Cloudera DataFlow

Release Notes

Date published: 2021-04-06
Date modified: 2024-12-03

https://docs.cloudera.com/

Legal Notice

Contents

What's new

Learn about major features, updates, and fixes for Cloudera DataFlow service.

February 05, 2025

Release 2.9.0-h5-b2 of Cloudera DataFlow fixes an issue where the Projects page and the Target Project dropdown in Flow Design displayed only 10 projects - even when more were available.

Fixed issues

Fixed paging of project listing

Cloudera Control Plane was affected by an issue where the Cloudera DataFlow Projects page displayed a maximum of 10 projects, even when more were available. This limitation also affected Flow Design when selecting a Target Project for new drafts.

- Environments running Cloudera DataFlow version 2.9.0-h4-b3 do not require this upgrade and will not have the upgrade banner.
- Environments running Cloudera DataFlow version 2.9.0-h3-b1 or lower, will need to upgrade to 2.9.0-h5-b2.

December 03, 2024

This hotfix release (2.9.0-h4-b3) of Cloudera DataFlow resolves an issue where environments with a high number of deployments could unpredictably fail to upgrade during a standard upgrade process.

Fixed issues

Fixed an issue, identified during testing, where environments with 50+ flow deployments sometimes encountered incomplete upgrade processes. This release ensures that upgrades can complete successfully, even for environments with a large number of deployments.

November 20, 2024

This release (2.9.0-h3-b1) of Cloudera DataFlow fixes a control plane issue that is not visible to Cloudera customers.

- Environments running Cloudera DataFlow version 2.9.0-h2-b1 do not require this upgrade and will not have the upgrade banner.
- Environments running Cloudera DataFlow version 2.9.0-h1-b1 or lower, will need to upgrade to 2.9.0-h3-b1.

Fixed issues

See 2.9.0 hotfix 2.

Older releases

Overview of new features, enhancements, and changed behavior introduced in earlier releases of Cloudera DataFlow.

November 15, 2024

This release (2.9.0-h2-b1) of Cloudera DataFlow fixes an issue preventing the running of test sessions, and an upgrade issue affecting single to Flexible Server migration in Azure environments.

Fixed issues

- Fixed an issue that was causing test sessions in Flow Designer to either take a long time to start, or not start at all.
- Fixed an issue that was causing the 2.9.0 upgrade to fail for some customers running in Azure Environments.

November 11, 2024

This release (2.9.0-h1-b1) of Cloudera DataFlow fixes issues where the version 2.9.0 control plane had problems interacting with older workloads.

Fixed issues

- Fixed an issue where the Resources view handled workloads with version lower than 2.9.0 incorrectly.
- Fixed an issue where Flow Designer became inaccessible in workloads with version lower than 2.9.0.
- Fixed an issue where publishing a flow to the Catalog in a workload with version lower than 2.9.0 failed.

October 31, 2024

This release (2.9.0-b383) of Cloudera DataFlow increases developer productivity through the introduction of Parameter Groups which can be shared between flow drafts. Developers can now also create NiFi 2.0 flows in the Designer leveraging new Cloudera exclusive processors for building RAG data pipelines. Deployments can now be configured with a Prometheus endpoint that allows scraping Apache NiFi metrics. Cloudera DataFlow's service and deployment events and alerts now support in-app and email notifications.

What's new

Latest NiFi version

Flow Deployments and Test Sessions now support the latest Apache NiFi 1.27 release.

Build NiFi 2.0 flows in Flow Designer [Technical Preview]

You can now select NiFi 2.0 when creating drafts and start test sessions including the ability to configure your test session to use your custom Python based processors.

New Cloudera-exclusive AI processors for NiFi 2 [Technical Preview]

You can now implement RAG pipelines by using new processors to parse, chunk and vectorize data, bringing context to their LLMs. The following processors are now available with NiFi 2:

- PartitionPdf, PartitionHtml, PartitionText, PartitionDocx, PartitionCsv
- ChunkData
- EmbedData
- InsertToMilvus, LexicalQueryMilvus, VectorQueryMilvus
- PutChroma, QueryChroma
- PutOpenSearchVector, QueryOpenSearchVector

Bedrock Parameter Groups

You can now centrally define and manage parameter groups in a workspace and re-use them for multiple drafts, eliminating tedious copy-and-pasting of parameters and their values.

New Resources page

Users can now easily view and manage all their workspace resources like deployments, drafts, parameter groups, inbound connections, custom NAR/Python configurations in a single place.

Notifications via App and Email for Cloudera DataFlow service and deployment events

You can now receive real-time notifications for all events related to a Cloudera DataFlow Service and its deployments through the Cloudera Management Console, under the **Notifications** tab, and through email.

For more information, see Setting up service and deployment notifications

NiFi metrics can now be exposed via a Prometheus endpoint

You can now configure deployments to expose NiFi metrics through a Prometheus endpoint. Once set up, you can configure your Prometheus instances to scrape these endpoints, consume relevant metrics and build custom dashboards.

For more information, see Configuring access for NiFi metrics scraping.

New ReadyFlows

- ADLS to Pinecone
- S3 to Pinecone
- ADLS to Milvus
- S3 to Milvus
- RAG Query Milvus

New Kubernetes version support

Cloudera DataFlow now supports EKS/AKS 1.29

Changes and improvements

- As part of the upgrade process to Cloudera DataFlow 2.9.0, the Azure Postgres database is migrated from a single server to a flexible server deployment.
- Improved asset handling for deployments makes deployment creation more robust in cases where many deployments are created at the same time.
- Kubernetes scale up events could result in the Cloudera DataFlow application container being rescheduled causing Cloudera DataFlow to become unavailable. Additional restrictions for rescheduling the Cloudera DataFlow application were added to avoid downtime.
- Dependencies have been updated to Java 21, Spring 6 and Spring Boot 3.

Fixed issues

- NiFi cluster failed to auto scale with a UDP inbound connection configured
- NiFi node failed to start up due to custom Kubernetes cluster domain name
- MiniFi logging failed to clean up a full content volume
- Vault failed to start up due to insufficient wait in its postStart script
- Auto scaling driven by flow metrics did not kick in

June 25, 2024

This release (2.8.0-h2-b2) of Cloudera DataFlow introduces bug fixes for RAZ-enabled CDP environments running Apache NiFi 1.25.

Fixed issues

- CDPDFX-8844 - RAZ is broken on AWS with NiFi 1.25 in Cloudera DataFlow

June 03, 2024

This release (2.8.0-h1-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces bug fixes which fix issues with upgrades and changing flow versions.

Fixed issues

- CDPDFX-8851 - fixes an issue where a CDF upgrade on Azure could fail due to MiNiFi pods getting stuck.
- CDPDFX-8826 - fixes an issue where a CDF upgrade could fail due to an existing deployment being in ‘FAILED_TO_IMPORT’ state.
- CDPDFX-8803 - fixes an issue where deployments failed due to too long HTTP timeouts during asset loader API calls.

- CDPDFX-8828 - fixes an issue where certain configuration files represented by the ‘CDP Environment’ parameter were uploaded multiple times when change flow version was used.
- CDPDFX-8760 - fixes an issue where CDF was sending pod deployment status even when the pod count was zero.

May 20, 2024

This release (2.8.0-b274) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces the ability to change the flow definition version of running deployments, a new Overview page with tutorials and shortcuts, the ability to create NiFi 2.0 deployments (Technical Preview), filtering ReadyFlows by different categories, new GenAI ReadyFlows, and supports new Kubernetes versions.

What's new

Latest NiFi version

Flow Deployments and Test Sessions now support the latest Apache NiFi 1.25 release.

NiFi 2.0 Tech Preview

You can now select NiFi 2.0 (based on upstream Apache NiFi M2 release with critical fixes from M3) when creating deployments including the ability to configure your deployment to use your custom Python based processors.

Change flow definition version of existing deployments

You can now change the version of your flow definition for existing deployments. This eliminates the need to recreate deployments whenever a new version of your flow is available. Depending on your needs, you can choose from three different strategies when changing the flow definition version. Learn more about changing flow definition versions.

New Overview Page

When navigating to CDF, users now start on the new **Overview** page. The **Overview** page helps new users getting started with guides and documentation, informs administrators about recent releases and offers shortcuts to power users.

Custom NiFi Node sizing (requires entitlement)

When creating a deployment, CDF now supports specifying custom core/memory settings. This feature requires an entitlement. Reach out to your Cloudera team to request access.

ReadyFlow Gallery filtering

The ReadyFlow Gallery now supports filtering available ReadyFlows by four different categories: Use case category, Source, Destination and compatible NiFi version.

New ReadyFlows

- DB2 CDC to Iceberg (Technical Preview)
- DB2 CDC to Kudu
- MySQL CDC to Iceberg (Technical Preview)
- Oracle CDC to Iceberg (Technical Preview)
- PostgreSQL to Iceberg (Technical Preview)
- Slack to Pinecone (NiFi 2.0)
- SQL Server CDC to Iceberg (Technical Preview)

Support for new Kubernetes versions

CDF now supports Kubernetes 1.28 on EKS and AKS.

