column

Bug Number: 30542368

Fix

The Database target table can have CLOB and BLOB columns. You can either map text fields or specify no mapping for these columns.

1.1.3.5 Release 19.1.0.0.1 - Corrected Problems

This section lists the issues that were fixed in Oracle Stream Analytics release 19.1.0.0.1.

1.1.3.5.1 Full Values of Dropdown list Cannot be seen if Live Out Stream table has One Row

Bug Number: 27473663

Fix

You can now see the scroll bar with all the right-click menu options for columns in a Live Output table.

1.1.4 Bug Fixes and Enhancements

This chapter describes the bug fixes and enhancements for a specific release.

The Bug number is the number of the BugDB ticket. For questions on specific tickets or issues, consult Oracle Support.

1.1.4.1 Release 19.1.0.07_MLR36064807 - December 2023

THROUGHPUT EFFICIENCY IN CQL ENGINE

While running a multi-query stage pipeline on multiple executors, no linear throughput efficiency increase, with increase in source partitions and executors aligned with the partition count.

This issue is fixed.

1.1.4.2 Release 19.1.0.07_MLR35766877 - September 2023

1.1.4.2.1 Kafka SASL Support for SCRAM-SHA-256 and SCRAM-SHA-512

Kafka connection using SASL supports the following Security settings:

- **Security Protocol:** SASL_PLAINTEXT
- **Security Mechanisms:**
 - SCRAM-SHA-256
 - SCRAM-SHA-512

1.1.4.3 Release 19.1.0.0.6_MLR33750853 - January 2022

Bug 33675135 - FIX THE KAFKA BOOTSTRAP PORT THAT IS PRE-SEEDED TO DB IN THE MARKETPLACE INSTANCE

Issue

The `init-web-tier.sh`, a utility that is used while creating a production Kafka cluster, had defaulted the Kafka broker port to 2181 instead of 9092.

Fix

Update the port from 2181 to 9092.

Bug 33388524: WRITING TO NOTIFICATION TARGET (ONS), ON THE MARKETPLACE INSTANCE FAILS DUE TO LIBRARY VERSION CONFLICT ERROR

This issue is fixed.

Bug 33617140: CANNOT EXPORT PIPELINES WITH CQL STAGE

Fixed the issue where exporting pipelines with a custom CQL stage was failing.

Bug 33638523: UNABLE TO RESET CONTENT TYPE USING CUSTOM HEADER IN REST TARGET

The fix is to set the default content type before custom headers are set.

Bug 33688240: CANNOT PUBLISH A PIPELINE WHICH HAS A KAFKA SOURCE WITH A SPACE IN ITS DISPLAY NAME

This issue is fixed. You can publish a pipeline which has a Kafka source with space in its display name.

Bug 33692708: EXCEPTION WHILE WRITING TO OCI OBJECT STORAGE TARGET WITH JSON_ROW FORMAT

Fix is to disable the `includeColumnNames` and `disableescaping` fields.

Bug 33728331: CVE-2021-44832: APACHE LOG4J UPDATE TO AT LEAST 2.17.1

Issue

CVE-2021-44832

Fix

Upgrade Apache Log4j to 2.17.1.

Bug 33592687: PIPELINE GETS TERMINATED WITHOUT LOGGING ANY EXCEPTION OR ERROR MESSAGE IN GGSA LOGS.

The fix is to revert logging from `Log4j 2` to `Log4j 1.2.17-16`, in the GGSA Runtime.

1.1.4.4 Release 19.1.0.0.6_MLR33251626 - September 2021

Bug 33172932: MISSING TUPLES WHEN PIPELINE DEPLOYED ON MULTIPLE EXECUTORS

Fixed the missing tuples issue, when a pipeline is deployed on multiple executors.

Bug 33193157: DRUID CUBE CREATION FAILS WHEN KAFKA TARGET USES A KAFKA CONNECTION CONFIGURED WITH BOOTSTRAP SERVER

Fixed the Druid cube creation failure issue, when a Kafka target uses a Bootstrap Kafka connection.

Bug 32964239: DIFFERENTIATE PIPELINE TERMINATION BETWEEN SOFT EXCEPTION AND RUNTIME EXCEPTION

Issue

The pipeline terminates due to incompatible input data

Fix

Based on your requirement, you can terminate the pipeline, or continue by dropping the erroneous data. Update the `osa.continueon.data.exception` flag in the

`OSA_SYSTEM_PROPERTY` table:

- If the flag is set to true, the pipeline will continue to run, dropping the faulty event.
- If the flag is set to false, pipeline will terminate.
- By default the value is set to true.

Bug 32964187: CONFIGURATION PARAMETER ISSUE TO ALLOW PIPELINE TO CONTINUE OR TERMINATE ON EXCEPTION

Issue

The pipeline terminates due to incompatible input data

Fix

Based on your requirement, you can terminate the pipeline, or continue by dropping the erroneous data. Update the `osa.continueon.data.exception` flag in the

`OSA_SYSTEM_PROPERTY` table:

- If the flag is set to true, the pipeline will continue to run, dropping the faulty event.
- If the flag is set to false, pipeline will terminate.
- By default the value is set to true.

