

Cloudera DataFlow for Data Hub Release Notes

Date published: 2019-12-16

Date modified: 2021-09-09



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

What's New in Cloudera DataFlow for Data Hub 7.2.11.....	4
What's New in Flow Management.....	4
What's New in Streams Messaging.....	5
What's New in Streaming Analytics.....	7
Component Support in Cloudera DataFlow for Data Hub 7.2.11.....	7
Supported NiFi Extensions.....	7
Supported NiFi Processors.....	7
Supported NiFi Controller Services.....	9
Supported NiFi Reporting Tasks.....	11
Unsupported Features in Cloudera DataFlow for Data Hub 7.2.11.....	11
Unsupported Flow Management features.....	11
Unsupported Streams Messaging features.....	12
Unsupported Streaming Analytics features.....	12
Known Issues In Cloudera DataFlow for Data Hub 7.2.11.....	13
Known Issues in Flow Management.....	13
Known Issues in Streams Messaging.....	14
Known Issues in Streaming Analytics.....	19
Fixed Issues in Cloudera DataFlow for Data Hub 7.2.11.....	20
Fixed Issues in Flow Management.....	20
Fixed Issues in Streams Messaging.....	23
Fixed Issues in Streaming Analytics.....	25

What's New in Cloudera DataFlow for Data Hub 7.2.11

Cloudera DataFlow for Data Hub 7.2.11 includes components for Flow Management, Streaming Analytics, and Streams Messaging. Learn about the new features and improvements in each of these components.

What's New in Flow Management

Learn about the new Flow Management features in Cloudera DataFlow for Data Hub 7.2.11.

The Flow Management clusters in CDP Public Cloud introduces the following features:

ASN1 Reader

A new Controller Service uses the record API while reading ASN1 data.

Consume Kinesis Stream processor

New processor to consume data from AWS Kinesis Stream.

Decrypt and Encrypt Content PGP processors

New processors to decrypt and encrypt data using PGP.

NAR hot loading from object stores

Custom NARs can be loaded from object stores such as HDFS, S3, ADLS, GCS, and similar where NiFi is running. NAR files can be dropped into a single location and will be automatically retrieved and loaded by NiFi across all nodes.

For more information, see *Hot Loading Custom NARs*.

New Cloudera Kafka processor

A new Kafka processor is now available to support the Kafka client version provided in CDP Public Cloud 7.2.11 and to provide out-of-the-box integration with SMM by configuring the metrics interceptors.

For more information, see *Ingesting Data into Apache Kafka in CDP Public Cloud*.

Scale up and down NiFi clusters

You can now scale up and down your NiFi clusters to match your resources needs.

For more information, see *Scaling up or down a NiFi cluster*.

Object Store Processors

The following new processors are now available to interact with the object store of the underlying cloud provider in RAZ-enabled CDP environments:

- List
- Fetch
- Put
- Delete

The processors also integrate with RAZ which enables you to define Object Store access policies directly in Ranger.

RAZ is currently available for CDP Public Cloud environments on AWS and Azure.

Related Information

[Hot Loading Custom NARs](#)

[Ingesting Data into Apache Kafka in CDP Public Cloud](#)

[Scaling up or down a NiFi cluster](#)

What's New in Streams Messaging

Learn about the new Streams Messaging features in Cloudera DataFlow for Data Hub 7.2.11.

Kafka

New performance related health tests for Kafka

Two new health tests based on the `kafka_request_handler_avg_idle_1min_rate` and the `kafka_network_processor_avg_idle` metrics are added for Kafka in Cloudera Manager. The health tests added are the following the following:

- Request Handler Capacity

This health test checks the most recent value of the `kafka_request_handler_avg_idle_1min_rate` metric and sends a warning if less than 30% of request handler capacity is available. Additionally, the warning recommends that users increase the number of I/O threads using the `Number of I/O Threads (num.io.threads)` property.

- Network Processor Capacity

This health test checks the most recent value of the `kafka_network_processor_avg_idle` metric and sends a warning if less than 30% of network processor capacity is available. Additionally, the warning recommends that users increase the number of network threads using the `Number of Network Threads (num.network.threads)` property.

For more information, see [Cloudera Manager Health Tests Reference](#).

Schema Registry

Schema Registry Confluent API needs a compatibility endpoint

A new endpoint that checks compatibility between schemas is now available to the Confluent-compatible API.

The endpoint is as follows:

```
/api/v1/confluent/compatibility/subjects/[***SCHEMA**]/versions/[***VERSION***]
```

Where:

- [***SCHEMA**] is the name of the subject/schema.
- [***VERSION***] is the version that the new schema text is compared to. The version can be latest or a valid version ID.

When using the endpoint, you must specify the schema, version, and the schema text that you want to compare. The schema text can be sent as a data parameter. For example:

```
curl -X POST "http://\hostname\':9090/api/v1/confluent/compatibility/subjects/schemaname/versions/latest" -H "accept: application/json" -H "Content-Type: application/json" -d '{"schema\":"{\\"fields\\":[{\\"default\\":\\"yellow\\",\\"name\\":\\"color\\",\\"type\\":\\"string\\"}],\\"name\\":\\"schemaname\\",\\"type\\":\\"record\\"}]'}
```

The endpoint responds with a compatibility result. This can either be true or false.

If the schema is compatible, the endpoint returns a true response. For example:

```
{"compatible":true,"errorMessage":null,"errorLocation":null,"schema":{"type\":"record","name\":"compatible","fields":{"
```

```
\ "name\":"compatible\","type\":[\ "string\"],\ "default\":"test\
\"}]"} }
```

If the schema is not compatible, the endpoint returns a false response. This response contains an error message as well as the location where the error was encountered. For example:

```
{"compatible":false,"errorMessage":"reader union lacking writer
type: STRING","errorLocation":"/fields/0/type/0","schema":{"\ "ty
pe\":"record\","name\":"compatible\","fields\":[{\ "name\":"
compatible\","type\":[\ "string\"],\ "default\":"test\"}]"} }
```

Streams Messaging Manager

The SMM API now hides email notifier SMTP passwords in its response

Previously, the /notifiers endpoint returned the full configuration of the notifier. In the case of email notifiers, the configuration included the password of the SMTP server. API responses from now on do not include the password. As result of this change, the PASSWORD field of existing email notifiers is left blank when you edit them. If you decide to edit the notifier you must reenter the password.