Include JVM Heap and Thread Dump in UDX

Users can trigger a Diagnostic Bundle collection to include heap and thread dump of all the NiFi nodes of a flow deployment.

Changes and improvements

- The **Dashboard** page has been renamed to **Deployments** and continues to be the single pane of glass to monitor all existing deployments.
- Diagnostic Bundle collection now includes additional information and allows collection of heap dumps for faster troubleshooting.
- **Flow Details** of the **Catalog** page now support text search and filtering by tags.
- When enabling CDF on Azure, CDF now provisions an Azure Database for PostgreSQL instead of a single server. This ensures supportability as single servers get phased out by Microsoft.
- CDF on AWS now supports RDS certificate rotation ensuring compatibility with planned certification rotation changes from AWS.

Removed NiFi components

In this release of CDF-PC, a number of deprecated NiFi processors and controller services have been removed from the porduct. For a list of removed components and suggested replacements, see Removed processors and Removed controller services, respectively.

Fixed issues

- CDPDFX-8257 - Updated sensitive parameter values are now propagated to NiFi
- CDPDFX-8513 - Hello World ReadyFlow corrected
- CDPDFX-8605 - Alert mapping for PROMETHEUS_MEMORY_WARNING_ALERT added so deployment termination won't get stuck
- CDPDFX-8624 - Reassign Workload Resource now updates project id in global deployment table for test sessions so after reassign projects can be deleted
- CDPDFX-8268 - False alerts during vertical auto scaling of Prometheus instances

Technical Service Bulletins

- • TSB 2024-782: Cloudera DataFlow for Public Cloud on AWS requires upgrade before the Root Certificate expires

January 9, 2024

This release (2.7.0-b190) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces new Change Data Capture processors, flow version tagging, deployment configuration reuse, NiFi bulletin monitoring and supports new Kubernetes versions.

What's new

Latest NiFi version

Flow Deployments and Test Sessions now support the latest Apache NiFi 1.24 release.

New Debezium CDC processors

You can now build and deploy flows with Debezium based CDC processors for MySQL, Postgres, Oracle, SQLServer and DB2 databases.

Deployment alerts for NiFi bulletins

CDF now detects bulletin error messages in NiFi flow deployments and displays them as a warning in the dashboard.

Flow Definition versions can now be tagged

You can now use tags to identify versions easier. Common use cases for tags are applying a custom versioning scheme, or labeling flow versions as 'ready for deployment'. The CLI introduces new commands to search flow definitions for tags and streamlines the CI/CD process.

Deployment configurations can now be exported and reused

Customers can now export a configuration archive of existing deployments and reuse them when creating new deployments. This speeds up redeploying new versions of the same flow through the Deployment Wizard by filling in parameter values, KPI configurations etc. from the deployment configuration.

Support for new Kubernetes versions

AKS 1.27 / EKS 1.27

New ReadyFlows

- MySQL CDC to Kudu (Tech Preview)
- Postgres CDC to Kudu (Tech Preview)
- Oracle CDC to Kudu (Tech Preview)
- SQL Server CDC to Kudu (Tech Preview)
- ADLS to Databricks
- S3 to Databricks
- HuggingFace dataset to S3/ADLS
- S3 to IBM watsonx

Changes and improvements

- The HelloWorld ReadyFlow has been updated. You need to deploy Version 2 to run it with NiFi 1.24.0 or higher.
- Underlying AKS creation error during a CDF enablement, if available, is now reported together with the enable failed status event instead of showing up a few minutes later.
- You can now download Client Certificate and Private Key for an Inbound Connection in a single action.

Fixed issues

- CDPDFX-8057: Fixed an issue where TLS certificate renewal was not correctly handled by VPA deployment.
- CDPDFX-7969: Fixed an issue where upgrade fails if there is a deployment using an unsupported NiFi version.
- CDPDFX-8162: Fixed an issue where a deployment creation can be initiated with an existing inbound connection endpoint.
- CDPDFX-8102: Fixed an issue where suspended deployments are reported as not healthy after upgrade.
- CDPDFX-8202 and CDPDFX-8203: CLI improvements when listing projects and deployment archives.
- CDPDFX-8057: Introduced a common component to address TLS certificate reload after periodic renewal.
- CDPDFX-6745: Re-instated bounded automatic helm operation retry during upgrade.
- CDPDFX-8096: Rolling upgrade for NiFi cluster when there is no core NiFi version change.

December 5, 2023

This release (2.6.1-h1-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces a fix that prevents potential NullPointerExceptions (NPEs).

Fixed issues

- CDPDFX-8136 Fixed an issue where the Kubernetes PATCH API response potentially caused an NPE during upgrades and Environment configuration updates in private clusters.

November 1, 2023

This release (2.6.1-b92) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces fine-grained access controls over resources within an environment, support for creating and managing reporting tasks through the CLI, besides other improvements and fixes.

Note: This release of CDF supports deployments running NiFi 1.18.0.2.3.7.0-100 or newer. If your DataFlow service has older NiFi versions, you can perform a Change NiFi Version to bring each into compliance or select to update to the latest as part of DataFlow Upgrade.

What's new

- Support for granular access controls over resources have been introduced to the DF Service. On the new Projects page users can define Projects that limit the visibility of Flow Drafts, Deployments, Inbound Connections, and Custom NAR Configurations within an Environment. Dashboard, Flow Designs and Workspace Resources pages now have new filters and controls to organize resources into Projects.
- Added support for creating and deleting reporting tasks in a CDF deployment through the CLI. Listing reporting tasks is also available in the UI under the NiFi Configuration tab in deployment manager.

Changes and improvements

- Third party and base images updated to address CVEs.
- Improvements for larger scale CDF clusters of up to 50 nodes, including how Prometheus instances are monitored and vertical scaling cluster resources in terms of CPU.
- View NiFi UI is a valid action when the deployment is in a failed to upgrade state.
- Port 80 is removed from the security group for the load balancer that is created by DFX on AWS. It redirected requests to port 443, but port 80 was not necessary to be open.

October 5, 2023

This release (2.6.0-b311) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces NiFi 1.23, vertical scaling, several cost optimization features and the ability to collect Diagnostic Bundles through the UI.

Note: This release of CDF supports deployments running NiFi 1.18.0.2.3.7.0-100 or newer. If your DataFlow service has older NiFi versions, you can perform a Change NiFi Version to bring each into compliance or select to update to the latest as part of DataFlow Upgrade.

What's new

- Added support for downloading NiFi application logs directly from CDF for deployments and test sessions.
- Added support for suspend and resume CDF Deployments. Suspending a Deployment temporarily terminates the NiFi cluster resources and billing activities while maintaining flow persistence.
- Added support for resizing an existing CDF Deployment, which allows increasing and decreasing the memory and CPU allotted to an existing NiFi cluster
- Added support for selecting the Node Storage Profile when creating a CDF Deployment. This allows the use of economical or high performance persistent storage for the NiFi cluster.
- New flow deployments on AWS will now use GP3 volumes instead of GP2 providing a better cost/performance ratio.
- Added active monitoring of NiFi cluster health for CDF Deployments which raises and resolves an alert automatically when a NiFi cluster encounters and recovers from issues impacting its running nodes.
- Added ability to create sensitive dynamic properties in Flow Designer.
- Diagnostic Bundles can now be requested via the Unified Diagnostics UI in the Cloudera Manager. Bundles will be uploaded directly to the selected support case.
- Added Workspace Resources page to workload instances. This provides authorized users with a view of deployments, draft flows, inbound connections, and custom NAR configurations in the workspace.
- Introduced new deployment alerts for Kubernetes component failures like NiFi pods or supporting statefulsets.

Changes and improvements

- Support for updated Kubernetes server versions: AKS 1.26 / EKS 1.25.
- Improved role assignment processing in the DataFlow workload application so role updates are automatically updated within 5 minutes by regular queries to FreeIPA. The user no longer has to log out and log back in to refresh their roles.
- Overhauled NiFi node offloading behavior to be more robust and fault resistant in a variety of cases, including when offloading is disrupted by Kubernetes events.

- Completed significant reliability improvements for larger scale CDF clusters up to 50 nodes, including a variety of efficiency refinements and support for vertical scaling for essential cluster resources.
- Renamed existing Suspend Flow action to Stop Flow. Renamed the resulting deployment state Suspended to Flow Stopped. A CDF Deployment in Flow Stopped state does not remove NiFi cluster resources and the deployment remains billable.
- Released the latest version of NiFi 1.23 which includes the following improvements and CDF integrations:
  - Kubernetes Leader Election and State Management - eliminates the need for a ZooKeeper pod in each CDF Deployment cluster.
- Automatic restart the deployment under the hood if NiFi properties are changed.
- Vertical auto scaling for Prometheus to automatically adjust memory to support a larger number of deployments and metrics.
- Improved stability of helm upgrade during CDF upgrade process.
- Retaining more NiFi provenance data given the current storage capacity allocated for that purpose.

