



Open Location Platform

Release Notes

Version 1.5.0

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 1.5.0

Document

Name: Open Location Platform Release Notes

ID: 1a771bf-1522784084-827c58ac

Status: FINAL

Date: 2018-Apr-03, 19:35 (GMT)

Contents

Chapter 1: Overview	6
Highlights.....	7
Chapter 2: Changes	8
New Features.....	9
Functional and Behavioral Changes.....	10
API Changes.....	11
Resolved Issues.....	11
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.

- **Notebooks**
 - Visually analyze polylines and GeoJSON-formatted data in Notebooks. For more information, see *Visualize Data* in the Notebooks User Guide.
 - Simplify data access in Notebooks through an updated Spark Data Service Connector. This version is compatible with the latest versions of the Data Service Client and Location Library, and removes the need to separately retrieve the user's access token to authenticate with the Data Service. For more information, see *Upgrade and Install Libraries* in the Notebooks User Guide.
- **Data**
 - Access essential HERE map data through the HERE Map Content catalog, which replaces the HERE Reality Index Base catalog. For details on this new catalog, see *Functional and Behavioral Changes* on page 10.
 - Access HERE Real Time Traffic catalog with global coverage.
 - Improved error messages, particularly for consistency.
 - Enabled encryption of stream data while stored on disk for the entire retention period.
 - Enabled TLS data encryption between Data Services.
 - Added reference example in Java for the Data Service Spark connectors for the Pipeline Service.
 - Added an API Lookup Service with the ability to get the base URLs for all Data Service APIs per catalog. For more information, see *Get API Base URLs* in the Data API Developer's Guide.
 - Added the ability to completely delete catalogs, including full data cleanup.
 - Updated the **Data List** view, resulting in changes in the Data landing page experience.
 - Added additional attributes on the **Layer Details** page: *layer version*, *content type*, *last modified date*, and *schema version*.
 - Added the ability to view a list of up to twenty thousand (20,000) partitions with their data size in bytes.
 - Added the ability to search for partitions by ID.
 - Improved the data decoding and download experience.
 - Added hint texts on the catalog and layer configuration pages.
- **Monitoring & Alerting**
 - Added the ability to create monitoring dashboards and configure alerts using the hosted Grafana tool.
 - Added new Data ingestion and access dashboards that you can use to monitor your data-related processes. This feature includes metrics for both levels: Catalog and Layer.
- **Documentation**
 - Added new User Guides for the following topics:
 - Data
 - Notebooks
 - Metrics and Logs
 - Setting Up Teams and Permissions
- **SDK**
 - Added Visualization Library to visually inspect your data and test the outputs from your data processing pipelines.

- Added the Open Location Platform Command Line Interface (OLP CLI) to interact with the Data Service.
- Location Library
 - Improved Java support: most of the `core` and `integration` packages provide a `javads1` sub-package meant to hold Java friendly wrappers on top of the Scala algorithms and data structures.
 - High Level API: ability to pass Data Client Library settings in `StandaloneCatalogFactory`.
 - Graph API: `outEdges` is `Iterable`; `RangeBasedPropertyMap` finds the last containing range by offset instead of the first one.
 - Map Matching API: new `TransitionProbabilityStrategies` factory methods provide state-of-the-art configurations for the underlying HMM-based map matcher, including `distanceWithPaths`, which is a pivotal strategy for dealing with sparse probes.
 - Added numerous new examples in either Java and/or Scala, including `SparsePathMatcherExample` to demonstrate how to reconstruct a partially matched path, as well as a story-telling example "Infer stop signs from sensors" for Spark.
- Added Java bindings for the Data Processing Scala Library
- Added a Data Validation Library version compare functionality
- Added a Java streaming archetype and a Scala `MapGroup` compiler archetype

Functional and Behavioral Changes

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

- **Changes from HERE Reality Index Base to HERE Map Content:**
 - Renaming of layers
 - Split of the 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
 - Environmental zones
 - Junction Views
 - Segment to Core Map Link references, which is a new catalog in preview 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
- **Notebooks**
 - Notebooks no longer require additional paragraphs to retrieve the user's access token for authentication to the Data Service. This functionality is now included in the Spark Data Service Connector.