Streams Replication Manager

The SRM Service role tries to recover automatically if errors are encountered

The SRM Service role might encounter errors that make metrics processing impossible. An example of this is when the target Kafka cluster is not reachable. If such an error is encountered, the SRM Service role now tries to recover automatically. If recovery is successful, the SRM Service role continues to monitor replications and displays as healthy in Cloudera Manager. However, during recovery, until the recovery is successful, the role displays as unhealthy.

New health tests for the SRM Service role

New health tests are introduced for the SRM Service role. These health tests describe the state of the SRM Service role. If the SRM Service role encounters an error that makes metrics processing impossible, Cloudera Manager now correctly displays the SRM Service role as unhealthy.

The wait time before starting new connectors is now configurable

A new configuration property, connect.start.task.timeout.ms, is added. This property controls the timeout of the tasks executed when starting connectors. The default value of the property is 20000 ms. You can configure the property on a replication level through the Streams Replication Manager's Replication Configs Cloudera Manager property. For example:

```
[ ***ALIAS*** ]->[ ***ALIAS*** ].connect.start.task.timeout.ms=25000
```

Custom lists of supported/excluded cipher suites and TLS/SSL protocols are configurable for the SRM Service role

A number of new properties related to TLS/SSL are introduced for the SRM Service role. These properties allow users to customize which cipher suites and TLS/SSL protocols should be supported or excluded by the SRM Service role. The properties added are as follows:

- Supported SSL/TLS Cipher Suites
(streams.replication.manager.ssl.supportedCipherSuites)
- Excluded SSL/TLS Cipher Suites
(streams.replication.manager.ssl.excludedCipherSuites)
- Supported SSL/TLS Protocols
(streams.replication.manager.ssl.supportedProtocols)
- Excluded SSL/TLS Protocols
(streams.replication.manager.ssl.excludedProtocols)

What's New in Streaming Analytics

Learn about the new Streaming Analytics features in Cloudera DataFlow for Data Hub 7.2.11.

The following new features are introduced in Streaming Analytics CDF for Data Hub 7.2.11:

- SQL Stream Builder support
- Google Cloud Platform (GCP) support

Component Support in Cloudera DataFlow for Data Hub 7.2.11

Cloudera DataFlow for Data Hub 7.2.11 includes the following components.

Flow Management clusters

- Apache NiFi 1.13.2
- Apache NiFi Registry 0.8.0

Streams Messaging clusters

- Apache Kafka 2.5.0
- Schema Registry 0.10.0
- Streams Messaging Manager 2.1.0
- Streams Replication Manager 1.0.0

Streaming Analytics clusters

- Apache Flink 1.12

Supported NiFi Extensions

Apache NiFi 1.13.2 ships with a set of Processors, Controller Services, and Reporting Tasks, most of which are supported by Cloudera Support. Review the supported extensions and avoid using any unsupported extensions in your production environments.

Supported NiFi Processors

This release ships with Apache NiFi 1.13.2 and includes a set of Processors, most of which are supported by Cloudera Support. You should be familiar with the available supported Processors, and avoid using any unsupported Processors in production environments.

Additional Processors are developed and tested by the Cloudera community but are not officially supported by Cloudera. Processors are excluded for a variety of reasons, including insufficient reliability or incomplete test case coverage, declaration of non-production readiness by the community at large, and feature deviation from Cloudera best practices.

AttributesToCSV	GetAzureQueueStorage	PutCouchbaseKey 1
AttributesToJSON	GetCouchbaseKey 1	PutDatabaseRecord
Base64EncodeContent	GetFile	PutDistributedMapCache
CalculateRecordStats	GetFTP	PutDynamoDB 1
CaptureChangeMySQL	GetHBase	PutElasticsearch 1
CompressContent12	GetHDFS	PutElasticsearchHttp 1