Fixed issues

- CDPDFX-7455: Fixed issue where underlying failure to create AKS cluster on CDF is masked by a 404 error. The actual Azure activity log failure is now raised in a status update event soon after the 404 failure occurs.
- CDPDFX-7540: Fixed issue preventing Kafka to Kudu ReadyFlow from being opened in Flow Designer.
- CDPDFX-7828: Improved sensitivity of CPU based autoscaling behavior when flow metrics autoscaling feature is enabled.
- CDPDFX-7400: Fixed issue rendering an apostrophe in Flow Designer Data Viewer.
- CDPDFX-7399: Fixed issue with MiNiFi pod failing to recover after a CDF upgrade encounters an issue and triggers rollback.
- CDPDFX-7815: Fixed issue with validation of Custom NAR Configuration in RAZ enabled environments.
- CDPDFX-7399: Fixed MiNiFi logging data migration.
- CDPDFX-7704: Fixed racing condition among containers in DFX-Local pod.
- CDPDFX-7737: Fixed issue where UDX collection is not associated with case properly.

August 1, 2023

This release (2.5.0-h2-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud resolves an issue with NiFi cluster autoscaling behavior that can cause unnecessary flow disruptions.

Fixed issues

- CDPDFX-7567 - Fixed unnecessary NiFi cluster pod restarts triggered by autoscaling operations.
  - This issue impacted only those running DataFlow 2.5.0-h2-b1 and deployments configured with autoscaling enabled.
  - Any scale-up or scale-down operation would trigger unnecessary restarts of existing pods resulting in temporary flow disruptions.
  - Upgrading to this new DataFlow version automatically resolves the issue.

Technical Service Bulletins

- TSB 2023-688: Incorrect Java heap memory allocation when creating Small, Medium or Large NiFi flow deployments
- TSB 2023-690: Auto-Scaling operations triggering NiFi pod restarts causing temporary flow disruptions

July 18, 2023

This release (2.5.0-h1-b6) of Cloudera DataFlow (CDF) on CDP Public Cloud resolves a critical bug in JVM heap allocation for newly created NiFi deployments besides other minor bug fixes.

Fixed issues

- CDPDFX-7507 - NiFi deployments assigned incorrect JVM heap settings:
  - This issue only impacts deployments created after upgrading to CDF release 2.5.0-b210.
  - SMALL, MEDIUM and LARGE deployments appear to be utilizing the appropriate amount of heap memory for their allocated cluster resources, but they are only allocated 2GB of heap memory. This results in a possible limitation of flow functionality.
  - Upgrading to CDF release 2.5.0-h1-b6 automatically resolves this issue.
- CDPDFX-7461 - introduces a minor fix to show the proper axis label on metric charts.

Technical Service Bulletins

- TSB 2023-688: Incorrect Java heap memory allocation when creating Small, Medium or Large NiFi flow deployments

June 21, 2023

This release (2.5.0-b210) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces NiFi 1.21, flow metrics driven auto-scaling, advanced UIs for JoltTransformJSON and UpdateAttribute processors, bulk actions in the Flow Designer, UDR support on Azure and makes in-place upgrades generally available to all customers.

New features and changes

- Flow Deployments and Test Sessions now support the latest Apache NiFi 1.21 release.
- In-place upgrades are now generally available to all customers. For more information, see Service upgrade.
- Flow Deployments now support Flow Metrics Scaling in addition to CPU utilization based auto-scaling. For more information, see Auto-scaling flow deployments.
- When selecting Private Cluster during enablement on Azure, CDF no longer provisions a public load balancer and now supports User Defined Routes (UDR).
- Users can now provide their own values for the Kubernetes Pod CIDR Range and Kubernetes Service CIDR Range during enablement.
- CDF now inherits load balancer configuration including available subnets from its associated CDP environment during the enablement process.

Flow Designer

- The Flow Designer now supports multi-selection on the canvas and bulk actions for Start, Stop, Enable, Disable, Move, Change parent group, Copy/Paste, and Delete.
- The Flow Designer now supports the advanced configuration UI for UpdateAttribute.
- The Flow Designer now supports the advanced configuration UI for JoltTransformJson.
- The Flow Designer now supports Birdseye and Zoom controls.
- The Flow Designer now supports Processor Diagnostics with an active Test Session.

ReadyFlows

The following new ReadyFlows have been added to the ReadyFlow Gallery:

- CDW Ingest
- CDP Kafka to Snowflake
- Slack to S3
- Updated Confluent Cloud to Snowflake using new Snowpipe processors

Fixed issues

- CDPDFX-7231 - Fixed an issue where asset uploads were not supported in the Deployment Wizard following publishing from the Flow Designer.
- CDPDFX-7355 - Fixed issue where CDF Upgrade can lead to a stuck active alert indicating metrics are unavailable if it was present prior to the upgrade attempt.

- CDPDFX-7042 - Fixed lack of –no-paginate support on some df list calls in CDP CLI.
- CDPDFX-7060 - Fixed CDP CLI handling of large asset uploads as part of create deployment.

May 02, 2023

This release (2.4.1-h1-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces various improvements and bug fixes for Flow Designer and DataFlow enablement.

Fixed issues

- CDPDFX-7144 and CDPDFX-7116 - Addressed performance issues when determining referencing components of Parameters and Controller Services.
- CDPDFX-7150, CDPDFX-7148, CDPDFX-7151, and CDPDFX-7153 - Made the loading of the Flow in the Flow Designer more efficient by eliminating unnecessary details in the response.
- CDPDFX-7147 and CDPDFX-7167 - Fixed skeleton loader in the Flow Listing.

April 27, 2023

This release (2.4.1-b22) of Cloudera DataFlow (CDF) on CDP Public Cloud supports new NiFi and Kubernetes versions, includes usability improvements for Flow Designer, resolves an issue with deployment ZooKeeper persistence and introduces various bug fixes across the platform.

New features and changes

- CDF now creates Kubernetes clusters with version 1.24 in AWS EKS and 1.25 in Azure AKS.
- The PutIceberg processor is now considered GA (requires selecting NiFi version 1.20.0.2.3.8.2-2).
- Improves event age off logic for environments and deployments that produce a large number of events.
- In the Flow Designer, when deleting a Parameter referencing components are populated in the affected components dialog.
- In the Flow Designer, the Inbound Connection Details button visibility was fixed when a Test Session successfully starts.
- After upgrading CDF, aligned the CFM version of any existing Test Session configurations with the minimum CFM version supported by the Flow Designer.

Fixed issues

- CDPDFX-7081 - Fixed DataFlow Update sometimes timing out when a new AMI is available and AWS performs rolling upgrade of cluster nodes.
- CDPDFX-7018 - Fixed missing persistent volume support for ZooKeeper in deployments.
- CDPDFX-7095 - Fixed missing vault missing Flow Designer policies after CDF upgrade (Relevant only for environments that previously ran with CDF 2.1.x).

April 05, 2023

This release (2.4.0-h1-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces various improvements and bug fixes for Flow Designer and DataFlow enablement.

Fixed issues

- CDPDFX-6970 - Fixed an issue where using the Convert to Parameter feature in Flow Designer cleared all sensitive parameter values.
- CDPDFX-7068 - Added missing support for skipPreflightChecks to Enable and Update DataFlow CLI calls
- CDPDFX-6991 - Fixed an issue with Enable DataFlow for Azure where the cloud provider sometimes reflects clusters in an unexpected UPDATING state.
- CDPDFX-7057 - Fixed Enable DataFlow request for AWS sometimes timing out during initial validation phase due to cloud provider API latency with resources in some regions.

- CDPDFX-6534 - Fixed concurrency issues with uploads of large assets via CLI. Assets of up to 150 MB are now supported.
- CDPDFX-7052 - Fixed an issue with validation of Flow Designer parameters where validation was too strict and could prevent certain values from being entered. The validation is now consistent with parameter values from the Deployment Wizard.

March 14, 2023

This release (2.4.0--b317) of Cloudera DataFlow (CDF) on CDP Public Cloud makes the Flow Designer generally available to all CDP Public Cloud customers, introduces a streamlined way to collect log files, supports new NiFi and Kubernetes versions and introduces several usability improvements when creating and managing flow deployments through the UI.

New features and changes

- The ability to design flows with the new Flow Designer is now generally available to all CDP Public Cloud customers.
- The CDP cli now supports commands to generate diagnostic bundles for CDF for troubleshooting.
- Flow Deployments and Test Sessions now support the latest Apache NiFi 1.20 release.
- When creating a deployment through the deployment wizard, users can now filter for sensitive or empty parameters.
- CDF now creates Kubernetes clusters with version 1.23 in AWS EKS and 1.24 in Azure AKS.
- Upgrading and restarting deployments now triggers a NiFi Health Alert if the NiFi cluster does not return to a healthy state. This can be cleared by performing a Restart Deployment once the underlying issue has been fixed.
- Restart Deployment now supports force suspending the NiFi Flow in case the running flow itself is impairing the health of the NiFi cluster.
- Enable DataFlow action now performs validation steps immediately and fails fast if there are issues with the request.
- Some enablement failure messages are improved when the failure is due to common issues surrounding volume provisioning.
- The following ReadyFlows have been added to the ReadyFlow Gallery:
  - HubSpot to S3/ADLS
  - JDBC to JDBC
  - Kafka to Apache Iceberg
  - Shopify to S3/ADLS

Fixed issues

- CDPDFX-6416 - Reduced default tags applied to cloud resources to avoid tag limit concerns
- CDPDFX-6740 - Fixed private cluster support for AKS
- CDPDFX-6542 - Fixed tags not being applied to Load Balancer created for DataFlow ingress
- CDPDFX-6872 - Fixed issue with control plane missing alert resolutions
- CDPDFX-6591 - DFX Local stability improvements
- Various upgrade reliability improvements

February 13, 2023

This release (2.3.0-h3-b4) of Cloudera DataFlow (CDF) on CDP Public Cloud addresses various bug fixes and improvements in Flow Designer, private cluster support, upgrades and general stability.

