

# Workday Data Cloud

Product Summary

August 7, 2026



# Contents

|                                                                    |    |
|--------------------------------------------------------------------|----|
| Workday Data Cloud.....                                            | 3  |
| Setup Considerations: Workday Data Cloud.....                      | 4  |
| Data Out.....                                                      | 9  |
| Data Cloud Catalog.....                                            | 9  |
| Workday Live Data Query.....                                       | 11 |
| Get Started with Workday Live Data Query.....                      | 11 |
| Set Up Live Data Query ISU Authentication in Workday.....          | 12 |
| Set Up Live Data Query User-as-Self Authentication in Workday..... | 14 |
| Set Up Live Data Query in DBeaver with JDBC Driver.....            | 16 |
| Set Up Live Data Query Using Python Client and Driver.....         | 19 |
| Set Up Live Data Query in Tableau with JDBC Driver.....            | 21 |
| Set Up Live Data Query in Power BI Desktop.....                    | 24 |
| Generate Public and Private Keys for Workday Live Data Query.....  | 25 |
| Query Workday Data Using Live Data Query.....                      | 26 |
| Troubleshooting: Workday Live Data Query.....                      | 28 |
| Workday Data Lake.....                                             | 29 |
| Get Started with Workday Data Lake.....                            | 29 |
| Set Up the Data Cloud Functional Area.....                         | 30 |
| Set Up Integration System Users for Data Lake.....                 | 31 |
| Register API Client for Data Lake (JWT Bearer Grant).....          | 33 |
| Register API Client for Data Lake (Refresh Token Grant).....       | 34 |
| Set Up Security for Data Lake Admin Console.....                   | 35 |
| Add Tables and Columns to Data Lake.....                           | 36 |
| Configure Data Lake Access Roles.....                              | 37 |
| Query Workday Data Lake.....                                       | 38 |
| Data In.....                                                       | 40 |
| Get Started with Zero-Copy Data In.....                            | 40 |
| Set Up Workday Security Policies for Zero-Copy Data In.....        | 41 |
| Set Up Zero-Copy Data In with Databricks.....                      | 43 |
| Set Up Zero-Copy Data In with Google Cloud.....                    | 45 |
| Set Up Zero-Copy Data In with Salesforce.....                      | 49 |
| Set Up Zero-Copy Data In with Snowflake.....                       | 51 |
| Enable External Data for Analysis.....                             | 55 |
| Blend External Data with Workday Data.....                         | 55 |
| Create Visualizations.....                                         | 57 |
| Troubleshooting: Zero-Copy Data In.....                            | 58 |

# Workday Data Cloud

## Introduction

Note: You might need to take additional steps to enable this feature depending on your organization's subscription service agreement. For more information, see this [Community](#) article.

Note: Workday Data Cloud is only available in public-cloud-hosted Implementation tenants.

Workday Data Cloud facilitates bidirectional data exchange between the Workday platform and external enterprise data environments. Data Cloud utilizes open standards like Apache Iceberg and Trino to allow external systems to access Workday data - and Workday to access external system data - without requiring time-consuming extraction, transformation, and loading (ETL) processes.

## Overview

| Data Exchange Direction                    | Data Cloud Component                    | Connection             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------|-----------------------------------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Out: Workday Data to External Systems | <a href="#">Workday Live Data Query</a> | JDBC, Python, Power BI | <p>Use for targeted queries requiring real-time freshness.</p> <p>Available Data</p> <p>Selected Workday business objects exposed as SQL tables. See <a href="#">Data Cloud Catalog</a> .</p> <p>Access</p> <ul style="list-style-type: none"> <li>• ISU authentication for table-level &amp; column-level access controls.</li> <li>• User-as-self authentication for fully enforced Workday security model.</li> </ul> <p>Additional Features:</p> <ul style="list-style-type: none"> <li>• Live data.</li> <li>• Open connectivity from Data Platforms and Analytics Tools.</li> </ul> |

