



Open Location Platform

Release Notes

Version 2.0.0

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 2.0.0

Document

Name: Open Location Platform Release Notes

ID: 1e5ab39-1538136349-827c58ac

Status: FINAL

Date: 2018-Sep-28, 12:06 (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.....	12
Known Issues.....	12
Limitations and Workarounds.....	13

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. For more information, see the [HERE Open Location Platform website](#).

Open Location Platform 2.0.0 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 (OLP) release contains the following new functionality or improvements:

HERE Marketplace

The inaugural release of the HERE Marketplace enables the secure, private exchange of location-related data between two parties.

In the Marketplace, data consumers can:

- Discover valuable location-related data for use in their own analyses or applications.
- Secure rights to license and use this data from HERE-approved data providers.
- Get programmatic access to this data via a standardized set of data APIs.

Data providers can use the Marketplace to:

- Create Listings that describe and promote their data. Data can be of 3 different types:
 - Stream - Stream of real-time, location-referenced event data.
 - Volatile - Real-time, current snapshot of frequently-changing data.
 - Versioned - Historical, location-referenced batches of data.
- Control who can discover their data by setting the visibility of a Listing—for all registered data consumers, or only for select data consumers.
- Control who has access to their data by explicitly granting entitlements to specific data consumers.
- View data consumption metrics.

NOTE: In order to get access to the Marketplace Listing features, existing OLP customers must subscribe to a Marketplace Provider plan.

Learn more about the Marketplace at <https://openlocation.here.com>, or by reading the documentation at <https://openlocation.here.com/resources/documentation>.

HERE Content

The HERE Map Content catalog has been updated with new layers, including:

- Truck Attributes (Premium Content) - Includes restrictions to the road network applicable to commercial vehicles and large trucks.
- 3D Buildings (Premium Content) - Includes three-dimensional models of buildings, and related attributes like height and ground clearance.
- Building Metadata (Standard Content) - Includes text descriptions of the categories that classify the buildings in the 3D Buildings layer.
- Cartography Metadata (Standard Content) - Includes text descriptions of the feature type that classify the cartographic features, and for display class.

In addition, the Building Footprints layer now contains 200M+ more buildings, which is four (4) to five (5) times more than before.

Catalog: HERE Map Content - External References

The HERE Map Content - External References catalog has been updated with a new field for road segment/link length.

Platform Status & Incident Reporting

View service availability and incident reports from the Status page at <https://status.openlocation.here.com>. This page shows current & historic availability of services, and allows you to confirm that platform services are healthy. This can be useful when debugging your processes, and to check status before reporting incidents to Technical Support.

Monitor Usage of Pipeline I/O

View and monitor usage of network I/O to and from Pipelines. These usage metrics have been added to the Current Usage dashboard in Grafana, and provides visibility into your applications' cost drivers.

Pipeline Logging-Level

Change the logging-level of a Pipeline Version using the Portal UI. Options include INFO, DEBUG, WARN, and ERROR modes, with or without restart. Enable DEBUG to troubleshoot Pipelines and related processes faster, and switch back to save logging costs, and to reduce log clutter.

Start & Stop Default Notebooks Cluster

The Spark cluster for Notebooks can be stopped and started from the new Cluster tab, accessible from the Portal UI > Notebooks. Manage costs to operate Notebooks by running the clusters only when you need them.

NOTE: Although HERE does not yet charge for the usage of Notebooks clusters, this will be activated in a future release.

Reduced Minimum Stream Time-To-Live (TTL)

Configure the TTL for a Stream Layer as low as 10 minutes, compared to the previous minimum TTL of 60 minutes. A longer TTL gives developers extra time to debug problems, but it also increases costs. A lower minimum TTL gives you more flexibility to control costs for when a longer TTL is not needed.

Enhanced Data Inspection for Traffic Flow, Building Footprints & Cartography Data

In the Data Inspect tab in the Portal UI, you can select visualized objects to expand the relevant section in the Decoded Data panel when viewing layers containing Traffic Flow, Building Footprints or Cartography data. This was previously only possible for Road Topology & Geometry data.

Search History in Visualizations

Find relevant data partitions faster, and with less manual input, by selecting from recently used partitions and lat/long coordinates. This is available in the Data Inspect tab in the Portal UI, and also in custom applications built using the Visualization Library in the SDK.

Version Picker in Visualizations

When browsing Versioned data, select different versions of layers—and list containing catalogs' dependencies—using the version selection control. This is available in the Data Inspect tab in the Portal UI, and also in custom applications built using the Visualization Library in the SDK. Use this to visualize historical versions of a data layer, to understand how the versions have changed over time. Understanding the dependency and version lineage of a Versioned Layer can help improve the debugging of data processing stages.