ConnectWebSocket	GetHDFSFileInfo	PutElasticsearchHttpRecord
ConsumeAMQP	GetHDFSSequenceFile	PutElasticsearchRecord
ConsumeAzureEventHub	GetHTMLElement	PutEmail 1
ConsumeEWS	GetHTTP	PutFile
ConsumeGCPubSub	GetIgniteCache	PutFTP
ConsumeJMS	GetJMSQueue	PutGCSObject
ConsumeKafka	GetJMSTopic	PutGridFS
ConsumeKafka_0_10	GetKafka	PutHBaseCell 1
ConsumeKafka_1_0	GetMongoRecord	PutHBaseJSON
ConsumeKafka_2_0	GetSFTP	PutHBaseRecord
ConsumeKafka_2_6	GetSolr	PutHDFS
ConsumeKafka2CDP	GetSplunk	PutHive3QL
ConsumeKafka2RecordCDP	GetSQS	PutHive3Streaming
ConsumeKafkaRecord_0_10	GetTCP	PutHiveQL
ConsumeKafkaRecord_1_0	GetTwitter	PutHiveStreaming
ConsumeKafkaRecord_2_0	HandleHttpRequest	PutHTMLElement
ConsumeKafkaRecord_2_6	HandleHttpResponse	PutInfluxDB
ConsumeKinesisStream	HashAttribute	PutJMS
ConsumeMQTT 1	HashContent	PutKafka
ConsumeWindowsEventLog	IdentifyMimeType	PutKinesisFirehose
ControlRate	InvokeAWSGatewayApi	PutKinesisStream
ConvertAvroSchema	InvokeGRPC	PutKudu
ConvertAvroToJSON	InvokeHTTP	PutLambda
ConvertAvroToORC	InvokeScriptedProcessor	PutMongoRecord
ConvertAvroToParquet	JoltTransformJSON	PutORC
ConvertCharacterSet	JoltTransformRecord	PutParquet
ConvertCSVToAvro	JsonQueryElasticsearch	PutRecord
ConvertJSONToAvro	ListAzureBlobStorage	PutRiemann
ConvertJSONToSQL	ListAzureDataLakeStorage	PutS3Object
ConvertRecord	ListCDPObjectStore	PutSFTP
CreateHadoopSequenceFile	ListDatabaseTables	PutSNS
CryptographicHashAttribute	ListenFTP	PutSolrContentStream
CryptographicHashContent	ListenGRPC	PutSolrRecord
DecryptContentPGP	ListenHTTP	PutSplunk
DeleteAzureBlobStorage	ListenRELP	PutSplunkHTTP 1
DeleteAzureDataLakeStorage	ListenSyslog	PutSQL
DeleteByQueryElasticsearch	ListenTCP	PutSQS
DeleteCDPObjectStore	ListenTCPRecord	PutSyslog
DeleteDynamoDB	ListenUDP	PutTCP
DeleteGCSObject	ListenUDPRecord	PutUDP
DeleteGridFS	ListenWebSocket	PutWebSocket 1
DeleteHBaseCells	ListFile	QueryCassandra
DeleteHBaseRow	ListFTP	QueryDatabaseTable
DeleteHDFS	ListGCSBucket	QueryDatabaseTableRecord
DeleteS3Object	ListHDFS	QueryElasticsearchHttp
DeleteSQS	ListS3	QueryRecord
DetectDuplicate	ListSFTP	QuerySolr
DistributeLoad	LogAttribute	QuerySplunkIndexingStatus
DuplicateFlowFile	LogMessage	QueryWhois
EncryptContent 2	LookupAttribute	ReplaceText 1

EncryptContentPGP	LookupRecord	ReplaceTextWithMapping
EnforceOrder	MergeContent 1	ResizeImage
EvaluateJsonPath	MergeRecord	RetryFlowFile
EvaluateXPath	ModifyHTML element	RouteHL7
EvaluateXQuery	MonitorActivity	RouteOnAttribute
ExecuteGroovyScript	Notify	RouteOnContent
ExecuteInfluxDBQuery	ParseCEF	RouteText
ExecuteProcess	ParseEvtx	SampleRecord 1
ExecuteScript	ParseSyslog	ScanAccumulo
ExecuteSQL	PartitionRecord	ScanAttribute
ExecuteSQLRecord	PostHTTP	ScanContent
ExecuteStreamCommand	PublishAMQP 1	ScanHBase
ExtractAvroMetadata	PublishGCPubSub 1	ScriptedTransformRecord
ExtractGrok	PublishJMS 1	ScrollElasticsearchHttp
ExtractHL7Attributes	PublishKafka	SegmentContent
ExtractImageMetadata	PublishKafka_0_10	SelectHive3QL
ExtractText	PublishKafka_1_0	SelectHiveQL
FetchAzureBlobStorage	PublishKafka_2_0	SplitAvro 1
FetchAzureDataLakeStorage	PublishKafka_2_6	SplitContent 1
FetchCDPObjectStore	PublishKafka2CDP	SplitJson 1
FetchDistributedMapCache	PublishKafka2RecordCDP	SplitRecord
FetchElasticsearch	PublishKafkaRecord_0_10	SplitText 1
FetchElasticsearchHttp	PublishKafkaRecord_1_0	SplitXml 1
FetchFile	PublishKafkaRecord_2_0	TagS3Object
FetchFTP	PublishKafkaRecord_2_6	TailFile
FetchGCXObject	PublishMQTT 1	TransformXml
FetchGridFS	PutAccumuloRecord	UnpackContent
FetchHBaseRow	PutAzureBlobStorage	UpdateAttribute
FetchHDFS	PutAzureCosmosDBRecord 1	UpdateCounter
FetchParquet	PutAzureDataLakeStorage	UpdateHive3Table
FetchS3Object	PutAzureEventHub 1	UpdateHiveTable
FetchSFTP	PutAzureQueueStorage	UpdateRecord
FlattenJson	PutBigQueryBatch	ValidateCsv
ForkRecord	PutBigQueryStreaming 1	ValidateRecord
GenerateFlowFile	PutCassandraQL 1	ValidateXml
GenerateTableFetch	PutCassandraRecord	Wait
GeoEnrichIP	PutCDPObjectStore	YandexTranslate
GeoEnrichIPRecord	PutCloudWatchMetric	
GetAzureEventHub		

Footnotes

- 1 – indicates a memory intensive processor
- 2 – indicates a CPU intensive processor

Supported NiFi Controller Services

This release ships with Apache NiFi 1.13.2 and includes a set of Controller Services, most of which are supported by Cloudera Support. You should be familiar with the available supported Controller Services, and avoid using any unsupported Controller Services in production environments.

Additional Controller Services are developed and tested by the Cloudera community but are not officially supported by Cloudera. Controller Services are excluded for a variety of reasons, including insufficient reliability or incomplete test case coverage, declaration of non-production readiness by the community at large, and feature deviation from Cloudera best practices.