| Data Exchange Direction             | Data Cloud Component              | Connection                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------|-----------------------------------|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                     | <a href="#">Workday Data Lake</a> | Apache Iceberg                | <p>Use for large-scale analytics or data warehousing where periodic refresh is acceptable.</p> <p>Available Data</p> <p>Selected Workday business objects exposed as Iceberg tables. See <a href="#">Data Cloud Catalog</a> .</p> <p>Access:</p> <ul style="list-style-type: none"><li>• Table-level access controls.</li></ul> <p>Additional Features:</p> <ul style="list-style-type: none"><li>• Daily refresh.</li><li>• More performant high-volume raw extracts.</li><li>• Data platform connectivity.</li></ul> |
| Data In: External Data into Workday | <a href="#">Zero-Copy Data In</a> | Prism Data Catalog connection | <p>Use to access zero-copy data from external systems directly within Workday.</p> <p>Access:</p> <ul style="list-style-type: none"><li>• Prism security domains and roles.</li></ul> <p>Additional Features:</p> <ul style="list-style-type: none"><li>• Prism Data Management capabilities (Examples: Joins, calculated fields, data lifecycle management)</li></ul>                                                                                                                                                 |

For further information on Use Cases, Questions to Consider, and Security, see [Setup Considerations: Workday Data Cloud](#) .

## Setup Considerations: Workday Data Cloud

Note: Workday Data Cloud is only available in public-cloud-hosted Implementation tenants.

You can use this topic to help make decisions when planning your configuration and use of Workday Data Cloud.

It explains:

- Why to set it up.
- How it fits into the rest of Workday.
- Downstream impacts and cross-product interactions.
- Security requirements and business process configurations.
- Questions and limitations to consider before implementation.

Refer to detailed task instructions for full configuration details.

## What It Is

Workday Data Cloud is a product suite that facilitates data flow between Workday and external data platforms.

It utilizes a zero-copy architecture, which enables external systems to access Workday data objects without physical data extraction, transformation, and loading (ETL) processes.

The suite provides bidirectional access, allowing you to query Workday data from external platforms and view external data within Workday.

## Business Benefits

Workday Data Cloud addresses data access challenges related to extraction volume and data replication.

With Workday Data Cloud, you can:

- **Reduce integration maintenance:** Lower the cost of integrations by reducing the need for ETL tools, storage, and operational overhead through the use of standard connectors like JDBC and Apache Iceberg.
- **Accelerate data access:** Provide data teams with direct access to Workday data tables, reducing the time required to stage data for analytics.
- **Minimize data replication:** Keep data stored within the Workday secure environment while allowing virtual access, reducing the creation of redundant data copies across the enterprise. Minimum data replication minimizes security risk by eliminating redundant copies and ensuring consistent security models.
- **Support advanced analytics:** Provide scalable access to granular Workday data for use in external machine learning models and enterprise data lakes, accelerating AI initiatives and innovations.

## Use Cases

### Data Out

In these scenarios, data in the Workday Data Cloud Catalog is accessed from an external platform:

| Use Case                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Unified enterprise analytics | <p>Organizations can blend trusted Workday data with operational data in external warehouses (Examples: Snowflake, Databricks, Google Cloud Platform, and others) to gain cross-functional insights.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• <b>Retail performance:</b> A retailer can combine employee engagement data from Workday with store sales data in Snowflake to identify which teams drive the highest revenue.</li> <li>• <b>Operational efficiency:</b> A manufacturing plant manager can view Workday metrics (overtime, training compliance) alongside factory output data (from a Manufacturing Execution System) in a third-party BI tool like Power BI. This provides a holistic view of how staffing affects production safety and efficiency.</li> </ul> |

| Use Case                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Advanced AI and Machine Learning | <p>Data scientists require access to granular, historical data to train predictive models. Workday Data Cloud provides governed access to the Data Cloud Catalog without requiring file extraction.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• <b>Talent retention models:</b> Scientists can access detailed employee performance metrics, compensation changes, and career progression paths to build models that predict attrition or identify skill gaps.</li> <li>• <b>Workforce planning:</b> Organizations can leverage rich semantic context from Workday to power AI initiatives in external platforms.</li> </ul> |
| Self-Service data discovery      | <p>Workday Data Cloud democratizes access to data schemas, reducing reliance on IT to build custom reports.</p> <p>Example:</p> <ul style="list-style-type: none"> <li>• <b>Ad-hoc analysis:</b> Data analysts or HR Business Partners can browse the Data Cloud Catalog to understand data relationships and schemas. They can directly query objects, such as compensation history or training records, using standard SQL in their preferred tools.</li> </ul>                                                                                                                                                                              |