Functional and Behavioral Changes

- Notebooks

- ◻ After new releases, or after maintenance windows, the Spark cluster for Notebooks will be stopped by default. Before submitting any Spark jobs, restart the cluster using the new Cluster tab in Notebooks.
- Pipelines
 - ◻ Pipeline Templates and Pipeline Versions are now limited to a maximum of 200 CPU cores and 1.4 TB RAM. See more details in the Quotas and Limits section of the Pipeline User Guide document.
- Portal
 - Usage of Time Series data storage is no longer tracked or reported in the [Credit Usage](#) dashboard, or in the Current Usage dashboard in Grafana.
 - The Resources menu has moved. You can now find these links in the support panel.
 - The SDK download page has moved. You can now find this page in the 'Tools' section.
 - The documentation page has been re-categorized. All of the documents remain there, but have new grouping categories.

API Changes

Data APIs

- The Data Service Publish API has been updated to a new version (v2). It provides a common publication flow for all layer types, and therefore we have made changes to some terminology:
 - ◻ “batches” is now “publications”
 - ◻ “batchToken” is now “publicationId”
- With this release we are deprecating publish v1, so you should modify your applications to use the publish v2 API as soon as possible. To update your applications, make the following changes for the endpoints that you use.
 - ◻ GET /batches
 - This endpoint, which was used to search for active batches, is no longer exposed in v2.
 - ◻ POST /batches
 - This endpoint has been replaced with POST /publications.
 - The `catalogVersion` field has been removed from the request body. This change is made to simplify the publication. If you require version validation, check the `catalogVersion` value in the response body.
 - The new array `layerIds` defines the list of layers to be updated in the publication. You must add this to the request body.
 - ◻ GET /batches/{batchToken}
 - This endpoint has been replaced with GET /publications/{publicationId}.
 - ◻ PUT /batches/{batchToken}
 - This endpoint has been replaced with PUT /publications/{publicationId}.
 - ◻ DELETE /batches/{batchToken}
 - This endpoint has been replaced with DELETE /publications/{publicationId}.
 - ◻ POST /layers/{layerId}/publications
 - This endpoint has been replaced with POST /layers/{layerId}/publications/{publicationId}/partitions.

- The request body has been simplified to remove unnecessary elements and labels. You should now provide an array of partitions, which replaces `PublishPartition`.
- NOTE: We also recommend that you update to the latest version of the Data Client available with the OLP 2.0 release.
- For more information about the new publish v2 API, see the API Reference section of the [Data API Developer's Guide](#).

Resolved Issues

The following section summarizes major issues resolved in this release:

- Notebooks have been upgraded to the official 0.8 release of Apache Zeppelin.
- Pipelines:
 - Fixed: Duplicated input catalog selection on pipeline creation in Portal UI.
 - Fixed: LEVEL column in Splunk is not populated with the logging-level of the Pipeline.
 - Fixed: A pipeline version takes some time to reach the running state even though the state shows Running.
 - Fixed: Sometimes, pausing a Stream Pipeline using Portal UI does not work.
 - Fixed: Pipeline Version activation results in a wrong logging URL.
- HERE Reality Index Base catalog has been depreciated since July 2018

Known Issues

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

Some downloadable PDF documents have formatting errors.

Notebooks only support Python 3.

Notebooks do not support Flink.

Notebooks do not contain support for Streaming Layers.

- As previously communicated, the following deprecated catalogs will be removed:
 - `hrn:here:data:::here-optimized-map-for-location-libraries-1`
 - `hrn:here:data:::here-optimized-map-for-location-library-1-1`
- Within the **HERE Live Weather** catalogs (North America and Western Europe), the *Pavement Temperature* and *Pavement Conditions* attributes have been suppressed and will no longer be populated.
- Within the **HERE Archive Weather** catalogs (North America and Western Europe), the *Pavement Temperature* and *Pavement Conditions* attributes remain from when previously archived. But any future archiving ends at September 14, 2018 when the suppression of these attributes in HERE Live Weather begins.

Some traffic incidents aren't currently published in Real-Time Traffic's Incident Layer due to a missing translation between map referencing systems

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.

In the Pipeline API, the character-limit for inputCatalogHRN and outputCatalogHRN fields is 62. The Portal UI currently allows a user to specify more than 62 characters in the Pipeline Configuration page.

When updating permissions, it can take up to an hour for changes to take effect.

Portal UI and Notebooks are not compatible with Internet Explorer 11.

Data visualization in the Portal UI is only available for versioned and volatile layers. You are welcome to create your own visualization for your own schemas by creating a GeoJSON renderer and making it part of that schema. Out of the box, however, the only data that can be visualized is data formatted according to the following schema formats:

- GeoJSON
- HERE Reality Index Topology
- HERE Reality Index Building Footprints
- HERE Reality Index Cartography
- HERE Traffic Flow
- SDII

Compute, Storage, and Transfer metrics are not fully available within Monitoring and Alerts.