Bug 32990445: IGNITE SERVER IS CALCULATING WRONG NUMBER OF WORKER NODES.

Fixed the ignite server calculating wrong number of worker nodes issue.

Bug 33208855: CANNOT REPUBLISH A PIPELINE WITH OFFSET THAT WAS KILLED DUE TO INPUT DATA EXCEPTION

Fixed issue related to republishing a pipeline that has an erroneous data at the beginning of the offset.

Bug 33236693: ISSUE WITH SYSTEM PROPERTY "SPARK.STREAMING.EXCEPTION.ISTERMINATE" SET TO FALSE

Fixed the issue in which setting `spark.streaming.exception.isTerminate` in the `OSA_SYSTEM_PROPERTY` table to `FALSE` does not keep the pipeline running when exceptions are encountered.

Bug 33272967: ENTIRE BATCH IS DROPPED WHEN THERE IS A FAULTY EVENT IN THE BATCH**Issue**

When `spark.streaming.exception.isTerminate` is set to `false`, and if any record failed to get inserted into the Oracle DB table, then the entire batch is skipped.

Fix

You can skip only the erroneous data that is failing insertion, by updating the `osa.jdbc.switchtononbatch.insert` to `true`, in the `OSA_SYSTEM_PROPERTY` table.

1.1.4.5 Release 19.1.0.0.6_MLR32849357 - June 2021

Bug 32845653: ENHANCE TOJSON PATTERN WITH OUTPUT JSON

Enhanced the ToJson pattern to upload a sample JSON file which will be used to infer JSON path for field mapping.

Bug 32647810: INCREASE THE STRING LENGTH IN CQENGINE

Fixed the CQEngine restriction of 1024 characters for attribute of type string. Increased the attribute length to allow 4096 characters.

Bug 32820171: NULL VALUES SEEN ON SELECTING ENABLE BATCHING IN TOJSON PATTERN

Fixed an issue in which the new event values showed null in the Live Output table, on clicking the Enable Batching option after adding a ToJson pattern to an event stream.

Bug 32845715: FIX TO IGNORE DELIMITED COLUMN HEADER AND DISABLE CHARACTER ESCAPING

Enabled an option to ignore the column header, while writing the delimited text using GGBD handler. This fix also disables escaping of characters while writing JSON string as a delimited text.

Bug 32846444: MAKE JDBC Sink QUERY TIMEOUT CONFIGURABLE

Set the JDBC sink query timeout as a configurable parameter from the `osa_system_property` table. The property that needs to be inserted into the `osa_system_property` table is `mkey=com.oracle.cep.adapters.jdbc.JdbcSink.QueryTimeout`.

Bug 32850867: OSA INSTALLER FAILS WITH ORACLE DB METADATA STORE

Fixed the issue where the metadata store user creation fails when the password has a special character like '@', while configuring Oracle metadata store. The metadata store password is the password that is configured in the `jetty-osa-datasource.xml` file.

Bug 32870972: WRITING TO KAFKA TARGET FAILS IF THE FIELD_PATH AND FIELD_NAME ARE DIFFERENT

Fixed the writing to Kafka target failure issue, when the `field_path` and `field_name` are different.

Bug 32872055: NULLPOINTEREXCEPTION AND NOSUCHELEMENTEXCEPTION WHILE COMMITTING THE OFFSET TO OSA INTERNAL TOPIC

Fixed the `NullPointerException` and `NoSuchElementException` issue arising while committing the offset to OSA internal topic.

Bug 32853397: STREAMING ANALYTICS FAILED TO START PIPELINE. : UNCAUGHT EXCEPTION. : 1024

Fixed the `Failed to start pipeline. : Uncaught exception. : 1024` error when configuring a pipeline.

Bug 32894915: KAFKA GROUPAUTHORIZATIONEXCEPTION: NOT AUTHORIZED TO ACCESS GROUP

Fixed the user authorization to access group issue with the error `GroupAuthorizationException: Not authorized to access group`.

Bug 32887340: ALLOW COLUMNS TO BE DROPPED IN THE DATABASE TARGET TYPE

Fixed the UI to allow columns to be dropped in the Database target type. At runtime the insert statement only specifies columns in the shape.

Bug 32754024: NO SPACE LEFT ON DEVICE- CHANGE SPARK WORKING DIRECTORY

The fix is to ensure that the correct pipeline status is reflected after changing the Spark working directory.