### Zero-Copy Data In

These scenarios use the data management tools in Workday Prism Analytics to access zero-copy external data within Workday.

| Use Case                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Financial reconciliation | <p>Finance teams often struggle to reconcile General Ledger (GL) accounts in Workday with high-volume transactions stored in external operational systems. Example:</p> <p><b>Automated matching:</b> Controllers can bring granular transaction data from external warehouses directly into Workday. This facilitates automated matching against the GL and creates an auditable trail from summary accounts back to source transactions.</p> |
| Compliance reporting     | <p>Data Subject Access Requests (DSAR) often require gathering data scattered across multiple systems. Example:</p> <p><b>Centralized retrieval:</b> Compliance officers can use Workday Data Cloud to pull activity logs or personal data stored in external warehouses (e.g., IT helpdesk tickets in Snowflake) to generate a complete compliance report for a specific employee.</p>                                                        |
| Talent scorecards        | <p>HR leaders can import business outcome data into Workday to evaluate employee impact beyond standard performance reviews. Example:</p> <p><b>Performance correlation:</b> By bringing in Customer Satisfaction (CSAT) scores from Salesforce or project profitability metrics from finance systems, organizations can link employee profiles directly to business outcomes for better succession planning.</p>                              |

## Questions to Consider

| Questions                                                                                 | Considerations                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Do you need real-time access or high-volume analytics?                                    | <p>Use <b>Workday Live Data Query</b> for transactional reads that require real-time data freshness.</p> <p>Use <b>Workday Data Lake</b> for large-scale analytics where periodic refreshes are acceptable.</p>                                                                                                                                                                                                                                                                                                                                                                                  |
| Does the external data require transformation?                                            | <ul style="list-style-type: none"> <li>• Use <b>Zero-Copy Data In</b> to reference or augment Accounting Center processes with external data that does not require complex manipulation.</li> <li>• Use <b>Prism Analytics</b> ingestion connectors when the use case requires multiple joins, heavy transformations, or large batch processing before use in Workday.</li> </ul>                                                                                                                                                                                                                |
| Which cloud infrastructure do you use?                                                    | Workday Data Cloud supports tenants running on the Public Cloud.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Which inbound connector is right for your use case: Prism Ingestion or Zero-Copy Data In? | <p>Use <b>Prism Ingestion Connectors</b> when you want to move and transform data into a centralized warehouse as the source-of-truth to enable high-performance, complex analytics. Best for high-performance analytics and complex data modeling where you need to move, transform, and store a permanent copy of the data locally.</p> <p>Use <b>Zero-Copy Data In</b> when you want a live virtual link to the source-of-truth data exactly where it lives, eliminating the need for expensive pipelines or duplicate storage. Best for real-time visibility and strict data governance.</p> |

## Recommendations

- **Evaluate data volume and complexity:** Use zero-copy access for simple joins or lookups.
- Continue to use traditional ingestion connectors if your use case involves complex multi-row transformations.
- **Plan for data responsibility:** While Workday secures data storage within Workday, you must configure security controls within the partner platform for data read out of Workday.
- **Utilize the Data Cloud Catalog:** Consult the [Data Cloud Catalog](#) for information about what objects are available.

## Requirements

Note: You might need to take additional steps to enable this feature depending on your organization's subscription service agreement. For more information, see this [Community](#) article.

- **Public Cloud:** Your Workday tenant must run on the Public Cloud.
- **Partner Contract:** You must hold a contract with a supported partner platform (Examples: Salesforce, Snowflake, Databricks, Google Cloud, and others).
- **Subscription:** You must subscribe to [Workday Flex Credits](#) to access Workday Data Cloud products.

## Limitations

- **Object availability:** Consult the [Data Cloud Catalog](#) for information about what objects are available.
- **Connectivity:** Zero-copy connectivity is currently limited to specific vendor partners and is not available for general-purpose JDBC connections to unsupported tools.

## Tenant Setup

You must configure your tenant to authorize connections with external partner platforms.

For an overview of setup instructions, see [Workday Data Cloud](#).