API Changes

Added

- API Lookup REST API added.

Changed

- None

Removed

- None

Resolved Issues

The following section summarizes major issues resolved in this release:

- When you delete a catalog, all the data and metadata for that catalog is also deleted (in the Open Location Platform 1.0 release, only the Portal reference was removed).
- For stream messages less than 1 MB in size where data is stored in blob storage, that data is now deleted from Blob per Blob data retention settings (previously, it wasn't deleted).

Known Issues

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

- Currently, the Open Location Platform version 1.5.0 does not contain end-to-end data encryption in place, in-flight or at rest.
- Within the **Resources** tab, the **Getting Started** and **Documentation** links lead to the same content.
- 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, as well as for data formatted according to SDII format, GeoJSON format or HERE Reality Index Topology layer format.
- Notebooks currently use a snapshot release of Apache Zeppelin version 0.8 instead of a fully-supported release (version 0.7.2). This is because map visualization requires version 0.8 or newer. There may be stability concerns with the 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 Library have not been updated for compatibility with version 1.5 of the library.
- Compute, Storage, and Transfer metrics are not fully available within Monitoring and Alerts.

- When choosing to run a Batch Pipeline now, instead of scheduling it, the User must explicitly provide the Catalog version for the Pipeline to use as input.
- Some traffic incidents aren't currently published in the Real-Time Traffic's Incident Layer due to a missing translation between map referencing systems.
- When a user has their `readResource` permission to a catalog revoked, the user can continue to read data from the catalog until the permission cache has been refreshed.
- After this release, the HERE HERE Reality Index Base catalog will no longer be updated. This catalog is succeeded by the HERE Map Content catalog. To continue receiving incremental updates, use the HERE Map Content catalog as soon as it is released. For details on this new catalog, see [Functional and Behavioral Changes](#) on page 10.

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.
Although catalogs can be shared, currently it is not possible to manage the permissions to enable a user to see with whom they have shared catalogs, as well as to edit or revoke these permissions.	To request for permission revocation, contact the Technical Support team.
Creating Catalogs and Layers using App or Group 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 a Spark Job while it is running in Zeppelin opens a new tab with an unreachable IP. This is a known issue with Zeppelin and reverse proxies.	The user should not click on the given link.
Pipelines cannot be created in http://platform.here.com .	Create Pipelines through the Command Line Interface (CLI). Note: This limitation will be addressed in a subsequent release.
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.

Open Location Platform Release Notes

► Changes



Limitation	Workaround
<p>Pipelines can only be shared with specific Users if these Users belong to the same Group.</p> <p>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.</p>	<p>Share the Pipeline with a Group containing the specific set of Users.</p> <p>Limit the Users in a Pipeline's Group to only those Users who should have full control over the Pipeline.</p>
<p>Pipelines are not automatically granted permissions to read Data Catalogs that the User owns.</p>	<p>Data owners must share Data Catalogs with the Group selected when the Pipeline was created.</p>
<p>Pipelines can only use a limited amount of RAM.</p>	<p>If a pipeline requires access to more RAM, contact the Technical Support team.</p> <p> Note: This limitation will be addressed in a subsequent release.</p>
<p>Pipelines can only be configured to run with a limited parallelism.</p>	<p>If a pipeline requires a higher parallelism, contact the Technical Support team.</p> <p> Note: This limitation will be addressed in a subsequent release.</p>
<p>The pipeline status seen on http://platform.here.com 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.</p>	<p>Refresh the Pipeline details page in the Portal to see latest status.</p>
<p>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.</p>	<p>The User can set up an alert in Monitoring and Alerts to send a notification if a Pipeline fails.</p>