



# Open Location Platform

## Release Notes

Version 1.5.1

# Important Information

## Notices

---

**Topics:**

This section contains document notices.

- [Legal Notices](#)
- [Document Information](#)

## Legal Notices

---

© 2018 HERE Global B.V. and its Affiliate(s). All rights reserved.

This material, including documentation and any related computer programs, is protected by copyright controlled by HERE. All rights are reserved. Copying, including reproducing, storing, adapting or translating, any or all of this material requires the prior written consent of HERE. This material also contains confidential information, which may not be disclosed to others without the prior written consent of HERE.

### Trademark Acknowledgements

HERE is trademark or registered trademark of HERE Global B.V.

Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

### Disclaimer

This content is provided "as-is" and without warranties of any kind, either express or implied, including, but not limited to, the implied warranties of merchantability, fitness for a particular purpose, satisfactory quality and non-infringement. HERE does not warrant that the content is error free and HERE does not warrant or make any representations regarding the quality, correctness, accuracy, or reliability of the content. You should therefore verify any information contained in the content before acting on it.

To the furthest extent permitted by law, under no circumstances, including without limitation the negligence of HERE, shall HERE be liable for any damages, including, without limitation, direct, special, indirect, punitive, consequential, exemplary and/ or incidental damages that result from the use or application of this content, even if HERE or an authorized representative has been advised of the possibility of such damages.

## Document Information

---

### Product

Name: Open Location Platform

Version: Version 1.5.1

### Document

Name: Open Location Platform Release Notes

ID: 9667afa-1524132572-827c58ac

Status: FINAL

Date: 2018-Apr-19, 10:10 (GMT)

---

# Contents

---

<b>Chapter 1: Overview</b> .....	<b>6</b>
Highlights.....	7
<b>Chapter 2: Changes</b> .....	<b>8</b>
New Features.....	9
Functional and Behavioral Changes.....	9
API Changes.....	9
Resolved Issues.....	10
Known Issues.....	11
Limitations and Workarounds.....	12

# Chapter 1

## Overview

---

### Topics:

- [Highlights](#)

This document provides a summary of important changes for this version. Apart from newly resolved issues and added or improved functionality, this document also includes information about known issues remaining in this release as well as any existing limitations and workarounds.

## Highlights

---

The Open Location Platform powers collaborative location intelligence for the next generation of products and services. The platform accelerates development of applications through Big Data processing and integrated access to rich location content and services.

Open Location Platform provides the following capabilities:

- **Data:** Store and access location related data, including built-in HERE map data, as well as your own ingested streams and batches of data.
- **Pipelines:** Develop and deploy processing logic that runs at scale, using Apache Flink for stream processing, Apache Spark for batch processing, or HERE data processing and location enrichment libraries for spatio-temporal processing.
- **Notebooks:** Perform ad-hoc analysis using Apache Zeppelin and HERE location data.
- **Schemas:** Declare the structure of your data, so that it can be used by others.

Developers can access and manage these capabilities in one of the following ways:

- **Portal:** Interactively manage Data and Pipelines, and access developer resources, such as the SDK and documentation, via a browser.
- **SDK:** Download libraries, tools, and examples to assist with local development of pipelines, and to transmit data to and from the Platform via external systems.
- **APIs:** Programmatically deploy and manage Platform resources such as Data and Pipelines through REST interfaces.

Developers can monitor and debug deployed pipelines, data ingestion, and other operational processes through:

- **Metrics:** See operational metrics in standard dashboard views.
- **Logs:** Search and analyze pipeline and process logs using a hosted Splunk instance.

Account Administrators can control Platform access for their account by managing Users, Groups, and their corresponding Permissions.

## Chapter 2

# Changes

---

### Topics:

- [New Features](#)
- [Functional and Behavioral ...](#)
- [API Changes](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Limitations and Workaround...](#)

This section documents major changes to the release that may require the updating of affected applications.



## New Features

---

This Open Location Platform release contains the following new functionality or improvements.

- **Data**
  - Notifications:
    - Programmatically subscribe to notifications for catalog version changes to find out when data in a catalog changes.
- **Pipelines**
  - Deploy, manage and monitor a pipeline via the Portal. This feature enables users to get started quickly without having to write any additional code to deploy and manage a pipeline via the API or CLI.
  - Delete Pipelines along with their Pipeline Versions and Pipeline Templates. Users can focus on the most important Pipelines and create CI/CD processes to save time.
  - System-generated version numbers for Pipeline Versions. This feature reduces the need for users to manually supply and maintain a version number for their pipelines.
  - View logs for each data processing job. This feature allows users to have granular details to aid with troubleshooting and improving their pipelines.
  - Updated CLI for Pipelines in the OLP SDK.
  - New Pipelines User Guide and an updated Pipelines Developer's Guide. This content helps users to learn about the pipeline concepts, their relationship to other OLP services, deployment and monitoring processes, as well as best practices.