## Security

| Domains                                                                                                                                                  | Considerations                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <i>Integration Security</i> domain in the Integration functional area.<br><i>Security Configuration</i> domain in the System functional area.            | Enables administrators to configure Integration System Users.                                         |
| <i>Set Up: Tenant Setup - Security and Security Administration</i> domains in the System functional area.                                                | Enables administrators to register API Clients.                                                       |
| <i>Data Cloud: Manage Data Lake Access Policies</i> domain in the Data Cloud functional area.                                                            | Enables administrators to create access roles and user provisioning groups for Data Lake.             |
| <i>Data Cloud: Manage Data Lake Application Configurations</i> domain in the Data Cloud functional area.                                                 | Enables administrators to configure which Workday objects and fields are made available in Data Lake. |
| <i>Manage: Datalake User Provisioning</i> domain in the Data Cloud functional area.                                                                      | Enables administrators to create user provisioning groups for Data Lake.                              |
| <i>Prism: Manage Connection</i> domain in the Prism Analytics functional area.                                                                           | Enables administrators to create a connection in the Prism Data Catalog.                              |
| <i>Prism: External Catalog Create, External Catalog Manage, and Prism: External Catalog Owner Manage</i> domains in the Prism Analytics functional area. | Enables administrators to create and manage external catalogs in the Prism Data Catalog.              |
| <i>Prism: External Tables Manage and Prism: External Tables Owner Manage</i> domains in the Prism Analytics functional area.                             | Enables administrators to enable external tables for analysis in the Prism Data Catalog.              |

## Reporting

Workday Data Cloud allows you to:

- Query Workday data using standard SQL from external BI tools and data platforms.



- This replaces the need to create custom Reports-as-a-Service (RaaS) for some data extraction scenarios.
- Use Prism Analytics reporting and analytics tools to access blended external and Workday data.

## Integrations

Workday Data Cloud provides an alternative to traditional integrations.

| Integrations                   | Considerations                                                                                           |
|--------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Workday Live Data Query</b> | Provides real-time access to data via SQL.                                                               |
| <b>Workday Data Lake</b>       | Uses Apache Iceberg table formats to share data files without a need to persist data in other platforms. |
| <b>Zero-Copy Data In</b>       | Uses Prism Data Catalog for zero-copy access to external data within Workday.                            |

## Connections and Touchpoints

| Features                       | Considerations                                                                                                                   |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Workday Prism Analytics</b> | Use Prism Data Management tools for accessing zero-copy external data within Workday.                                            |
| <b>Accounting Center</b>       | Can ingest data via Prism connectors or reference external data via Data Cloud for journal entry generation.                     |
| <b>Workday Build</b>           | Workday Data Cloud acts as the data foundation for the Workday Build portfolio, supporting app development and interoperability. |

# Data Out

## Data Cloud Catalog

The Data Cloud Catalog contains information about the Workday business object tables available for Live Data Query and Data Lake. You can download the catalog in these formats:

- .xlsx: [Workday\\_Data\\_Cloud\\_Catalog\\_GA.xlsx](#)
- .pdf: [Workday\\_Data\\_Cloud\\_Catalog\\_GA.pdf](#)

The objects and fields included in the catalog represent a curated set of business objects and fields in Workday. You can make requests to add additional columns or tables to the catalog by contacting Workday Customer Care.

### How to read the catalog

There are 2 types of tables in the catalog:

| Table type                  | Description                                                                                                                                                       |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Core business object tables | Tables holding your organization's business data (Examples: Worker, position, candidate). They contain many fields and link to other tables through foreign keys. |

| Table type                       | Description                                                                                                                                                                                                  |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Reference business object tables | Lookup tables with minimal field exposure in the catalog (Examples: Country, currency, language). Core tables use foreign keys to join to reference tables and return readable values instead of stored IDs. |

For each table, we list the domain that governs access to the whole table. Some individual fields are secured by separate domains so one table can expose different fields to different users. Example: In the *worker* table, most fields are secured by the *Worker Data* domain while others are secured by the *Public Reporting Items* domain.

