

Impact Study: Accelerating Interoperability with Palantir Foundry

How a U.S. Government agency leveraged Palantir Foundry's capabilities to facilitate enterprise connectivity

Copyright © 2022 Palantir Technologies Inc. and/or affiliates ("Palantir"). All rights reserved. The information in this document is proprietary and provided for informational purposes only and shall not create a warranty of any kind. Any data contained herein is notional.



Palantir Foundry – An Open and Interoperable Platform

Foundry can:

- Accelerate the creation of an enterprise data infrastructure and prevent lock-in
- Interoperate with existing and future data landscape systems

Introduction

Palantir Technologies Inc.’s (“Palantir,” “we” or “our”) Palantir Gotham Platform configured with Foundry (“Foundry”)—a FedRAMP Moderate-authorized Software as a Service (SaaS)—serves as the data infrastructure for hundreds of government, commercial, and non-profit organizations. Foundry supports critical missions across the U.S. Government (USG) as an enterprise data infrastructure. Foundry was built to connect an organization’s disparate systems and operations and offers a variety of capabilities for interoperability, such as data and logic storage in open formats, open APIs, ease of connection and export of data, and access to transformation logic. Importantly, our customers retain full control over their data and logic in Foundry: at any point in time, a customer can query, copy, export, delete, or revoke access to their data. In this document, we present a case study from a major USG agency that validated Foundry’s interoperability through testing performed by a third party.

Background

A large USG agency has an ever-growing data ecosystem comprising of disparate, heterogenous systems. The agency wanted to ensure access to and proper use of data across the organization to advance its mission. The agency’s siloed environment created an urgent need for interoperability between its data and analytics platforms. To help ensure connectivity across its IT landscape, the agency established an interoperability framework and processes to implement, operationalize, and maintain interoperable data pipelines.

As part of its efforts to advance interoperability, the agency conducted a Foundry proof-of-concept and sought to evaluate whether Foundry could meet the agency’s stringent interoperability requirements. The agency needed assurance that integrated data transformation pipelines configured in Foundry could be reused outside of the platform, should it choose to acquire Foundry to support its data and operational needs. This case study describes the evaluation process and how the agency came to recognize Foundry as a fully open and interoperable solution.

Overview

To rigorously validate that Foundry facilitates interoperability, the agency sought the help of a third-party software consultancy (“validator”) to perform testing of both Foundry and an open architecture component platform (“third-party platform”). The agency wanted to ensure Foundry and the platform can bidirectionally exchange data and information effectively. In order to avoid vendor lock-in, the agency wanted to confirm that the underlying data, the logic used to transform it, and the metadata used to describe it could be extracted and used in both platforms.¹ This process included two phases, described below.

PHASE 1

In Phase 1, the validator tested the exchange of data, code, and metadata from Foundry to the third-party platform. During Phase 1, Foundry’s APIs enabled the extraction of data and metadata, and facilitated rapid extraction of logic, alongside Foundry’s Git-backed logic repositories. The validator **proved Foundry’s interoperability** by:

- Successfully exporting a subset of the agency-specific configured data model logic and associated data and metadata from Foundry.
- Rebuilding and running the pipeline logic in the third-party platform.
- Confirming that the original logic written in Foundry and the logic rebuilt in the third-party platform produced the same datasets and results.

PHASE 2

In Phase 2, the validator tested the exchange of data and metadata from the third-party platform to Foundry. During Phase 2, the validator assessed that one-time extracts were successfully extracted from the third-party platform, which included data, its properties, and other metadata. This information was easily copied into Foundry through a configured data connector.² Unlike Foundry, the third-party platform did not support the export of logic during the assessment—only the export of data and metadata.

At the conclusion of this evaluation, Foundry was not only deemed open and interoperable, but its capabilities set the standard for interoperability by which all other IT investments at the agency are being evaluated.

¹ Code interoperability from the third-party platform to Foundry was not validated as it was challenging to extract a comprehensive code example from the third-party platform.