Catalogs not associated with a realm will not be visible in OLP.

The OLP SLA does not cover the ability for users to write Data to a versioned or stream layer.

Portal UI: the Pipelines list page might be slow to load all the elements.

- Upon activating a Pipeline Version, the corresponding job has a null state till it starts running.
- Portal UI: For a pipeline version running for more than a day, the elapsed time doesn't show the Day count.
- Portal UI: It is not possible to create a pipeline version if the billing tag was specified and then removed.
- Portal UI: While editing a pipeline's logging levels, the logger name in the table is inappropriately capitalized.

The system may not be able to detect the stream layer type and display the text pub-UNKNOWN on the Marketplace usage dashboard

Limitations and Workarounds

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

Table 1: Limitations and Workarounds

Limitation	Workaround
Marketplace users do not receive stream data usage metrics when reading data from Kafka direct.	Marketplace users must write data to a Stream layer using the Ingest REST API in order to receive usage metrics. Marketplace users must read data from a Stream layer using the Data Client Library configured for HTTP Connector in order to receive usage metrics.
When a Marketplace listing is published, the listing cannot be edited.	Users should deactivate the published listing and then create and publish a new listing.

Open Location Platform Release Notes

► Changes



Limitation	Workaround
When a Marketplace listing is deactivated from the Marketplace, it stays in the Marketplace Provider UI as inactive. Inactive listings cannot be permanently deleted.	To permanently delete unwanted listings, please contact technical support.
A Marketplace Provider can only send invitations to view listings during the listing creation.	Contact technical support to resend or send more invitations for a published listing.
When the Splunk server gets busy, it is possible for the server to lose usage metrics.	If users suspect they are losing usage metrics, contact technical support as soon as possible. We may be able to help rerun queries and validate data.
If you are a Workspace user and have pipelines directly connected to a stream, no messages-in and bytes-in metrics can be collected.	Users can instrument their pipelines with custom messages-in and bytes-in metrics. Contact technical support if needed.
After a Marketplace Consumer Manager has subscribed to data, if a Consumer Manager granted data access to their internal users, they cannot revoke access grant.	A Consumer Manager should first create a group, add users to the group and then grant data access to the group. A Consumer Manager can add and remove users from the group as needed.
After a Catalog is marked "Marketplace Ready", the Catalog may take up to 2 hours to show in the Provider Management Group's list of Catalogs.	Wait 2 hours. Contact technical support if a Provider Management Group cannot see a Marketplace Ready catalog after more than 2 hours has passed.
Occasional HTTP 5xx failures when attempting to read and write from a Volatile and Versioned Layer.	Retry the request.
Some older catalogs cannot be shared temporarily because the catalog owners don't have permissions assigned to them yet. This limitation only applies to older catalogs.	Report each occurrence as a bug.
Incomplete e2e in-flight data encryption: <ul style="list-style-type: none">In-flight data encryption is not consistently implemented across all component services.	Customers should first evaluate the risk associated with sending sensitive / personal information to OLP.
Creating Catalogs and Layers using App credentials via the Data API or Client Library 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.
A finite number of access tokens (~250) are available for each app or user. Depending on the number of resources included, this number may be smaller.	Create a new app or user if you reach the limitation.
Limited data encryption. Encryption is limited to: <ul style="list-style-type: none">Versioned data at rest for new versioned layersData retained in the stream layer per the user's data retention configuration (Time to Live)NotebooksIn-flight data is not encrypted end-to-end, throughout the service	Customers sending sensitive data or personal information to OLP at this time should opt out of doing so. Existing versioned catalog data will be encrypted in an upcoming OLP release with no additional work necessary by customers.
Creating Catalogs and Layers using App credentials via the Data API, the Client Library 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

Open Location Platform Release Notes

► Changes



Limitation	Workaround
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.
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 the 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 the API to force a Batch Pipeline Version to run or wait for the input catalogs' version to update. For more details, see the 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.	For the property value, there is no workaround. Users must stay within this limitation. For property names, users can define a shorter name in the configuration and map that to the actual, longer name within the pipeline code.
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.
Only a finite number of permissions are allowed for each app or user in the system across all services. It will be reduced depending on the inclusion of resources and types of permissions.	Delete Pipelines/Pipeline Templates to recover space.
The Pipeline Status Dashboard can be edited by users. Any changes made by the user will be lost when updates are published in future releases. Also, the dashboard will not be available for user edits in a future release.	Duplicate the dashboard or create a new dashboard.
Pipeline can still be activated after a catalog is deleted.	The pipeline will fail when it starts running and will show an error message about the missing catalog. Re-check the missing catalog or use a different catalog.