Each column in a table is described by these columns:

| Column      | Description                                                                                                                                                                                                                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data Type   | The SQL-compatible field type:<br><br>BIGINT: Whole number. Used for IDs and foreign keys.<br><br>ARRAY(BIGINT): List of numeric foreign keys for a repeating field.<br><br>VARCHAR: Text string.<br><br>DECIMAL (precision, scale): Decimal number with fixed precision and scale.<br><br>DATE: Calendar date.<br><br>TIMESTAMP: Date and time value.<br><br>BOOLEAN: True or false. |
| Key         | PK: Primary key.<br><br>FK: Single foreign key.<br><br>FK[ ]: Array of foreign keys.                                                                                                                                                                                                                                                                                                  |
| FK Target   | The table(s) a foreign key points to. In the PDF version of the catalog, the FK target is included in the Key column.                                                                                                                                                                                                                                                                 |
| Nullable    | Whether the column can be empty.                                                                                                                                                                                                                                                                                                                                                      |
| Description | The business meaning of the column, matching the Workday report-column definition.                                                                                                                                                                                                                                                                                                    |
| Domain      | The security domain(s) that govern access to the column.                                                                                                                                                                                                                                                                                                                              |

### ARRAY(BIGINT) columns

Workday models multi-instance relationships as array columns on the parent table, rather than as separate child tables with one row per value. Each element is a foreign key to the FK Target table. You can join to that table to return readable values instead of stored IDs.

For query examples, see [Query Workday Data Using Live Data Query](#) , and [Query Workday Data Lake](#) .

## Workday Live Data Query

### Get Started with Workday Live Data Query

#### Overview

Note: You might need to take additional steps to enable this feature depending on your organization's subscription service agreement. For more information, see this [Community](#) article.

Workday Live Data Query (LDQ) enables real-time SQL access from external systems to Workday's core business objects without using ETL or data replication. External clients query data within Workday on demand through a JDBC or Python driver.

#### Live Data Query Setup

##### Authentication

You can set up your LDQ connection with either of these authentication methods:

| Authentication Method         | Description                                                                                                                                                                                                                                                                                    | Setup Instructions                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| Integration System User (ISU) | Allows the external system to authenticate as a dedicated, non-person system account. Access to data is limited by the security permissions configured for the Integration System User (ISU), ensuring the connection can only access the specific tables and columns it is authorized to see. | <a href="#">Set Up Live Data Query ISU Authentication in Workday</a>          |
| User As Self                  | Allows the external system to read data on behalf of a Workday user. Enforces Workday's full security model, including row level controls and contextual security, based on the permissions of the user who authorizes the connection.                                                         | <a href="#">Set Up Live Data Query User-as-Self Authentication in Workday</a> |

##### External Client Setup

Live Data Query offers these combinations of external clients and drivers:

| External Client | Driver | Setup Instructions                                                     |
|-----------------|--------|------------------------------------------------------------------------|
| DBeaver         | JDBC   | <a href="#">Set Up Live Data Query in DBeaver with JDBC Driver</a>     |
| Python          | Python | <a href="#">Set Up Live Data Query Using Python Client and Driver</a>  |
| Snowflake       | Python | <a href="#">Set Up Live Data Query in Snowflake with Python Driver</a> |
| Power BI        | .mez   | <a href="#">Set Up Live Data Query in Power BI Desktop</a>             |

To download the drivers, see [Live Data Query Downloads](#).

Note: It's possible to use LDQ drivers in clients other than those documented here.

### Using Live Data Query

After you've established a data connection from the external client with the driver, you can run SQL statements to query Workday data. See [Query Workday Data Using Live Data Query](#).

### Guardrails and Limitations

- These guardrails and limitations are subject to change:

|                                               | Maximum Value        |
|-----------------------------------------------|----------------------|
| Query Runtime (Execution)                     | 30 minutes           |
| Query Runtime (Queueing, Execution, Data Out) | 60 minutes           |
| Query Concurrency (Per Tenant)                | 90 queries           |
| Query Concurrency (Per Environment)           | 100 queries          |
| Query Queue Length (Per Tenant)               | 800 queries          |
| Query Queue Length (Per Environment)          | 1000 queries         |
| Memory per Query                              | 20 GB                |
| Query Request Size (SQL Length)               | 1,000,000 characters |
| Query Response Size (Per Chunk)               | 16 MB                |
| Query Response Size (Total)                   | No limit             |
| Query Response Cells (Rows, Columns)          | No limit             |

- Available objects are currently limited to Workforce and Talent. See [Data Cloud Catalog](#).

## Set Up Live Data Query ISU Authentication in Workday

### Prerequisites

- Determine the required Integration System User (ISU) accounts you need for Live Data Query.