AccumuloService	HiveConnectionPool
ActionHandlerLookup	HortonworksSchemaRegistry
ADLSCredentialsControllerService	IPFIXReader
ADLSIDBrokerCloudCredentialsProviderControllerService	IPLookupService
AlertHandler	JASN1Reader
AvroReader	JMSConnectionFactoryProvider
AvroRecordSetWriter	JndiJmsConnectionFactoryProvider
AvroSchemaRegistry	JsonPathReader
AWSCredentialsProviderControllerService	JsonRecordSetWriter
AWSIDBrokerCloudCredentialsProviderControllerService	JsonTreeReader
AzureBlobIDBrokerCloudCredentialsProviderControllerService	KafkaRecordSink_1_0
AzureCosmosDBClientService	KafkaRecordSink_2_0
AzureStorageCredentialsControllerService	KafkaRecordSink_2_6
AzureStorageCredentialsControllerServiceLookup	KeytabCredentialsService
CassandraDistributedMapCache	KuduLookupService
CassandraSessionProvider	LoggingRecordSink
CouchbaseClusterService	LogHandler
CouchbaseKeyValueLookupService	MongoDBControllerService
CouchbaseMapCacheClient	MongoDBLookupService
CouchbaseRecordLookupService	ParquetReader
CSVReader	ParquetRecordSetWriter
CSVRecordLookupService	PrometheusRecordSink
CSVRecordSetWriter	ReaderLookup
DatabaseRecordLookupService	RecordSetWriterLookup
DatabaseRecordSink	RecordSinkHandler
DBCPCConnectionPool	RecordSinkServiceLookup
DBCPCConnectionPoolLookup	RedisConnectionPoolService
DistributedMapCacheClientService	RedisDistributedMapCacheClientService
DistributedMapCacheLookupService	RestLookupService
DistributedMapCacheServer	ScriptedActionHandler
DistributedSetCacheClientService	ScriptedLookupService
DistributedSetCacheServer	ScriptedReader
EasyRulesEngineProvider	ScriptedRecordSetWriter
EasyRulesEngineService	ScriptedRecordSink
ElasticSearchClientServiceImpl	ScriptedRulesEngine
ElasticSearchLookupService	SimpleDatabaseLookupService
ElasticSearchStringLookupService	SimpleKeyValueLookupService
EmbeddedHazelcastCacheManager	SimpleScriptedLookupService
ExpressionHandler	SiteToSiteReportingRecordSink
ExternalHazelcastCacheManager	StandardHttpContextMap
FreeFormTextRecordSetWriter	StandardPGPPrivateKeyService
GCPCredentialsControllerService	StandardPGPPublicKeyService
GrokReader	StandardProxyConfigurationService
HadoopDBCPCConnectionPool	StandardRestrictedSSLContextService
HazelcastMapCacheClient	StandardS3EncryptionService

HBase_1_1_2_ClientMapCacheService	StandardSSLContextService
HBase_1_1_2_ClientService	Syslog5424Reader
HBase_1_1_2_ListLookupService	SyslogReader
HBase_1_1_2_RecordLookupService	VolatileSchemaCache
HBase_2_ClientMapCacheService	WindowsEventLogReader
HBase_2_ClientService	XMLReader
HBase_2_RecordLookupService	XMLRecordSetWriter
Hive3ConnectionPool	

Supported NiFi Reporting Tasks

This release ships with Apache NiFi 1.13.2 and includes a set of Reporting Tasks, most of which are supported by Cloudera Support. You should be familiar with the available supported Reporting Tasks, and avoid using any unsupported Reporting Tasks in production environments.

Additional Reporting Tasks are developed and tested by the Cloudera community but are not officially supported by Cloudera. Reporting Tasks are excluded for a variety of reasons, including insufficient reliability or incomplete test case coverage, declaration of non-production readiness by the community at large, and feature deviation from Cloudera best practices. Do not use these features in your production environments.

- AmbariReportingTask
- ControllerStatusReportingTask
- MetricsEventReportingTask
- MonitorDiskUsage
- MonitorMemory
- PrometheusReportingTask
- QueryNiFiReportingTask
- ReportLineageToAtlas
- ScriptedReportingTask
- SiteToSiteBulletinReportingTask
- SiteToSiteMetricsReportingTask
- SiteToSiteProvenanceReportingTask
- SiteToSiteStatusReportingTask

Unsupported Features in Cloudera DataFlow for Data Hub 7.2.11

Some features exist within Cloudera DataFlow for Data Hub 7.2.11 components, but are not supported by Cloudera.

Unsupported Flow Management features

There are no unsupported Flow Management features in Cloudera DataFlow for Data Hub 7.2.11

NiFi

There are no updates for this release.

NiFi Registry

There are no updates for this release.

Related Information

[Cloudera Community Forum](#)

Unsupported Streams Messaging features

Some Streams Messaging features exist in Cloudera DataFlow for Data Hub 7.2.11, but are not supported by Cloudera.

Kafka

The following Kafka features are not ready for production deployment. Cloudera encourages you to explore these features in non-production environments and provide feedback on your experiences through the *Cloudera Community Forums*.

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- While Kafka Connect is available as part of Runtime, it is currently not supported in CDP Public Cloud. NiFi is a proven solution for batch and real time data loading that complement Kafka's message broker capability. For more information, see [Creating your first Flow Management cluster](#).
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.

Schema Registry

There are no updates for this release.

Streams Messaging Manager

There are no updates for this release.

Streams Replication Manager

There are no updates for this release.

Related Information

[Cloudera Community Forum](#)

[Creating your first Streams Messaging cluster](#)

Unsupported Streaming Analytics features

Some Streaming Analytic features exist in Cloudera DataFlow for Data Hub 7.2.11, but are not supported by Cloudera.

Flink

The following Flink features are not ready for production deployment. Cloudera encourages you to explore these features in non-production environments and provide feedback on your experiences through the *Cloudera Community Forums*.

- Apache Flink batch (DataSet) API
- GPU Resource Plugin
- Application Mode deployment
- SQL Client

- The following features are not supported in SQL and Table API:
 - HBase Table Connector
 - Old Planner
 - Non-windowed (unbounded) joins, distinct

SQL Stream Builder

- INSERT INTO statements are not supported for SQL Stream Builder. Cloudera recommends to use sink tables instead of them.

Related Information

[Cloudera Community Forum](#)

Known Issues In Cloudera DataFlow for Data Hub 7.2.11

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera DataFlow for Data Hub 7.2.11.