² The data connector can also support ongoing interoperability by ingesting either scheduled or triggered copies of datasets.

Foundry’s interoperability capabilities set the standard for interoperability evaluations at the agency.

Foundry Interoperability Testing: Objectives, Approach, and Outcomes

The interoperability evaluation tested Foundry’s ability to prevent lock-in as well as ensure flexibility and openness needed using the approach described below.

DATA AND METADATA MIGRATION

To prevent lock-in, customers can leverage Foundry’s:

- Comprehensive data lineage and provenance tracking
- Open APIs and API documentation

Objective: The agency wanted to ensure data, metadata, and schema can be copied from Foundry into another platform in a reproducible and consistent manner—including the ability to copy and track changes to the data and metadata in the third-party platform.

Approach: The validator tested Foundry’s ability to facilitate data and metadata migration with the following process:

1. The validator used PIV-derived Foundry user credentials in a tool built in the third-party platform to interface with Foundry REST APIs. These APIs can be used for data discovery, download, and storage.
2. Next, the third-party platform tool extracted a set of transactions of the relevant subset of modeled data and any source datasets.³
3. Using a Foundry API, the third-party platform tool then aggregated, staged, and compared every transaction file into one accessible file. This file could now be used for comparison between Foundry and the third-party platform outputs.
4. Finally, the validator used a Foundry API to retrieve unique and versioned metadata for the select subset of the ontology, source datasets, and relevant builds.

Outcome: The validator concluded that Foundry met the agency’s objective: Foundry enables effective migration of data, metadata, and logic out of the platform. The third-party platform successfully leveraged Foundry’s APIs to export the pipeline in its entirety, including relevant data and metadata.

³ Foundry’s API models data as part of a job in a build. The generation of a dataset produces a transaction of this dataset, which an external tool can extract.

PIPELINE REPRODUCTION USING MIGRATED TRANSFORMATION LOGIC

To prevent lock-in, the Agency leaned on Foundry's use of:

- Open-source languages for logic (such as Python, SQL, and Java)
- Git-backed repositories with open storage formats

Objective: The agency wanted to validate that logic in Foundry is consistently reproducible in other environments—such that, once copied out of the platform, the logic can a) be run, and b) outputs of the logic match those generated in Foundry.

Approach: To evaluate Foundry's logic interoperability, the validator took the following steps:

1. The validator copied the relevant subset of the Foundry ontology logic into a private repository within the agency's GitHub.⁴
2. Next, the validator copied source datasets and the logic repositories into the third-party platform.
3. The validator then used another tool in the third-party platform to read and execute the ontology logic subset as-is (to ensure maximal interoperability without requiring logic changes). The copied logic enabled the visualization of data flows, identification of source datasets, and the correct execution sequence for transforms in the pipeline in order to reconstruct it.
4. Finally, the validator leveraged developer documentation in consultation with Palantir to mirror libraries to load, run, and chain the ontology subset-related transforms to validate output consistency.

Outcome: The validator affirmed that Foundry pipeline logic can be:

- Exported to another platform; and
- Executed effectively and consistently.

Most importantly, the validator found no differences between the ontology subset re-generated in the third-party platform and the original Foundry ontology subset.

⁴ Since Foundry pipelines are written in an authoring environment backed by a secure Git repository, Foundry and any authorized external Git repository can be synced to capture changes in logic within Foundry.

Conclusion

A major U.S. Government agency acquired the services of an independent third party to conduct comprehensive testing of Foundry interoperability as part of a Foundry proof-of-concept. During testing, the third-party validator successfully exported data, code, and metadata into a third-party platform through open REST APIs and open Git-backed logic. The validator affirmed that the subset of the ontology exported into the third-party platform was readily and accurately reproducible. The agency therefore found Foundry to be open and interoperable, and its interoperability capabilities now serve as the standard by which all other IT investments at the agency are evaluated.