Fixed issues

Stability Improvements

- CDPDFX-6591 - Added liveness probe to dfx-local to influence a restart if the container becomes unhealthy.

- CDPDFX-6701 - Fixed a resource leak that would eventually lead to degraded behavior of the application like interruption of workload communication (through heartbeats) to the control plane or the Flow Designer not displaying feedback for running test sessions and preventing processors from being started.

Flow Designer

- CDPDFX-6641 - Updated the redis client in the workload to use a connection pool that was previously aggressively allocating new threads.
- CDPDFX-6635 - Fixed an issue with Flow Designer UI that would continually attempt to reconnect/resubscribe even when their authentication expires.
- CDPDFX-6636 - Fixed an issue with Flow Designer test session event blocking due to infinite event retry.
- CDPDFX-6632 - Ensure additional access rights are enforced when using Flow Designer.
- CDPDFX-6695 - Fixed issues in the Flow Designer that prevented applying Parameter changes when an asset was removed.
- CDPDFX-6695 - Fixed an issue in the Flow Designer where editing a parameter with multiple assets does not trigger a message to apply changes.

Private Cluster Support

- CDPDFX-6740 - Fixed an issue causing AKS private clusters to fail when waiting for load balancer IP via cluster proxy service.

Other Issues

- CDPDFX-6731 - Fixed an issue with CDF 2.1.0-b123 or older not being able to perform workload access certificate renewal.
- CDPDFX-6634 - Fixed issue that would result in the user being redirected to the wrong URL following successful authentication on the workload.
- CDPDFX-6655 - Fixed issue that would load the app switcher in the sidebar in the Deployment Manager and Flow Designer with the incorrect options.
- CDPDFX-6567 - Fixed an issue with logging in AWS AP region 3 (Jakarta).
- CDPDFX-6693 - Fixed navigation when using the Home button from the Manage Deployment page.
- CDPDFX-6713 - Fixed DataFlow configuration update failures when attempting to add new authorized IP ranges without changing minimum or maximum cluster capacity.

Technical Service Bulletins

- TSB 2023-648: ApacheNiFi pods not properly configured with fully qualified domain name causing communication errors
- TSB 2023-649: File watcher container in MiNifi logging pod causing process ID leaks

Behavioral changes

If you defined a custom NAR storage location during flow deployment in earlier releases, the setting got stored as the default value for that environment. CDF 2.3.0-h3 changes this behavior: custom NAR configuration does not get stored for the environment any more. This means every time you deploy a flow that requires custom NARs, you need to configure the access credentials and the storage location.

January 25, 2023

This release (2.3.0-h2-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces fixes and improvements for etrics collection, enablements, upgrades and general stability.

Fixed issues

- CDPDFX-6458 - Fixed issue where NiFi flow deployents failed at host name verification due to a startup script returning the wrong host name.
- CDPDFX-6671 - Reduced metric collection sample sizes in Cluster-level Prometheus, resulting in lower memory footprint and faster query times without regression or reduction of performance for existing environment metrics.

- CDPDFX-6677 - Fixed issue where the import of very large flows failed with timeout.
- CDPDFX-6683 - Fixed issue where a TimeoutException was thrown without its cause being logged.

January 20, 2023

This release (2.3.0-h1-b8) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces fixes and improvements for Flow Designer, enablements, upgrades and general stability.

Fixed issues

- CDPDFX-6593 - Fixed issue where test session reports the wrong state after being started and stopped.
- CDPDFX-6582 - Fixed issue where processor property configuration does not support special characters.
- CDPDFX-6596 - Fixed errors which occurred in certain interdependent controller services configurations.
- CDPDFX-6594 - Fixed issue where suspend updates flag is not cleared after an enable attempt fails, but retry enable succeeds. This results in newly created deployments not displaying properly in the deployment dashboard because heartbeats are being ignored.
- CDPDFX-6578 - Fixed upgrade issue resulting in larger deployment sizes having insufficient java heap memory allocated to NiFi pods, resulting in Out of Memory issues.
- CDPDFX-6561 - Fixed issue preventing retry upgrade when previous upgrade attempt failed after kubernetes upgrade to 1.23 but before liftie service finished upgrading deployments..
- CDPDFX-6479 - Fixed redis pods being unable to recover after non-graceful shutdown
- CDPDFX-6512 - Improved minifi logging stability by defining memory needs for toybox container

January 03, 2023

This release (2.3.0-h1-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces a fix for loading custom NARs.

Fixed issues

- NIFI-10981 - Fixed an issue where NarAutoLoader started before the provider retrieved all NARs.

December 08, 2022

This release (2.3.0-b347) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces the technical preview of Flow Designer, adds several new ReadyFlows, improves upgrade reliability and fixes issues with stability for clusters with high utilization.

New features and changes

- Flow Designer [Technical Preview]

  Developers can now build new data flows from scratch using the integrated Designer. Flow Drafts can be tested using Test Sessions before they are published to the Catalog. Once they are in the catalog, users can create Flow Deployments using the Deployment Wizard or CLI.
- Added support for certificate renewal of inbound connection endpoints
- Six new ReadyFlows have been added to the ReadyFlow Gallery:
  - Airtable to S3/ADLS
  - Box to S3/ADLS
  - Dropbox to S3/ADLS
  - Google Drive to S3/ADLS
  - MQTT to Kafka
  - Salesforce to S3/ADLS
- Apache Iceberg Integration [Technical Preview]

  Developers can now use the PutIceberg processor to write data to Apache Iceberg in CDP.

- DataFlow is now also available in the EU and APAC regional CDP control planes.

  Note: DataFlow Functions is an exception, as it is only available in the US control plane.

- Updated to support using Kubernetes 1.23 EKS and AKS clusters.

  If you use restricted IAM roles in your EKS clusters, this impacts the upgrade procedure. For more information, see the CDF Service Upgrade documentation.
- General improvements to DataFlow upgrade reliability to allow for improved success rates.
- Includes new NiFi version 1.18.0.2.3.7.0 for flow development and flow deployment.
- Updated DataFlow Function binaries to use NiFi version 1.18.0.

Fixed issues

- CDPDFX-5884 - Internal Logging Improvements - fixed logging issues sent to S3 (AWS) and Azure Blob Storage (Azure)
- CDPDFX-6308 / CDPDFX-6323 - updated legacy inbound connection endpoints to make them compatible with new certificate renewal capability
- CDPDFX-6305 - Fixed issue with inbound connection endpoints created prior to DFX 2.2.0 preventing upgrade
- CDPDFX-5689 - Expanded character limit on deployment parameters to allow for parameters of up to 100000 characters.
- CDPDFX-5942 - Fixed a rare issue where a cluster update request can fail due to cloud API latency leading to timeouts.
- CDPDFX-6334 - Fixed an issue where a large deployment count with active metrics unavailable alerts can lead to a deadlock which prevents upgrade
- CDPDFX-6345 - Fixed issue which causes existing deployments using assets to no longer allow changes to those assets after upgrade due to a certificate mismatch
- CDPDFX-6364 - Fixed CLI issue where some listen components are ignored when creating deployments with multiple listen components
- CDPDFX-6378 - Fixed issue where a post-upgrade cleanup effort can fail with a warning if a customer’s outbound network restrictions do not allow retrieval of a liquibase schema file.
- CDPDFX-5731 - Cert-manager OOM - increased memory limits for cert-manager to ensure certificates can still be issued as deployment numbers increase
- CDPDFX-6371 - Correct nifi.provenance.repository.debug.frequency to a parseable default value (default 1000000)
- CDPDFX-6309 - ZK-Operator OOM - prevent operator from running OOM due to it scan/filtering of multiple secrets
- CDPDFX-6118 - Suppress deprecation logging in NiFi
- CDPDFX-6387 - Increased cpu limit for cer-manager-webhook container
- CDPDFX-6235 - Fix errors seen during NiFi autoscaling due to inconsistent nifi identity configmap updates
- CDPDFX-6230 - Fix to prevent metering alert from firing when DataFlow service not fully enabled

October 21, 2022

This release (2.2.0-h4-b2) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces compatibility fixes around upgrades and public ingress certificate renewals.