## Functional and Behavioral Changes

---

The following list contains functional and behavioral changes in this current release:

- **Pipelines**
  - Ability to allocate the size of each Supervisor (Spark Driver or Flink JobManager) and Worker (Spark Executor or Flink TaskManager) for each Pipeline. You can allocate between 7 GB to 105 GB of RAM per Supervisor or Worker.
  - Ability to define a Pipeline's parallelism by choosing the number of Workers (Spark Executor or Flink TaskManager).
  - Force-run pipelines via the Pipeline API using existing catalog versions. This feature enables users to test pipelines without waiting for changes in the input catalog.
  - Any pipeline version that fails for 12 times consecutively will be de-activated.
  - View Pipelines' metrics within OLP dashboards powered by Grafana instead of Splunk.

## API Changes

---

### Added

- Pipeline Template

## ► Changes

- Re-usable objects that contain logic and runtime details. Most fields of the Template are derived from either Packages or Pipeline Version v1.
- These templates are immutable objects.
- New fields:
  - `runtimeEnvironment` - this parameter replaces 2 parameters in v1 - `Type` and `Framework`. The supported run-time environments are `Stream-1.5.0` for Flink (Stream) and `Batch-1.5.0` for Spark (Batch) pipelines.
  - `clusterConfiguration` - this set of 3 parameters in v2 replaces the 4 parameters that were used in v1 to specify the cluster configuration. More information about these parameters is available in documentation.
  - `inputCatalogIds` - the identifiers used in the JAR file for input catalogs are defined as part of the Pipeline Template. The actual catalog HRNs corresponding to these identifiers are defined only when a template is used in a Pipeline Version.

## Changed

- Package
  - All endpoints other than POST-ing a JAR file are now deprecated.
  - POST only takes the JAR file. All other fields available for a POST in v1 are not supported in v2.
  - The UUID returned should be used when you create the Pipeline Template.
  - Main class and other entry point classes defined in the manifest will also be returned (to be used when you create the Pipeline Template).
- Changed fields:
  - `defaultRuntimeConfiguration` - this parameter replaces the separate endpoint in v1 that was used to upload the run-time configuration (content of `application.properties` file) of the pipeline.
- Pipeline
  - Pipeline type (Streaming or Batch) and framework type (Flink or Spark) are moved from Pipeline to the Pipeline Template.
  - State and Status are removed. State is now supported at Jobs-level (individual execution instances of a Pipeline Version).
  - `LoggingURL` is moved to Jobs.
  - The flag that indicates whether a pipeline is scheduled to run on catalog changes or run immediately is moved to Pipeline Version.
  - Pipeline Delete, which deletes all Pipeline Versions and any related information.

## Removed

- None

## Resolved Issues

---

The following section summarizes major issues resolved in this release:

- Notebook permissions can be overridden by a non-owner.
- Pipelines cannot be created in the Portal UI.
- Pipelines can only use a limited amount of RAM.
- Pipelines can only be configured to run with a limited parallelism.

- Pipeline status seen in the Portal UI may get out of sync with the actual Pipeline status. This can occur if another User makes a change after the Pipeline details page loads. This issue could subsequently cause an error message to display if the User takes an action that is not valid for the current Pipeline status. For example, pausing a Pipeline that has already been paused by another User.
- An unhandled error in a User's Pipeline logic could cause a Pipeline to end up in a Failed state, without any automated alert or notification to the User.

## Known Issues

---

The following list contains issues known to be present in the current release:

- Notebooks currently use a snapshot release of Apache Zeppelin version 0.8 instead of a fully-supported release (version 0.7.3). This is because map visualization requires version 0.8 or newer. There may be stability concerns with initial Notebook release.
- Notebooks only support Python 3.
- Notebooks do not support Flink.
- Notebooks do not contain support for Streaming Layers.
- Notebook sessions are valid for one hour. After one hour, users will need to return to the Portal to refresh their session.
- Sample notebooks for the Location Libraries have not been updated for compatibility with version 1.5 of the libraries.
- Some traffic incidents aren't currently published in Real-Time Traffic's Incident Layer due to a missing translation between map referencing systems.
- HERE Reality Index Base catalog will stop updating after this release. Users should start using HERE Map Content catalog as soon as it becomes available to continue receiving incremental update.
- Changes from HERE Reality Index Base to HERE Map Content are:
  - Renaming of layers
  - Split of HERE Reality Index Base Road layer into
    - HERE Map Content: Road Attributes
    - HERE Map Content: Navigation Attributes
    - HERE Map Content: Advanced Navigation Attributes
  - Removal of following conditions from HERE Reality Index Base: Road
    - Truck-restrictions (future release candidate)
    - Environmental zones (future release candidate)
    - Junction Views (future release candidate )
    - Segment to Core Map Link references (new catalog available for early adopters)
  - Move of the following HERE Reality Index Base: Lane attributes to the HERE Map Content: Navigation Attributes layer
    - Lane Category
    - Through Lane Count
    - Physical Lane count
- OLP does not have full e2e data encryption in place, in-flight or at rest.
- Pipelines can't be Paused, Resumed, Upgraded from Portal, CLI, or API (V2).
- Pipelines can access the Internet outside of the platform. We don't officially support this behavior, and may prevent this in a future release. So, please take caution in connecting pipelines to external resources.
- The character-limit for Input Catalog and Output Catalog fields is 62. The Portal UI currently allows a user to specify more than 62 characters.