Things to consider:

- ISUs can be configured with unique table and column access controls based on their security group membership.
- A unique ISU will likely be required for each unique configurable security combination. Example: If all downstream systems require access to all data, one ISU is sufficient. If some use cases require general information (such as secured to 'Public Reporting') and others require privileged information, two or more ISUs will likely be necessary.
- For each ISU, determine the security group and which domains to grant. Use the Domain column in the [Data Cloud Catalog](#) to identify domains for the tables and fields the ISU must access. The required group type is *Integration System Security Group (unconstrained)*.
- Generate the public and private key. See [Generate Public and Private Keys for Workday Live Data Query](#).

## Context

Authenticating your Live Data Query connection with an ISU allows the external system to authenticate as a dedicated, non-person system account.

Access to data is limited by the security permissions configured for the ISU, ensuring the connection can only access the specific tables and columns it is authorized to see.

## Steps

1. Access the **Create Integration System User** task.

Security: *Integration Security* domain in the Integration functional area.

If an ISU has already been created for this purpose, skip this step.

For each ISU needed for your use case, complete the task:

| Option                               | Description                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------|
| User Name                            | Enter a unique ISU user name.                                                        |
| New Password<br>New Password Verify  | Set a password.                                                                      |
| Require New Password at Next Sign In | Don't select the check box.                                                          |
| Session Timeout Minutes              | Keep the value at zero to prevent the integration system user session from expiring. |
| Do Not Allow UI Sessions             | Select the check box.                                                                |

2. Click **OK**.

3. Access the **Create Security Group** task.

Security: *Security Configuration* domain in the System functional area.

If you're using an existing unconstrained integration system security group, skip this step.

For each ISU needed for your use case, complete the task:

| Option                          | Description                                                 |
|---------------------------------|-------------------------------------------------------------|
| Type of Tenanted Security Group | Integration System Security Group (Unconstrained)           |
| Name                            | Enter a unique group name based on your naming conventions. |
| Integration System Users        | Add the ISU user(s).                                        |

Click **OK**.

Note: If you make security group changes after the client connection has been established, subsequent permission changes might take up to 15 minutes to propagate before users can access newly authorized tables or columns.

4. For each ISU security group, grant the View Only report/task permission on the required domains. Use the Domain column in the [Data Cloud Catalog](#) to identify domains for the tables and fields the ISU must access.

To mass update security policies for a security group, run the **Maintain Permissions for Security Group** task. For details, see [Maintain Security Group Permissions](#).

5. Access the **Activate Pending Security Policy Changes** task.

Security: *Security Configuration* domain in the System functional area.

- a) Add an appropriate comment and click **OK**.
- b) Review the changes, check the **Confirm** check box, and click **OK**.

The security domain permissions and user security groups are now active.

6. Access the **Register API Client** task.

Security: *Set Up: Tenant Setup - Security* and *Security Administration* domains in the System functional area.

Complete the task:

| Option                      | Description                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client Name                 | Enter a unique API client name.                                                                                                                                                                                                                                                                                                                          |
| Client Grant Type           | Select Jwt Bearer Grant.                                                                                                                                                                                                                                                                                                                                 |
| x509 Certificate            | <p>If you have an existing x509 certificate, select it from the drop-down list.</p> <p>To create a new one, select the Create x509 Public key option from the drop-down list. Provide a name and add the public key.</p> <p>For information about the public key, see <a href="#">Generate Public and Private Keys for Workday Live Data Query</a> .</p> |
| Integration System User     | Add the ISU user.                                                                                                                                                                                                                                                                                                                                        |
| Scope                       | Select any value. This field is required but not utilized by LDQ ISU authentication.                                                                                                                                                                                                                                                                     |
| Include Workday Owned Scope | Select the check box.                                                                                                                                                                                                                                                                                                                                    |

7. Click **OK**.

The **Register API Client** task displays the API Client information.

8. Save the API Client information from the **Register API Client** page.

You'll need this information to configure Workday connector properties in the partner platform:

| Information             | Example or Details                                     |
|-------------------------|--------------------------------------------------------|
| Token Endpoint          | <code>https://{HOST1}/ccx/oauth2/{TENANT}/token</code> |
| Integration System User | demo-isu                                               |
| Client ID               | MThmMWI3YWEtZDZjMS00MzRmLTliZWBeMzg4ZjI2NmI1M          |

**Next Steps**

Set up Live Data Query in the partner system or analytics tool of your choice. See the External Client Setup section of [Get Started with Workday Live Data Query](#) .