Fixed issues

- CDPDFX-6119 - Fixed an issue where a failed attempt to renew public ingress certificates incorrectly fires a certificate expiring soon warning.
- CDPDFX-6189 - Fixed an issue where public ingress certificate renewal fails for workloads that are not running the latest DataFlow version.
- CDPDFX-6191 - Fixed an issue where DataFlow upgrade fails if a public ingress certificate renewal has ever been attempted on the DataFlow service.

- CDPDFX-4177 - Improved reuse of recently generated public ingress certificate upon repetitive retries to avoid a Let's Encrypt certificate rate limit of 5 per 7 days.

October 10, 2022

This release (2.2.0-h3-b2) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces compatibility fixes around upgrades.

Fixed issues

- CDPDFX-6101 - Fixed an issue with DataFlow upgrade leaving deployments using custom NARs not able to retrieve the NAR components from S3. This results in custom processors becoming non-functional.
- CDPDFX-6120 - Fixed an issue with DataFlow Upgrade where the internal self-signed CA certificate is being unnecessarily renewed in upgrade scenarios where the certificates are not also renewed because they were already renewed in prior upgrades. This issue would have arisen for customers upgrading from DFX 2.2.0-h1 or 2.2.0-h2 to this version.
- CDPDFX-6025 - Fixed a documented known issue where a CDP environment configured with a storage base location with an S3 bucket but no root directory results in workload logs not being uploaded.
- CDPDFX-6055 - Fixed an issue where Update DataFlow fails on DataFlow services running version 2.0.0, 2.1.0 and 2.1.0-h1.
- CDPDFX-5944 - Added support to allow certificate renewal action on a DataFlow service in an 'Update Failed' state.

September 29, 2022

This release (2.2.0-h2-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces fixes improving performance of control plane database queries.

Fixed issues

Control Plane bug fixes

September 28, 2022

This release (2.2.0-h1-b11) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces fixes for certificate renewal issues that could cause the deployment manager and NiFi canvas to become inaccessible.

Fixed issues

- CDPDFX-5941 - Fixed an issue where CDF did not restart automatically to pick up renewed certificates resulting in the deployment manager and NiFi canvas becoming inaccessible.
- CDPDFX-5943 - Fixed an issue where the trusted certificate CA was not updated for vault after renewing certificates resulting in the deployment manager and NiFi canvas becoming inaccessible.
- CDPDFX-5953 - Fixed an issue where CDF provisioned certificates with an incorrect CA lifetime resulting in the deployment manager and NiFi canvas becoming inaccessible after 90 days .
- CDPDFX-5954 - Fixed an issue where prometheus and alertmanager pods did not correctly reload certificates upon renewal.
- CDPDFX-5961 - Added renewal of all certificates during CDF upgrades to prevent any further certificate related issues.
- CDPDFX-6025 - Fixed an issue where no CDF logs were written to S3 when CDP’s storage base location contains only a bucket name.
- CDPDFX - 5990 - Fixed an issue where a ‘metrics unavailable’ alert was incorrectly fired during CDF upgrades.
- CDPDFX-5965 - Fixed an issue where flow definitions failed to deploy successfully due to their size.

Technical Service Bulletins

- TSB-2022-621: Certificate Authority (CA) certificate lifespan being too short causing authentication problems with the Cloudera DataFlow Service

September 08, 2022

This release (2.2.0-b194) of Cloudera DataFlow (CDF) on CDP Public Cloud improves the first time user experience, introduces more ReadyFlows, and supports renewing certificates for Inbound Connection.

New features and changes

- A streamlined first time user experience

  You can use the new Hello World ReadyFlow to create your first flow deployment without dependencies on source or target systems. New landing pages for the Dashboard and Catalog allow you to get started with just one click.

- Four additional new ReadyFlows are available in the ReadyFlow Gallery
  - Non-CDP ADLS to S3/ADLS
  - Non-CDP S3 to S3/ADLS
  - ListenTCP filter to S3/ADLS
  - ListenSyslog filter to S3/ADLS
- Endpoint Hostnames certificates can now be renewed when using Inbound Connections

Fixed issues

- CDPDFX-3348 - Support custom KMS encryption key defined at CDP environment level
- CDPDFX-5126 - Improve tag validation for DataFlow service enablement
- CDPDFX-5449 - Fixed an issue where Data offloading during scale down operations is blocking reconciliation in CDF causing deployment requests to time out
- CDPDFX-5494 - Fixed an issue where Kubernetes node scaling events are not surfaced in the event history for DataFlow services
- CDPDFX-5555 and CDPDFX-5556 - Improved reliability of Kubernetes infrastructure reuse when DataFlow enablement fails
- CDPDFX-5577 - Fixed an issue where UDP Inbound Connections fail in AWS due to incorrect health check configuration for inbound data load balancer
- CDPDFX-5607 - Improve visibility of helm release details when enable or upgrade DataFlow action fails to complete

August 16, 2022

This release (2.1.0-h1-b10) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces improvements for Prometheus performance, reliability and performance improvements for upgrades, as well as a number of additional bug fixes.

Improvements and fixed issues

- Fixed regression in DFX 2.1.0 where new deployments contain NiFi pods with misconfigured cpu/memory limits. These limits are lower than necessary, leading to the possibility of NiFi pod crashing due to off-heap storage growing large enough to cause an Out of Memory error.
- CDPDFX-5483 - Fixed a bug where disabling DataFlow from certain failure states can potentially leave behind some cloud resources deployed by DataFlow.
- CDPDFX-5486 - Several settings for Prometheus instances running on the workload cluster have been adjusted to improve performance. This includes a reduction in the amount of metrics collected by Prometheus instances and enabling write ahead log compression.
- CDPDFX-5487 - Improved Enable, Upgrade/rollback and Create deployment performance to more reliably interpret status of helm releases and NiFi resources for improved success rates.
- CDPDFX-5494 - Fixed an issue where cluster node scaling events are not being properly surfaced to the user.

- CDPDFX-5495 - Fixed a bug where upgrades failed if a deployment was created with autoscaling disabled, but since was updated in any way (node count, params, and so on).
- CDPDFX-5498 - Improved upgrade timeout configurations to account for delays upgrading clusters with high node counts.
- CDPDFX-5568 - Fixed a bug causing some environments not to show in the filter options on the Deployment Dashboard.
- CDPDFX-5571 - Fixed potential for up to 10-minute delay during upgrades where helm repository sync is unnecessarily delayed. This could sometimes lead to timeout failures during upgrade.
- CDPDFX-5572 - Fixed a bug in the user interface on the Deployment Dashboard when an environment filter option is selected, but becomes no longer valid.
- CDPDFX-5606 - Fixed a bug in retrying a failed DataFlow enablement where a new enable attempt fails fast and old infrastructure is eventually deleted. Old infrastructure details are not cleared from the database, so subsequent retries attempt and fail to clean up infrastructure that is already deleted.
- CDPDFX-5676 - Fixed an issue where in-cluster certificates are renewed, but certain cluster resources are not automatically restarted to start using the new certificates.
- CDPDFX-5736 - Fixed an issue limiting upload of flow definition files to the catalog to 1MB in size. Flow uploads are now supported up to 25MB in size.
- CDPDFX-5780 - Fixed an issue preventing creation of deployments with large flow definitions.

June 28, 2022

This release (2.1.0-b123) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces support for DataFlow service upgrades as a technical preview feature, Inbound Connection configuration through the create-deployment CLI command, Kubernetes 1.22 as well as additional improvements and fixes.

New features and changes

- [Technical Preview] CDF environment and deployment upgrades are now supported.

  CDF service upgrades require a specific entitlement that is granted by request only. Reach out to your Cloudera Account Representative or open a support ticket to request CDF service upgrades.
- The `create-deployment` CLI command now supports Inbound Connections.
- CDF now correctly configures the deployment when it detects any of the two reserved Controller Services Default NiFi SSL Context Service and Inbound SSL Context Service regardless of whether they are external references or not.
- Kubernetes Version 1.22 is now the default Kubernetes version for CDF. When CDF is enabled, it creates AWS EKS and Azure AKS clusters based on version 1.22.
- CDF now supports Data Lake scaling. After the Data Lake scale operation has completed, affected deployments will show an active alert and need to be manually restarted to put the new configuration in effect.
- The Event Hub to ADLS ReadyFlow now allows you to specify an Event Hub Service Bus Endpoint.

Fixed issues

- CDPDFX-5317 - Fixed an issue that could cause timeouts when uploading assets as part of the Deployment Wizard
- CDPDFX-5400 - Fixed several issues that could cause cluster and deployment Prometheus instances to run out of memory
- CDPDFX-5322 - Fixed an issue with caching of UI content that sometimes required user to perform a hard refresh

- CDPDFX-5261 - Fixed an issue that causes autoscaling CPU utilization target to change from 75% to 80% if user performs a deployment update

  CDPDFX-5325 - Use availability sets instead of zones for Azure to reduce the chance of a customer failing to enable DataFlow due to instance type resource constraints in certain regions

  CDPDFX-5319/CDPDFX-5385 - Improved resiliency when recovering from network disruptions and database failovers to improve the chances of deployment related requests succeeding

  CDPDFX-5260 - Expanded the types of cluster resource logs that are exported to cloud storage

  CDPDFX-5483 - Fixed an issue where some cloud resources are not deleted properly when disabling DataFlow after it has encountered an update failure or upgrade failure

May 17, 2022

This release (2.0.0-b302) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces support for Inbound Connection Endpoints, allowing you to stream data from external sources to Flow Deployments. It also adds support for Apache NiFi 1.16, adds seven new ReadyFlows and improves stability of the service.