- When updating permissions, it can take up to an hour for changes to take effect.
- Portal and Notebooks are not compatible with Internet Explorer 11.
- Data visualization in the Portal is only available for versioned and volatile layers and only for data formatted according to SDII format, GeoJSON format or HERE Reality Index Topology layer format.
- Compute, Storage, and Transfer metrics are not fully available within Monitoring and Alerts.
- Create a Schema section in Data User Guide refers to deploy your created schemas to the Open Location Platform artifact repository. This feature was erroneously documented for 1.5.1.
- When creating input catalogs, if the user creates two catalogs with the same key and then deletes one before fixing the validation, the user won't be able to add any new catalogs.
- Sometimes, a pipeline could be saved without valid catalogs.

## Limitations and Workarounds

This section lists known limitations and, if available, related workarounds in the current release:

**Table 1: Limitations and Workarounds**

Limitation	Workaround
A finite number of access tokens (~250) can be generated for each app or user. This number reduces, depending on the number of resources included.	Create a new app or user if you reach that limitation.
Limited data encryption in the 1.5 release. This encryption is limited to data retained in the stream layer per the user's data retention configuration. There is no encryption for data at rest. Additionally, in-flight data is not encrypted end-to-end, throughout the service.	Data requiring encryption should not be stored inside the HERE Open Location Platform at this time.
Catalogs can be shared but there is not yet any management of permissions to enable a user to see with whom they have shared catalogs and to edit or revoke those permissions.	The User can edit/revoke permissions using Platform Portal. The user can track the permissions manually or contact Platform Customer Support to request the list of permissions.
Creating Catalogs and Layers using App credentials in the SDK does not automatically enable a user who created those credentials to discover, read, write, manage and share those resources.	Using the credentials with which the catalog was created, the catalog can be shared with users, apps or groups.
Deleting a catalog doesn't delete permissions associated with that catalog so if you create a new catalog using the same catalog name, whomever had permissions to the original/deleted catalog still has permissions to the new catalog.	If catalog permissions are a concern, use new names for new catalogs until this issue is resolved.
Within Notebooks, clicking on Spark Job while the job is running in Zeppelin opens new tab with unreachable IP. This is a known issue with Zeppelin and reverse proxies.	The user should not click on the given link.
Pipelines are not private to a User by default; they must be shared with exactly one Group when created.	Share the Pipeline with a Group containing only one User.
Pipelines can only be shared with specific Users if these Users belong to the same Group.	Share the Pipeline with a Group containing the specific set of Users.

# Open Location Platform Release Notes

► Changes



Limitation	Workaround
All Users and Apps in a Group are granted permissions to perform all actions on a Pipeline associated with that Group. There is no support for Users or Apps with limited permissions, for example, a reduced role that can only view Pipeline status, but not start and stop a Pipeline.	Limit the Users in a Pipeline's Group to only those Users who should have full control over the Pipeline.
Pipelines are not automatically granted permissions to read Data Catalogs that the User owns.	Data owners must share Data Catalogs with the Group selected when the Pipeline was created.
Pipeline Templates can't be deleted from the Portal UI.	Use CLI or API to delete Pipeline Templates.
A Batch Pipeline Version can't be forced to run from the Portal UI. It will run on its own when the input catalog version is updated.	Use API to force a Batch Pipeline Version to run or wait for the input catalogs' version to update. For more details, see Troubleshooting section in the Pipelines Developer's Guide.
In the Portal UI, a Pipeline Version is not automatically updated after the Pipeline Template has been created. Therefore, it takes some time for the Activate button for that Pipeline Version to get enabled.	Wait for a few minutes and refresh the Pipeline Versions page. The Activate button will get enabled.
A Batch Pipeline Version can't be forced to run from the Portal UI. It will run on its own when the input catalog version is updated.	Use CLI or API to force a Batch Pipeline Version to run or wait for the input catalogs' version to update. For more details, see Troubleshooting section in the Pipelines Developer's Guide.
The custom run-time configuration for a Pipeline Version has a limit of 64 characters for the property name, and 255 characters for the value.	Not Applicable.
From the Pipeline list in the Portal UI, clicking on the Output catalog for a pipeline doesn't take the user to the catalog information and instead shows the Pipeline's details page with its version list.	Open the Output catalog link in a new tab to view the details.
A running or scheduled pipeline can't be deleted. Also, if the user tries to delete a running pipeline/version from the Portal UI, the error message is not displayed.	Cancel or Deactivate the pipeline and then delete.
In the Portal UI, when viewing the Jobs and Operations for a Pipeline Version, the new Jobs and Operations are not automatically added to the list while the list is open.	Refresh the Jobs and Operations pages to see the latest job or operation in the list.