Known Issues in Flow Management

Learn about the known issues in Flow Management clusters, the impact or changes to the functionality, and the workaround.

NiFi cannot connect to NiFi Registry

By default, NiFi is configured with a NiFi Registry client to interact with the NiFi Registry instance. The URL used to configure the Registry client may not be correct depending on your deployment model for CDP Public Cloud. For example:

```
https://***gateway***/.../.../cdp-proxy/nifi-registry-app/nifi-registry/
```

If the URL is not correct, you may face "connect timed out" errors when interacting with NiFi Registry from the NiFi UI.

You can manually change the configuration of the client and provide the right FQDN of the management node of the DataHub cluster where the NiFi Registry instance is installed. To update the NiFi Registry client, go into the top right Actions menu, and select Controller Settings | Registry Clients. A correct URL will look similar to:

```
https://***management0***/.../cdp-proxy/nifi-registry-app/nifi-registry/
```

NIFI-9054: Calling Nifi Registry's createExtensionBundleVersion REST endpoint will cause a NullPointerException

The /buckets/{bucketId}/bundles/nifi-nar API in NiFi Registry may throw a NullPointerException.

If you are using this API, contact Cloudera for a Hotfix.

JDK versions mismatch

If doing a software only upgrade for your Flow Management DataHub clusters and if repairing one of the NiFi nodes after the upgrade, you may be in a situation where the JDK used by NiFi is not the same across the nodes. In such a case, this may cause issues in the NiFi UI and you may get an "Unexpected error" message.

Ensure that the same JDK is used across the NiFi nodes and if there is a JDK versions mismatch, manually upgrade the JDK to match the JDK version being installed on the node that has been repaired.

Failed to import XML templates through the NiFi UI

When you try to import an XML template through the NiFi UI, you get an Invalid CORS request error.

Export or import flow definitions as JSON files. To import a JSON flow definition, drag and drop a process group on the canvas and upload the JSON file.

Technical Service Bulletins

TSB 2022-580: NiFi Processors cannot write to content repository

If the content repository disk is filled more than 50% (or any other value that is set in `nifi.properties` for `nifi.content.repository.archive.max.usage.percentage`), and if there is no data in the content repository archive, the following warning message can be found in the logs: "Unable to write flowfile content to content repository container default due to archive file size constraints; waiting for archive cleanup". This would block the processors and no more data is processed.

This appears to only happen if there is already data in the content repository on startup that needs to be archived, or if the following message is logged: "Found unknown file XYZ in the File System Repository; archiving file".

Upstream JIRA

- [NIFI-10023](#)
- [NIFI-9993](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2022-580: NiFi Processors cannot write to content repository](#)

TSB 2022-589: CVE-2022-33140 Apache NiFi ShellUserGroupProvider Vulnerability

The optional `ShellUserGroupProvider` in Apache NiFi 1.10.0 to 1.16.2 and Apache NiFi Registry 0.6.0 to 1.16.2 does not neutralize arguments for group resolution commands, allowing injection of operating system commands on Linux and macOS platforms. The `ShellUserGroupProvider` is not included in the default configuration. Command injection requires `ShellUserGroupProvider` to be one of the enabled User Group Providers (UGP) in the Authorizers configuration. Command injection also requires an authenticated user with elevated privileges. Apache NiFi requires an authenticated user with authorization to modify access policies in order to execute the command. Apache NiFi Registry requires an authenticated user with authorization to read user groups in order to execute the command. The resolution removes command formatting based on user-provided arguments.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2022-589: CVE-2022-33140 Apache NiFi ShellUserGroupProvider Vulnerability](#)

Known Issues in Streams Messaging

Learn about the known issues in Streams Messaging clusters, the impact or changes to the functionality, and the workaround.

Kafka

Learn about the known issues and limitations in Kafka in this release:

Known Issues

Topics created with the kafka-topics tool are only accessible by the user who created them when the deprecated --zookeeper option is used

By default all created topics are secured. However, when topic creation and deletion is done with the kafka-topics tool using the --zookeeper option, the tool talks directly to Zookeeper. Because security is the responsibility of ZooKeeper authorization and authentication, Kafka cannot prevent users from making ZooKeeper changes. As a result, if the --zookeeper option is used, only the user who created the topic will be able to carry out administrative actions on it. In this scenario Kafka will not have permissions to perform tasks on topics created this way.

Use kafka-topics with the --bootstrap-server option that does not require direct access to Zookeeper.

Certain Kafka command line tools require direct access to Zookeeper

The following command line tools talk directly to ZooKeeper and therefore are not secured via Kafka:

- kafka-reassign-partitions

None

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Custom Kerberos principal names cannot be used for kerberized ZooKeeper and Kafka instances

When using ZooKeeper authentication and a custom Kerberos principal, Kerberos-enabled Kafka does not start. You must disable ZooKeeper authentication for Kafka or use the default Kerberos principals for ZooKeeper and Kafka.

None

KAFKA-2561: Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

Limitations**Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade**

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Configuration.
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Schema Registry

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

Streams Messaging Manager

Learn about the known issues in Streams Messaging Manager in this release.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not support the log type used by SMM UI.

View the SMM UI logs on the host.

OPSAPS-59828: SMM cannot connect to Schema Registry when TLS is enabled

When TLS is enabled, SMM by default cannot properly connect to Schema Registry. As a result, when viewing topics in the SMM Data Explorer with the deserializer key or value set to Avro, the following error messages are shown:

- Error deserializing key/value for partition [***PARTITION***] at offset [***OFFSET***]. If needed, please seek past the record to continue consumption.
- Failed to fetch value schema versions for topic : '[***TOPIC***]'.

In addition, the following certificate error will also be present in the SMM log:

- javax.net.ssl.SSLHandshakeException: PKIX path building failed:...

Additional security properties must be set for SMM.