New features and changes

- Inbound Connection Endpoints for Flow Deployments

  You can now deploy NiFi flows containing processors like ListenHTTP, ListenTCP, ListenUDP, ListenSyslog etc. and let CDF-PC take care of exposing a stable and secure endpoint to connect your clients to.
- Flow Deployments now support the latest Apache NiFi 1.16 release
- Seven new ReadyFlows are available in the ReadyFlow Gallery
  - Azure Event Hub to ADLS ReadyFlow
  - Confluent Cloud to S3/ADLS ReadyFlow
  - Confluent Cloud to Snowflake ReadyFlow
  - JDBC to S3/ADLS ReadyFlow
  - ListenHTTP to CDP Kafka ReadyFlow
  - Non-CDP ADLS to CDP ADLS ReadyFlow
  - Non-CDP S3 to CDP S3 ReadyFlow
- Single AZ support for DataFlow clusters

  CDF Kubernetes nodes are now always deployed in a single Availability Zone (AZ). When enabling CDF, you still have to provide subnets in at least two Availability Zones to satisfy Kubernetes control plane requirements.
- Value of javax.net.ssl.trustStore property no longer set to existing trust store on NiFi node

  The javax.net.ssl.trustStore system property is no longer set by default to the existing truststore on the NiFi node. Flows which rely on a truststore being available by default via the JVM should be updated to leverage the SSLContextService.

  This change may cause errors affecting older existing processors or custom processors that do not have a property to reference an SSL Context Service.
- Flow Deployments on Azure now use Premium SSD disks to store NiFi's various repositories
- CDF-PC integration into the CDP CLI is now generally available. Make sure you run the latest CDP CLI version to use cdp df commands.
- Python 3.9 is now installed on NiFi containers allowing users to run python scripts through NiFi processors. Requests and urllib3 libraries are also installed by default.
- Information level alerts are now displayed in the Dashboard to inform users about an action they should perform for a Flow Deployment without changing the Flow Deployment status to Concerning.
- Deployments can now be restarted from the Deployment Manager to perform actions that require a NiFi process restart.
- Filters that have been set in the Dashboard are now persisted after a user has navigated to the Deployment Manager and returned back to the Dashboard.

- When enabling DataFlow, you can now specify CIDRs as a comma separated list to configure Load Balancer and Kubernetes API Server Endpoint Access.

Fixed issues

This release includes all cumulative fixes released throughout 1.1.0-h1 to 1.1.0-h5 as well as additional stability and performance improvements.

April 06, 2022

This release (1.1.0-h5-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud fixes a deployment issue on Azure and introduces a new NiFi version (1.15.0.2.3.3.1-2) which fixes several issues that prevented deployments from downscaling.

Fixed issues

- CDPDFX-4440 – Fixed an issue that caused deployments to fail on Azure when a persistent storage volume was provisioned in a different Availability Zone than the NiFi pod.
- NIFI-9433 -- Fixed an issue that resulted in negative queue counts when using load balanced connections.
- NIFI-9835 -- Fixed an issue that prevented downscaling when queue count was negative.
- NIFI-9847 -- Fixed an issue where use of certain processors could prevent a NiFi node from offloading its data.

March 23, 2022

This release (1.1.0-h4-b2) of Cloudera DataFlow (CDF) on CDP Public Cloud fixes several issues around certificate renewal of long running DataFlow services and improves stability and performance for deployments with a large event history.

Fixed issues

- CDPDFX-4251 – Fixed an issue that could cause DataFlow service disable to fail when disabling was initiated through deleting the associated CDP environment.
- CDPDFX-4204 – Improved how CDF ages off historical heartbeats for increased stability and performance.
- CDPDFX-4203 – Reduced the retention time of events and metrics gathered in the DataFlow Prometheus instance for increased stability and performance.
- CDPDFX-4199 – Fixed an issue where DataFlow could become unresponsive due to the high number of events for a deployment.
- CDPDFX-4198 – Fixed an issue where DataFlow attempted to revoke already expired certificates causing the certificate renewal to fail.
- CDPDFX-4201 – Fixed an issue that prevented users from initiating certificate renewal for healthy DataFlow services when the associated CDP environment was in an unhealthy state.

March 8, 2022

This release (1.1.0-h3-b2) of Cloudera DataFlow (CDF) on CDP Public Cloud fixes an issue that caused DataFlow enablements to fail on AWS.

Fixed issues

CDPDFX-4010 – Fixed an issue that caused DataFlow enablements to fail on AWS due to requesting a deprecated Postgres version for RDS.

February 15, 2022

This release (1.1.0-h2-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces support for Kubernetes v1.21, removes disk-related metrics, and fixes a NiFi deployment issue.

New features and changes

Added support for Kubernetes v1.21

Version 1.21 is now the default Kubernetes version for CDF. When CDF is enabled, it creates AKS/EKS clusters based on version 1.21.

Removed disk related metrics from deployments and DataFlow service details

In previous versions of CDF, deployments and enabled DataFlow services showed disk capacity and disk usage metrics as part of their system metrics. You were also able to define KPIs and alerts on these metrics. Due to issues with the underlying metrics collection framework, the following metrics have been removed starting with CDF 1.1.0-h2-b1:

- Disk Capacity (DF Service Metric)
- Disk Capacity (Deployment System Metric)
- Disk Usage (Deployment System Metric)

Fixed issues

CDPDFX-4030 – Fixed an issue where NiFi pods were not correctly applying the heap size which the user selected during deployment.

Documentation updates

Removed disk related metrics documentation.

January 27, 2022

The 1.1.0-h1-b2 and 1.1.0-b127 releases of Cloudera DataFlow (CDF) on CDP Public Cloud introduce the following changes.

January 27, 1.1.0-h1-b2

Fixed issues

- Fixed an issue that prevented DataFlow from being enabled on Azure environments.

January 27, 1.1.0-b127

New features and enhancements

- The ability to deploy CDF on Azure is now generally available to all CDP Public Cloud customers.
- Flow Deployments now support the latest Apache NiFi 1.15 release.
- CDF now creates Kubernetes clusters with version 1.20 in AWS EKS and Azure AKS.
- CDF now supports username/password authentication for AWS environments with non-transparent proxy setups.
- When using the Default NiFi SSL Context Service in a flow deployment, the automatically generated truststore now contains the default cacerts from the local JDK. This ensures that the SSL Context Service can be used with third-party applications using certificates from common public certificate authorities.
- Users can now perform the Download Kubeconfig and Manage User Access actions when an enablement request fails. This allows in-depth troubleshooting of a failed enablement attempt.
- The following ReadyFlows have been added to the ReadyFlow Gallery:
  - ADLS to ADLS Avro ReadyFlow
  - Kafka to ADLS Avro ReadyFlow

Fixed issues

- CDPDFX-3888 - Fixed an issue that prevented CDF from being enabled in Azure regions where the multi-AZ feature is not available.
- CDPDFX-3873 - Fixed an issue where large file uploads could cause flow deployment requests to fail on Azure.

- CDPDFX-3909 - Fixed an issue that prevented using special characters like “@” as tag values when enabling CDF.

Documentation updates

- Documentation has been created for the new ADLS to ADLS Avro ReadyFlow and Kafka to ADLS Avro ReadyFlow.
- NiFi Endpoints Support Matrix has been updated.
- Troubleshooting information has been updated for Azure.
- Azure technical preview references have been removed.

December 21, 2021

This release (1.0.3-h2-b1) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

Fixed issues

- Fixed bug causing DataFlow to incorrectly appear to have been upgraded to the latest version.

  When you use Manage DataFlow to modify your DataFlow service, updates to the following values cause the Workload Version field for your service to be set to the current DataFlow application version.

  - Capacity
  - Networking
  - Kubernetes API Server Endpoint Access

  Once the update is started, it will incorrectly appear as if your DataFlow service is upgraded to the latest version.

  CDPDFX-3897

NOTE: There is no need to tear down and recreate your enabled DataFlow services to benefit from these updates.

December 17, 2021

This release (1.0.3-h1-b6) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

Fixed issues

- Upgraded all log4j and logback libraries present in DataFlow application to the latest version.
- Upgraded all log4j and logback libraries present in the DataFlow NiFi Runtime to the latest version.
- Removed requirement for setting unnecessary ingress rules on Azure Network Security Group during AKS cluster creation, which could trip policy rules preventing such NSG updates.
- Fixed a bug that prevented selecting specific Load Balancer subnets when enabling Azure DataFlow service.