To change the Spark working directory:

1. Stop Spark service:
 - a. `sudo systemctl stop spark-slave.service`
 - b. `sudo systemctl stop spark-master.service`
2. Create work directory under `u02`:
 - a. Navigate to `/u02`
 - b. `sudo mkdir spark`. Here `spark` is the work folder name, you can create folder with the name of your choice.
 - c. `chmod 777 spark`, to change permission.
3. Edit `spark-env.sh`
 - a. Navigate to `SPARK_HOME/conf`, and edit `spark-env.sh`
 - b. Add `SPARK_WORKER_DIR=/u02/spark` at the end of the file `spark-env.sh` to point to newly created folder under `/u02`
4. Start Spark service:
 - a. `sudo systemctl start spark-master.service`
 - b. `sudo systemctl start spark-slave.service`

You will see the application and driver data (files and logs) under `/u02/spark` when you publish the pipeline again.

Bug 32873823: OSA INSTALLER FAILS WITH ADW METADATA STORE WITH ERROR CODE ORA-64308

Fixed the issue where the OSA installation fails with ADW metadata configuration, while running the sql script `seedSchema-19.1.0.0.6.sql`, with error **ORA-64308: hybrid columnar compressed table cannot have column with LONG data type**.

1.1.4.6 Release 19.1.0.0.5_MLR32849350 - May 2021

Bug 32251429: ISSUE WITH PIPELINE DEPLOYMENT WITH MULTIPLE TARGETS

Fixed the pipeline deployment failure issue for pipelines with multiple targets.

Bug 32251367: HOT DEPLOY ISSUE WITH SSL ENABLED KAFKA SETTINGS

Fixed the issue where the hot deployment of pipeline fails with `SparkException`, when SSL is enabled for Kafka system settings.

Bug 32613021: PIPELINE TERMINATES WHEN NESTED JSON HAS NULL VALUE

Fixed the terminating of pipelines when the input nested JSON data has a null value.

Bug 32647925: PERFORMANCE ISSUE WITH DB TARGET AND REFERENCE

Enhanced the validation performance while writing to a Database Target and Reference.

Bug 32647546: WRITING TO OCI TARGET FAILED WITH JSON WITH NULL VALUES

Fixed issue where a pipeline fails with `NullPointerException` exception while writing to an OCI target, when input JSON data has null values.

Bug 32712524: PUBLISHED PIPELINE GETS KILLED WHILE WRITING TO REST ENDPOINT IF RECORD HAS DIACRITIC CHARACTERS

Fixed the issue where the published pipeline is killed while writing to a REST target, when data in input JSON contains a string value with characters with diacritics: `ěčřžýáíé` (valid UTF-8 characters).

Bug 32375359 - BIG DECIMAL ROUNDING NECESSARY EXCEPTION

Fixed the issue where a pipeline fails with an `NullPointerException`, when the input JSON has a big decimal value with scale greater than 56. Now the default supported precision and scale is 64.

Bug 32669536: INCORRECT QUERY GENERATED FOR CURRENT AND PREVIOUS PATTERN

Fixed the issue where an incorrect query is generated when more than one partition field is selected in the `Current and Previous Events` pattern.

Bug 32846444: MAKE JDBC Sink QUERY TIMEOUT CONFIGURABLE

Set the JDBC sink query timeout as a configurable parameter from the `osa_system_property` table. The property that needs to be inserted into the `osa_system_property` table is `mkey=com.oracle.cep.adapters.jdbc.JdbcSink.QueryTimeout`.

Bug 32850867: OSA INSTALLER FAILS WITH ORACLE DB METADATA STORE

Fixed the issue where the metadata store user creation fails when the password has a special character like '@', while configuring Oracle metadata store. The metadata store password is the password that is configured in the `jetty-osa-datasource.xml` file.

Bug 32872055: NULLPOINTEREXCEPTION AND NOSUCHELEMENTEXCEPTION WHILE COMMITTING THE OFFSET TO OSA INTERNAL TOPIC

Fixed the **NullPointerException** and **NoSuchElementException** issue arising while committing the offset to OSA internal topic.

Bug 32887340: ALLOW COLUMNS TO BE DROPPED IN THE DATABASE TARGET TYPE

Fixed the UI to allow columns to be dropped in the Database target type. At runtime the insert statement only specifies columns in the shape.

Bug 32754024: NO SPACE LEFT ON DEVICE- CHANGE SPARK WORKING DIRECTORY

The fix is to ensure that the correct pipeline status is reflected after changing the Spark working directory.

To change the Spark working directory:

1. Stop Spark service:
 - a. `sudo systemctl stop spark-slave.service`
 - b. `sudo systemctl stop spark-master.service`
2. Create work directory under u02:
 - a. Navigate to `/u02`
 - b. `sudo mkdir spark`. Here `spark` is the work folder name, you can create folder with the name of your choice.
 - c. `chmod 777 spark`, to change permission.
3. Edit `spark-env.sh`
 - a. Navigate to `SPARK_HOME/conf`, and edit `spark-env.sh`
 - b. Add `SPARK_WORKER_DIR=/u02/spark` at the end of the file `spark-env.sh` to point to newly created folder under `/u02`
4. Start Spark service:
 - a. `sudo systemctl start spark-master.service`
 - b. `sudo systemctl start spark-slave.service`

You will see the application and driver data (files and logs) under `/u02/spark` when you publish the pipeline again.