1. In Cloudera Manager, select the SMM service.
2. Go to Configuration.
3. Find and configure the SMM_JMX_OPTS property.

Add the following JVM SSL properties:

- Djavax.net.ssl.trustStore=[***SMM TRUSTSTORE LOCATION***]
- Djavax.net.ssl.trustStorePassword=[***PASSWORD***]

Streams Replication Manager

Learn about the known issues and limitations in Streams Replication Manager in this release:

Known Issues

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-14019: SRM may automatically re-create deleted topics

If `auto.create.topics.enable` is enabled, deleted topics are automatically recreated on source clusters.

Prior to deletion, remove the topic from the topic whitelist with the `srm-control` tool. This prevents topics from being re-created.

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1][TOPIC2]
```

CDPD-60823: Configuring the SRM Client's secure storage is mandatory for unsecured environments

In an unsecured environment the `srm-control` tool should not need any additional configuration to run. However, due to an issue with the automatic generation of the default configuration, configuring the SRM Client's secure storage is mandatory for the `srm-control` tool. This is true even if none of the clusters that the tool connects to are secured.

If a secure storage is not configured, the tool will fail with the following `NullPointerException`:

```
java.lang.NullPointerException
at com.cloudera.dim.mirror.SecureConfigProvider.retrievePassword(
SecureConfigProvider.java:99)
at com.cloudera.dim.mirror.SecureConfigProvider.configure(SecureConfigProvider.java:113)
at org.apache.kafka.common.config.AbstractConfig.instantiateConfigProviders(AbstractConfig.java:533)
at org.apache.kafka.common.config.AbstractConfig.resolveConfigVariables(AbstractConfig.java:477)
```

```

at org.apache.kafka.common.config.AbstractConfig.<init>(AbstractConfig.java:107)
at org.apache.kafka.common.config.AbstractConfig.<init>(AbstractConfig.java:142)
at org.apache.kafka.connect.mirror.MirrorMakerConfig.<init>(MirrorMakerConfig.java:88)
at com.cloudera.dim.mirror.MirrorControlCommand$SourceTargetCommand.init(MirrorControlCommand.java:97)
at com.cloudera.dim.mirror.MirrorControlCommand.issueCommand(MirrorControlCommand.java:369)
at com.cloudera.dim.mirror.MirrorControlCommand.main(MirrorControlCommand.java:346)

```

Configure a secure storage password and set it as an environment variable in your CLI session before running the srm-control tool.

1. In Cloudera Manager, select the Streams Replication Manager service.
2. Go to Configuration.
3. Find and configure the SRM Client's Secure Storage Password property.

Take note of the password that you configure.

4. Click Save changes.
5. Restart the SRM service
6. SSH into one of the SRM hosts in your cluster.
7. Set the secure storage password as an environment variable.

```

export [***SECURE STORAGE ENV VAR***]="[***SECURE STORAGE PASSWORD***]"

```

Replace `[***SECURE STORAGE ENV VAR***]` with the name of the environment variable you specified in Environment Variable Holding SRM Client's Secure Storage Password. Replace `[***SRM SECURE STORAGE PASSWORD***]` with the password you specified in SRM Client's Secure Storage Password. For example:

```

export SECURESTOREPASS="mypassword"

```

OPSAPS-61001: Saving configuration changes for SRM is not possible

Cloudera Manager incorrectly labels the SRM Client's Secure Storage Password property as mandatory. Moreover, it does not offer this property for configuration when SRM is installed with the Add Service Wizard.

As a result, it is possible to install and start SRM without configuring this property. However, in a case like this, making changes to SRM's configuration is not possible until the SRM Client's Secure Storage Password property is set.

Configure the SRM Client's Secure Storage Password property.



Important: Once the SRM Client's Secure Storage Password property is configured, you must set the password configured with the property as an environment variable in your CLI session before running the srm-control tool. The tool will fail to run if the password is not set as an environment variable. For more information see [Configuring srm-control](#).

OPSAPS-61814: Using the service dependency method to configure Kerberos enabled co-located clusters is not supported

Using the Streams Replication Manager Co-located Kafka Cluster Alias property to auto-configure the connection to a Kerberos enabled co-located Kafka cluster is not supported. In a case like this, the generated JAAS configuration contains host-specific configuration. This causes SRM to fail to connect to the co-located Kafka cluster on other hosts.

Define your co-located Kafka clusters using Kafka credentials. For more information, see [Defining co-located Kafka clusters using Kafka credentials](#). Alternatively, use the Streams Replication Manager's Replication Configs property to configure the connection to the co-located Kafka clusters.

OPSAPS-62546: Kafka External Account SSL keypassword configuration is used incorrectly by SRM

When a Kafka External Account specifies a keystore that uses an SSL key password, SRM uses it as the `ssl.keystore.key` configuration. Due to using the incorrect `ssl.keystore.key` configuration, SRM will fail to load the keystore in certain cases.

Workaround: For the keystores used by the Kafka External Accounts, the SSL key password should match the SSL keystore password, and the SSL keystore key password should not be provided. Alternatively, you can use the legacy connection configurations based on the `streams.replication.manager.configs` to specify the SSL key password.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manger.

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, consumer failover operations for transactional topics is not possible.

Known Issues in Streaming Analytics

Learn about the known issues in Streaming Analytics clusters, the impact or changes to the functionality, and the workaround.

There are no known Streaming Analytics issues in Cloudera DataFlow for Data Hub 7.2.11.

SQL Stream Builder

There is an authentication issue when accessing Streaming SQL Console

When using environments with assigned public IP addresses, Knox cannot authenticate the domain to reach the Streaming SQL Console. This causes an error, and you cannot access the User Interface of SQL Stream Builder.

For more information and for a workaround to solve this issue, contact your Cloudera sales representative.

CSA-1232: Big numbers are incorrectly represented on the Streaming SQL Console UI

The issue impacts the following scenarios in Streaming SQL Console:

- When having integers bigger than 253-1 among your values, the Input transformations and User Defined Functions are considered unsafe and produce incorrect results as these numbers will lose precision during parsing.
- When having integers bigger than 253-1 among your values, sampling to the Streaming SQL Console UI produces incorrect results as these numbers will lose precision during parsing.

None

CSA-1454: Timezone settings can cause unexpected behavior in Kafka tables

You must consider the timezone settings of your environment when using timestamps in a Kafka table as it can affect the results of your query. When the timestamp in a query is identified with `from_unixtime`, it returns the results based on the timezone of the system. If the timezone is not set in UTC+0, the timestamp of the query results will shift in time and will not be correct.

Change your local timezone settings to UTC+0.

CSA-3742: Catalogs are not working due to expired Kerberos TGT

When SSB is running for a longer period of time than the lifetime of the Kerberos Ticket Granting Ticket (TGT), authentication with the catalog services will fail and the catalogs stop working.

None

Fixed Issues in Cloudera DataFlow for Data Hub 7.2.11

Fixed issues represent selected issues that were previously logged through Cloudera Support, but are addressed in the current release. These issues may have been reported in previous versions within the Known Issues section; meaning they were reported by customers or identified by Cloudera Quality Engineering team.

Review the list of issues that are resolved in Cloudera DataFlow for Data Hub 7.2.11.

Fixed Issues in Flow Management

Review the list of Flow Management issues that are resolved in Cloudera DataFlow for Data Hub 7.2.11.

NIFI-8957

Possibility to set a description when creating a bucket in the NiFi Registry UI.

NIFI-8942

NiFi Registry: flow description cannot be selected and copied on the UI.

NIFI-8937

Show component name and version in configure dialog's title bar.

NIFI-8928

Upgrade Jetty to 9.4.43.v20210629.

NIFI-8788

Upgraded dependencies and removed unnecessary log4j test dependency.

NIFI-8787

Wrapped `hdfs.exists()` call in `UGI.doAs()` in `GetHDFS` processor.

NIFI-8782

Added Rate-Limiting for Access Token Requests.

NIFI-8770

Use queue `drainTo()` on shutdown in `HandleHttpRequest`.

NIFI-8768

Incorrect Date Parsing from String in Record Readers.

NIFI-8764

Refactor UnpackContent to use Commons Compress and Zip4j.

NIFI-8762

ADLSCredentialControllerService does not support EL for Storage Account name.

NIFI-8759

ExecuteSQL and ExecuteSQLRecord unnecessarily fall back to default decimal scale.

NIFI-8756

Upgraded AngularJS to 1.8.2 and JQuery to 3.6.0.

NIFI-8748

PutKudu Incorrect Date Conversion from String.

NIFI-8737

Incorrect provenance data in HDFS processors with ADLS destination.

NIFI-8730

Invalidate instead of evaluating empty script in scripted components.

NIFI-8724

Bouncy Castle 1.69.

NIFI-8723

Jackson 2.12.3.

NIFI-8718

Apache Commons IO 2.10.0.

NIFI-8717

Refactoring PutHDFS processor.

NIFI-8708

Spring Framework 5.3.8 for extension components.

NIFI-8705

Jetty 9.4.42.

NIFI-8704

Spring Framework 5.3.8.

NIFI-8699

Lucene 8.8.2.

NIFI-8682

opencsv 5.4.

NIFI-8662

Failed to parse AWS region from VPCE endpoint URL in AbstractAWSProcessor.

NIFI-8661

Update Record Reader/Writer lookup services to not require specific attributes exist.

NIFI-8659

JoltTransformRecord should support transformation of one record to multiple output records.

NIFI-8658

Allow for filtering functions to be used as top-level functions for RecordPath.

NIFI-8656

Support expression language for SAS Token in the ADLS Gen2 processors.

NIFI-8642

Select the default Old Gen Memory Pool for Memory Reporting Task.

NIFI-8630

Upgraded javax.mail 1.4.7 to jakarta.mail 2.0.1 for PutEmail.

NIFI-8627

Apache Derby 10.14.2.0.

NIFI-8625

ExecuteScript processor always stuck after restart or multi thread.

NIFI-8614

Spring Framework 5 issue with Cluster Node Firewall Bean.

NIFI-8604