December 1, 2021

This release (1.0.3-222) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

New features and enhancements

[Technical Preview] CDF can now be deployed on Azure

Deploying CDF on Azure requires a specific entitlement that is granted by request only. Reach out to your Cloudera Account Representative or open a support ticket to request CDF on Azure.

Custom Processors and Controller Services can be used in NiFi flow deployments.

CDF supports deploying NiFi flows that require custom components such as processors or controller services. Users can store their custom NiFi Archive (NAR) files in an object store location in S3 or ADLS and use the Deployment Wizard or the CLI to create their flow deployments accordingly. Once configured, CDF ensures that the custom NARs are available to all NiFi pods as the deployment scales up and down.

Non-transparent proxy support in AWS environments

If you have a non-transparent proxy registered for your CDP environment, CDF reuses the existing configuration and routes its outbound traffic through the proxy.

Note: Depending on the NiFi flows you are using, you might have to configure additional proxy settings if you require the NiFi data to go through the proxy as well.

Private EKS/AKS Kubernetes API Endpoints

CDF now supports private EKS/AKS Kubernetes API Endpoints for CDP environments which have been setup using Cloudera Connectivity Manager (CCM v2). When enabling CDF, you need to select a checkbox to create a private endpoint.

Enabling CDF for an environment with a minimum of one EKS/AKS instance

CDF can now be enabled with a minimum of one EKS/AKS instance and will scale up as required until the specified maximum number of Kubernetes nodes is reached. Previously the minimum number of nodes was set to three.

CDF instance sizing changes

CDF now uses c5.4xlarge instances on AWS and Standard_D16s_v4 on Azure to provide more efficient infrastructure usage.

NiFi node sizing changes

NiFi Node Sizes in the Deployment Wizard have also been updated to use available resources more efficiently. The new options are:

- Extra Small – 2 vCores, 4GB memory
- Small – 3 vCores, 6GB memory
- Medium – 6 vCores, 12GB memory
- Large – 12 vCores, 24GB memory

Flow deployment security certificates

Security certificates for deployments are now automatically renewed one month before their expiration. A manual restart is required to pick up the new certificates.

CDF UI guidance on using CDP CLI

The CDF UI now provides guidance on how to use the CDP CLI to enable and disable CDF, upload and delete flow definitions as well as editing existing flow deployments.

CDP Beta CLI support for editing flow deployments

The CDP Beta CLI now includes functionality to edit existing deployments allowing configuration changes such as parameter and KPI updates, scaling configuration, or NiFi version upgrades.

New events and alerts for pod scaling

Introduced deployment pod scaling alerts and events to reflect when a deployment auto scales and when scaling limits are approached.

Fixed issues

- CDPDFX-3568 – Fixed an issue where no tooltip data was shown for the Data Throughput charts in the Dashboard.
- CDPDFX-3588 – Fixed an issue where CDF did not write logging data to non us-west-1 S3 buckets.
- NIFI-9217 – Fixed an issue that prevented scale down operations due to NiFi nodes getting stuck in “offloading” mode.
- CDPDFX-3624 – Fixed an issue where deployment auto-scaling value changes were not saved.

October 13, 2021

This release (1.0.2-244) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

Fixed issues

- Fixed an issue that caused some DataFlow CLI commands to be unavailable in the Beta CDP CLI.

Documentation updates

The following documentation has been updated with CLI instructions.

- Enable DataFlow for an Environment
- Disable DataFlow for an Environment
- Deploy a flow
- Importing a flow definition to Cloudera DataFlow
- Importing a new version of a flow definition
- Adding a ReadyFlow to the Catalog

September 29, 2021

This release (1.0.2-b240) of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

New features and enhancements

- CLI support has been added for creating flow deployments and flow definition management.

  DataFlow CLI is available in the Beta CDP CLI.
- Added the S3 to S3 ReadyFlow leveraging S3 bucket notifications.
- DataFlow now supports Cloudera RAZ (fine-grained access control) for object store access.
- The Dashboard, Catalog, and Environments pages now persist search queries and applied filters for a better user experience.
- When deploying a flow definition, you can now select whether the NiFi flow should be started automatically.

Fixed issues

- Fixed an issue where Alerts would not correctly become inactive after the alert condition is no longer met.

  CDPDFX-3396
- Fixed an issue where Alerts were still generated when a deployment was suspended.

  CDPDFX-3419
- Fixed an issue that prevented users from using carriage return in the Flow Definition description.

  CDPDFX-3294
- Fixed an issue that resulted in environment system metrics being rendered multiple times.

Documentation updates

- ReadyFlow documentation
  - New ReadyFlow documentation has been developed for all ReadyFlows

    Using ReadyFlows
  - Documentation for the new S3 to S3 Avro with S3 Notifications ReadyFlow

    S3 to S3 Avro with S3 Notifications ReadyFlow Overview
- CLI documentation
  - Enable DataFlow for an Environment
  - Disable DataFlow for an Environment
  - Deploy a flow

- Updated ReadyFlows for RAZ integration

  The three AWS ReadyFlows include instructions on configuring access to AWS buckets when your environment is RAZ enabled

  - Kafka to S3 Avro
  - S3 to S3 Avro
  - S3 to S3 Avro with S3 Notifications
- NiFi Endpoints Support Matrix

August 24, 2021

This release of Cloudera DataFlow (CDF) on CDP Public Cloud introduces these changes.

Fixed Issues

- Fixed a bug that prevented DataFlow and associated deployments from being terminated when the associated CDP Environment had been deleted.

  CDPDFX-3327
- Fixed a bug that could result in duplicate system metrics being displayed in the Environment details view.

  CDPDFX-3318
- DataFlow now uses the single api.us-west-1.cdp.cloudera.com endpoint for Control Plane API access.

  CDPDFX-3288

August 16, 2021

This is the GA release of the Cloudera DataFlow service.

Cloudera DataFlow (CDF) is a CDP Public Cloud service that enables self-serve deployments of Apache NiFi data flows from a central catalog to auto-scaling Kubernetes clusters managed by CDP. Flow deployments can be monitored from a central dashboard with the ability to define KPIs to keep track of critical data flow metrics.

CDF eliminates the operational overhead that is typically associated with running Apache NiFi clusters and allows users to fully focus on developing data flows and ensuring they meet business SLAs.

Related Information

Cloudera DataFlow Overview

Cloudera DataFlow Quick Start

Supported component versions

Cloudera DataFlow includes support for the following component versions.

Cloudera DataFlow version	Release date	Cloudera DataFlow component	Component version
Cloudera DataFlow 2.9.0 hotfix 5, build 2	February 05, 2025	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0

Cloudera DataFlow version	Release date	Cloudera DataFlow component	Component version
Cloudera DataFlow 2.9.0 hotfix 4, build 3	December 03, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.9.0 hotfix 3, build 1	November 20, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.9.0 hotfix 2, build 1	November 15, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.9.0 hotfix 1, build 1	November 11, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.9.0, build 383	October 31, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.27.0 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0

Cloudera DataFlow version	Release date	Cloudera DataFlow component	Component version
Cloudera DataFlow 2.8.0 hotfix 2, build 2	June 25, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.8.0 hotfix 1, build 1	June 03, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.8.0, build 274	May 20, 2024	Apache NiFi	2.0.0-M2 [Technical Preview] 1.25.0 1.24.0 1.23.0 1.21.0 1.20.0 1.18.0
Cloudera DataFlow 2.7.0, build 190	January 9, 2024	Apache NiFi	1.24.0 1.23.0 1.21.0 1.20.0 1.18.0

Related Information

Cloudera Support lifecycle policy

Kubernetes Ingress NGINX Controller vulnerabilities

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname 'IngressNightmare'.

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions

- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

How do these vulnerabilities affect Cloudera DataFlow on cloud?

For mitigating CVE-2025-1974 on Cloudera DataFlow on cloud, refer to the information below.

Note: Cloudera recommends limiting direct access to cluster hosts to only authorized administrators and auditing all activity as a security best practice.

Mitigation of CVE-2025-24514, CVE-2025-1097, CVE-2025-1098, and CVE-2025-24513 is secondary to the previous CVE. They require no immediate action, as attackers can only exploit these with direct access to cluster hosts and privileges to create arbitrary ingress objects via the Kubernetes API.

Note: Cloudera has tested these mitigation steps only on the currently supported releases. Customers using older versions are advised to upgrade to a supported release before attempting the mitigation actions.

For the latest updates on this issue, see the corresponding Knowledge articles:

- TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication
- TSB 2025-839 Mitigation steps for Cloudera DataFlow on cloud

Instructions

1. Check your current version

To determine which version your cluster is running, use either the UI or the CLI:

DataFlow Service version discovery from the UI:

1. In Cloudera DataFlow, from the **Environments** page, select the environment you want to inspect.
2. Click Manage DataFlow from the **Environment Details** pane.
3. Look for “DATAFLOW VERSION”.

DataFlow Service version discovery from the CLI:

1. Run the cdp df list-services command to display a list of the environments with their workload versions included.
2. Run cdp df describe-service --service-crn <DataFlow CRN> to display more detailed information about an environment and display the DataFlow version.