Apache Accumulo 2.0.1.

NIFI-8538

Apache Commons IO 2.8.0.

NIFI-8522, NIFI-8640

NiFi can duplicate controller services when generating templates.

NIFI-8519, NIFI-8736, NIFI-8735

HDFS support for Hot Loading.

NIFI-8515

Apache Tika 1.26.

NIFI-8513

Spring Framework 4.3.30.

NIFI-8502

Spring Framework 5.3.6.

NIFI-8485

Jetty 9.4.40.

NIFI-8474

Add new Replacement Strategy for variable substitution in ReplaceText.

NIFI-8439

Handle parquet INT96 timestamps as byte-arrays (instead of exception).

NIFI-8435

Improve PutKudu memory consumption.

NIFI-8433

Add decommission command to nifi.sh.

NIFI-8429

DBCPCConnectionPool leaks registered drivers.

NIFI-8419

NPE when updating parameter context in a secure instance/cluster.

NIFI-8390

Handle HBase namespaces in Atlas reporting task.

NIFI-8388

Hazelcast 4.2.

NIFI-8344, NIFI-8458

Improve TailFile to continue tailing a file for some time after it has been rolled over.

NIFI-8343

Solr 8.8.2.

NIFI-8330

JythonScriptEngineConfigurator needs to recompile on init().

NIFI-8329

Updated dependencies with no build failures.

NIFI-8320

Fetching wrong schema from PostgreSQL DB.

NIFI-8251

Add Encrypt and Decrypt PGP Processors and Services.

NIFI-6061

PutDatabaseRecord does not properly handle BLOB/CLOB fields.

NiFi Registry dependencies upgrade

spring.boot.version to 2.5.1, spring.version to 5.3.8, jersey.server.version to 2.33, jetty.version to 9.4.42.v20210604, jackson.version to 2.12.3, bouncycastle.version to 1.69.

NIFIREG-395

Implemented the ability to import and export versioned flows in NiFi Registry UI.

Fixed Issues in Streams Messaging

Review the list of Streams Messaging issues that are resolved in Cloudera DataFlow for Data Hub 7.2.11.

Kafka

There are no fixed Kafka issues in this release.

Schema Registry

CDPD-26382: Schema Registry Client does not clear resources for daemon thread for Kerberos Login

Closing SchemaRegistryClient did not stop the threads managing KerberosLogin, so an application could end up with multiple open threads. This has been fixed and calling close() will also stop the Kerberos thread.

CDPD-25995: Fingerprint hashes are not consistent for Avro schemas

Previously, if there is a default null value in the schema text, the schema's fingerprint will be different with every restart of Schema Registry. With this fix, the default null values in the schema text will remain the same null values independent of restarting the Schema Registry service.

CDPD-25905: _orderByFields and name parameters on the API should not be mandatory

In the previous release, the "name" field was mandatory in the /search/schemas API endpoint. The intent was to make the API better, because searching without a "name" parameter returns an incorrect result. However, this was interpreted as a breaking change in the API and had to be reverted to maintain backward compatibility. Now it is again possible to send requests without providing a "name" parameter, and again this will return an incorrect result.

CDPD-25614: Clearing the cache throws a NPE

Fixed issue where invoking schemaRegistryClient.deleteSchema() caused a NullPointerException.

CDPD-21913: Rename properties in Schema Registry yaml file, remove the dots

Schema Registry uses the Dropwizard framework which allows overriding configuration properties from the command line. Due to implementation specifics, until now some of the properties could not be overridden. This issue is now resolved.

Schema Registry's StorageManager's init method's properties parameter has type StorageProviderConfiguration.

In Schema Registry's configuration file, StorageProviderConfiguration properties can be:

- DbType
- queryTimeoutInSecs

Properties that can be dataSourceClassName, dataSourceUrl, dataSourceUser, dataSourcePassword and connectionProperties can be:

- oracleNetSslVersion
- oracleNetSslServerDnMatch
- trustStore
- trustStoreType
- keyStore
- keyStoreType

CDPD-21617: Migrate lettuce to lettuce-io in cache module

The following classes have been removed from the cache module:

- | | |
|----------------------------------|------------------------|
| • CacheServiceRegistry | • RedisStringsCache |
| • CacheServiceLocalRegistry | • RedisHashesCache |
| • RedisCacheServiceBuilder | • RedisAbstractCache |
| • RedisCacheService | • GuavaCache |
| • DataStoreBackedCacheService | • PhoenixDataStore |
| • CacheServiceJsonFactory | • DataStoreWriter |
| • CacheServiceId | • DataStoreReader |
| • CacheService | • AbstractDataStore |
| • CacheWriterSync | • ViewConfig |
| • CacheWriterAsync | • TypeConfig |
| • CacheWriter | • DataStoreConfig |
| • CacheLoaderSyncFactory | • ConnectionConfig |
| • CacheLoaderSync | • CachesConfig |
| • CacheLoaderFactory | • CacheEntry |
| • CacheLoaderCallback | • CacheConfig |
| • CacheLoaderAsyncFactory | • StaticFactory |
| • CacheLoaderAsync | • Factory |
| • CacheLoader | • DataStoreBackedCache |
| • RedisConnectionPoolFactory | • CacheException |
| • RedisConnectionFactory | • LoadableCache |
| • AbstractRedisConnectionFactory | • AbstractCache |

CDPD-17358: Add logging to Ranger filtering operations in Schema Registry

Logging added to show which schemas have been filtered out.

OPSAPS-59993: Disable admin port in Schema Registry

The following admin ports are no longer enabled in Schema Registry:

- Port 7791 on a secure cluster.

- Port 7789 on a unsecured cluster.

OPSAPS-59972: Disable the TRACE method on all HTTP ports

In Streams Messaging Manager and Schema Registry, the allowed HTTP methods are changed to GET, POST, PUT, DELETE, HEAD, OPTIONS.

OPSAPS-60458: Knox principal is not over-ridable in SMM and Schema Registry

Custom Knox principal can be set for Schema Registry and SMM by setting the `knox_principal_name` property in Schema Registry Server Advanced Configuration Snippet (Safety Valve) for `registry.yaml` or Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for `streams-messaging-manager.yaml`.

Streams Messaging Manager**CDPD-26633: SMM API exposes email notifier SMTP password**

The SMM API no longer returns the SMTP password in its responses to avoid sensitive information leak. This leaves the SMTP password field blank in the SMM UI when editing an email notifier provider. The password needs to be re-entered when any changes are made to its configurations.

Streams Replication Manager**CDPD-13864 and CDPD-15327: Replication stops after the network configuration of a source or target cluster is changed**

This issue is fixed.

OPSAPS-60601: The SRM client's secure storage might become corrupted if the JAAS Secret properties are used

JAAS secrets can be used in Kafka External Accounts.

OPSAPS-60601: Replication does not start when the target cluster of the replication is unsecured

Unsecure clusters can now be targeted by replications.

OPSAPS-60775: Kafka External Accounts configurations are not generated for the SRM Service

The SRM Service's configuration now contains the Kafka External Accounts configuration, enabling SRM Service to access Kafka clusters defined through External Accounts.

Fixed Issues in Streaming Analytics

Review the list of Streaming Analytics issues that are resolved in Cloudera DataFlow for Data Hub 7.2.11.

7.2.11.12**CSA-3742: Catalogs are not working due to expired Kerberos TGT**

The issue regarding the expired Kerberos Ticket Granting Ticket (TGT) and catalog authentication has been fixed.

7.2.11

There are no fixed issues for Streaming Analytics in Cloudera DataFlow for Data Hub 7.2.11.