Note: For instructions on how to install and use the Cloudera CLI, refer to the CLI documentation.

2. Mitigation steps for Cloudera DataFlow on Cloud

Upgrade the “ingress-nginx” controller component deployed by Cloudera DataFlow to version v1.11.5 for all environments running Cloudera DataFlow versions 2.9.0 and lower. To apply the patch to the environment:

1. Go to Cloudera DataFlow in the Cloudera Management Console.
2. Select the target environment for the patch from the “Environments” list.
3. Click “Manage Cloudera Data Flow service” in the upper right corner.
4. Click the “Actions” dropdown menu, and select “Download Kubeconfig”.
5. A KubeConfig yaml file will be downloaded to your local machine. (Named “kubeconfig.yaml” by default.)
6. Depending on your operating system, follow the instructions at https://kubernetes.io/docs/tasks/tools/#kubectl to install the “kubectl” binary to your local machine.
7. Create a text file called bump-controller-version.yaml, add the following text, and save the file.

```
# kubectl -n nginx-ingress patch helmreleases.helm.toolkit.fluxcd.io ing
ress-nginx --patch-file admission-webhook-off.yaml --type merge
spec:

  values:

    ingress-nginx:

      controller:

        image:

          tag: v1.11.5
```

8. Open a terminal or command prompt, and run (from the directory you saved the yaml files) the following command to test your access to the Kubernetes API server:

   kubectl --kubeconfig kubeconfig.yaml get nodes

   The command prints out information about the list of worker nodes.
9. Execute the following command to apply the patch:

   kubectl --kubeconfig kubeconfig.yaml -n nginx-ingress patch helmreleases ingress-nginx --patch-file bump-controller-version.yaml --type merge

   It prints out helmrelease.helm.toolkit.fluxcd.io/ingress-nginx patched on successful execution.
10. Execute the following command and verify that the printout in the command line ends with v1.11.5:

    kubectl --kubeconfig kubeconfig.yaml -n nginx-ingress get deployment dfx-nifi-nginx-ingress-controller -o jsonpath='{.spec.template.spec.containers[?(@.name=="controller")].image}'
11. Execute to watch the progress of the upgrade as the ingress-nginx controller pod is restarted in the background:

    kubectl --kubeconfig kubeconfig.yaml -n nginx-ingress get deployments dfx-nifi-nginx-ingress-controller -w
12. Wait until the columns show the following:
    a. “READY” column showing “1/1”,
    b. “UP-TO-DATE” column showing “1”,
    c. “AVAILABLE” column showing “1”
13. Press “Ctrl-C” to return to the command line.

Note: While executing the upgrade, for a short duration (usually a second or two)users may briefly lose connectivity to the workload side UI for deploying new flows or accessing the NiFi Web UI of existing flows, but there’s no interruption to data processing in any existing flows.

14. Verify that the new “ingress-nginx” controller is working properly by accessing the NiFi Web UI of any existing flow.
15. Repeat the process for each environment where the patch needs to be applied.

Important:

1. The mitigation steps apply a patch to the DataFlow environments. Before upgrading the patched environment of Cloudera DataFlow to another version, the patch must be reverted temporarily. To revert the patch, execute the following command:

```
kubectl --kubeconfig kubeconfig.yaml -n nginx-ingress patch helmrele
ases ingress-nginx -p '{"spec":{"values":null}}' --type=merge
```

2. When creating a new deployment of Cloudera DataFlow version 2.9.0 and below, customers need to repeat the mitigation steps outlined above once the DataFlow environment is successfully created.

Known issues and limitations

You must be aware of the known issues and limitations, the areas of impact, and workarounds in Cloudera DataFlow.

Known issues

CDPDFX-9476: PutCDPObjectStore error with CFM 1.24 on AWS Jakarta region

In Amazon AWS in the ap-southeast-3 (Jakarta) region, the PutCDPObjectStore NiFi processor throws an error in Cloudera DataFlow deployments running on NiFi 1.24. The issue does not manifest in NiFi 1.25 or higher.

Upgrade the deployment to NiFi 1.25.

Cloudera DataFlow Functions: DEFAULT_PARAM_CONTEXT variable no longer works alone

This issue only occurs in AWS environments and affects configurations where there are no PARAM_CONTEXT_ variables defined, only a DEFAULT_PARAM_CONTEXT.

The DEFAULT_PARAM_CONTEXT configuration variable instructs Cloudera DataFlow Functions which default secret to use when there is no secret that matches the parameter contexts in the flow. This variable is now ignored.

Create an environment variable in Configuration called PARAM_CONTEXT_*[***NAME***]* where *[***NAME***]* is the user-defined name of the parameter context. Specify the name of the AWS Secret you want to use as the value of this variable.

NiFi 2.0 [Technical Preview] deployments fail to obtain authentication token in RAZ enabled AWS environments

NiFi 2.0 [Technical Preview] deployments in RAZ enabled AWS environments are not able to obtain an authentication token and therefore fail to read/write to RAZ protected sources/destinations. NiFi 2.0 [Technical Preview] deployments without any Ranger RAZ dependencies work as expected.

There is no workaround for this issue.

IAM Policy Simulator preflight check fails with resource policy validation

With all cross account policies in place, IAM Policy Simulator preflight check still fails with the following error message:

```
IAM Resource Policy validation failed on AWS. CrossAccount role
does not have permissions for these operations : : ssm:GetParame
ter, ssm:GetParameters, ssm:GetParameterHistory, ssm:GetParamete
rsByPath
```

This happens because even if a given cross account role is allowed to perform a certain action (granted through IAM Policies), an attached Service Control Policy (SCP) may override that capability if it enforces a Deny on that action. SCP takes precedence over IAM Policies. SCPs are

either applied at the root of an organization, or can be applied to individual accounts. A permission can be blocked at any level above the account, either implicitly or explicitly (by including it in a Deny policy statement).

As the IAM Simulator SDK does not have an option to include or exclude an organization's SCP policy, the preflight check will fail if an SCP policy is denying an action, even though the IAM role has the necessary permissions.

This is a known issue in AWS.

Do not select the Skip Validations option when enabling Cloudera DataFlow to bypass this issue. This bypasses all preflight validation checks. Instead, submit a request to add the LIFTIE_DISABLE_IAM_PREFLIGHT_CHECK entitlement to your account which ensures only the IAM Policy preflight validation check is skipped.

Limitations

Parameter groups that have referencing flow drafts, and flow drafts referencing a parameter group cannot be reassigned to another project

A flow draft referencing any parameter group cannot be reassigned to another project.

Project reassignment does not move assets. When reassigning a parameter group that includes a FILE type parameter (asset) reference to another project, that asset needs to be re-uploaded to the new project.

There is no workaround for this issue.

Duplication of parameter groups referencing assets does not result in asset duplication

After duplicating a parameter group, assets have to be re-uploaded manually. This means that if you reassign a parameter group to another project, duplicate a parameter group, or export a parameter group, the asset will not be moved or duplicated.

Duplication can only happen inside a given project. If the newly created group is to be assigned to another project, that step can happen only after the duplication concludes. (Another project cannot be targeted for duplication.)

There is no workaround for this issue.

Diagnostic Bundle collection through the Management Console is available on the US Control Plane only

There is no workaround for this issue.

Data Lineage information is not automatically reported to Atlas in the Data Catalog

Flow deployments created by Cloudera DataFlow do not come with a pre-configured `ReportLineageToAtlas` Reporting Task.

If you have been assigned the DFFlowAdmin role, you can manually create and configure the `ReportLineageToAtlas` Reporting Task in the NiFi canvas after a deployment is completed.

PowerUsers are not able to create flow deployments without additional Cloudera DataFlow roles

While the PowerUser role gives you the ability to view flow deployments in the Dashboard, view flow definitions in the Catalog, and allows the user to initiate flow deployments, the Deployment Wizard fails after selecting an environment for which the user does not have the DFFlowAdmin resource role assigned.

Assign the DFFlowAdmin role to the user of the environment to which they want to deploy flow definitions.

Cloudera DataFlow reports "Bad Health" during Data Lake upgrade

Cloudera DataFlow monitors the state of the associated Cloudera on cloud environment to decide which actions Cloudera DataFlow users can take. Cloudera DataFlow detects Data Lake upgrades of the associated Cloudera on cloud environment and puts the Cloudera DataFlow service into Bad Health for the duration of the upgrade blocking new deployments.

To work around this issue, wait for the Data Lake upgrade to complete before creating new flow deployments.

Deployments and Cloudera DataFlow Services are no longer visible in the Cloudera DataFlow Dashboard or Environments page when the associated Cloudera on cloud Environment has been deleted

If the associated Cloudera on cloud Environment is deleted while a Cloudera DataFlow Service is enabled, it will become orphaned. Orphaned resources are no longer visible to users without the `PowerUser` role.

To work around this issue, open the Environments or Dashboard page with a user who has been assigned the `PowerUser` role. `PowerUsers` are able to view orphaned deployments and Cloudera DataFlow services.

Non-transparent proxies are not supported on Azure

There is no workaround for this issue.