**Set Up Live Data Query User-as-Self Authentication in Workday**

**Prerequisites**

Security: *Set Up: Tenant Setup - Security* and *Security Administration* domains in the System functional area.

## Context

The user-as-self authentication method enforces Workday's full security model, including row-level controls and contextual security, based on the permissions of the user who authorizes the connection.

## Steps

1. Access the **Register API Client** task in Workday.

Complete the task:

| Option                          | Description                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client Name                     | Enter a unique API client name.                                                                                                                                                                                                                                                                         |
| Client Grant Type               | Select <b>Authorization Code Grant</b> .                                                                                                                                                                                                                                                                |
| Access Token Type               | Select <b>Bearer</b> .                                                                                                                                                                                                                                                                                  |
| Redirection URI                 | <ul style="list-style-type: none"> <li>• For Power BI: <a href="https://oauth.powerbi.com/views/oauthredirect.html">https://oauth.powerbi.com/views/oauthredirect.html</a></li> <li>• For all other systems: <a href="https://localhost:8888/callback?">https://localhost:8888/callback?</a></li> </ul> |
| Refresh Token Timeout (in days) | Keep the default value.                                                                                                                                                                                                                                                                                 |
| Scope                           | Select the functional scopes you would like exposed in the catalog to users when they authenticate using user-as-self authentication. For a list of the distinct functional scopes required to provide full access, see the XLSX version of the <a href="#">Data Cloud Catalog</a> .                    |
| Include Workday Owned Scope     | Select the check box.                                                                                                                                                                                                                                                                                   |

2. Click **OK**.

The **Register API Client** task displays the API Client information.

3. Save the API Client information from the **Register API Client** page.

You'll need this information to configure Workday connector properties in the partner platform:

| Information            | Example or Details                                                                                |
|------------------------|---------------------------------------------------------------------------------------------------|
| Client ID              | {Generated client ID}                                                                             |
| Client Secret          | {Generated client secret}                                                                         |
| Token Endpoint         | <a href="https://{HOST1}/ccx/oauth2/{TENANT}/token">https://{HOST1}/ccx/oauth2/{TENANT}/token</a> |
| Authorization Endpoint | <a href="https://{HOST2}/{TENANT}/authorize">https://{HOST2}/{TENANT}/authorize</a>               |

## Next Steps

Set up Live Data Query in the partner system or analytics tool of your choice. See the External Client Setup section of [Get Started with Workday Live Data Query](#) .

## Set Up Live Data Query in DBeaver with JDBC Driver

### Prerequisites

- Either:
  - [Set Up Live Data Query ISU Authentication in Workday](#)
  - [Set Up Live Data Query User-as-Self Authentication in Workday](#)
- Download the latest DBeaver client from DBeaver.
- Download the JDBC driver .jar file and save the file to a local directory. See [Live Data Query Downloads](#).

### Context

Through a JDBC connection, Live Data Query enables SQL access from DBeaver to Workday business objects.

### Steps

1. In DBeaver, add a driver to the [Driver Manager](#).
  - a) Configure the below details in the **Settings** tab:  
Leave all other fields in the **Settings** tab as default.
  - b) In the **Libraries** tab, upload the .jar file you downloaded in the Prerequisites section.
  - c) Click **Find Class** to verify the driver class is detected.  
Ensure `com.workday.dataservice.driver.api.DataServiceDriver` appears in the driver class dropdown.
2. Create a new [database connection](#) using the driver you created above.
  - a) In the **Connection Settings**, select **Main > Connect by URL**, and enter details relevant to the authorization flow you're using:

|          | ISU                                                                                                                                                   | User-As-Self                                                                                                                                                                                                                                                           |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JDBC URL | <code>jdbc:workday://{HOST}:443</code><br><br>Replace {HOST} with the {HOST1} value from the Token Endpoint created in the Live Data Query ISU setup. | <code>jdbc:workday://{HOST}:{PORT}?SSLVerification=NONE</code><br><br>Replace: <ul style="list-style-type: none"> <li>• {HOST} with the {HOST1} value from the Token Endpoint created in the User-As-Self Authentication setup.</li> <li>• {PORT} with 443.</li> </ul> |

- b) On the **Driver Properties** tab, add the properties relevant to the authorization flow you're using:

| Property                        | ISU                                                 | User As Self                                                |
|---------------------------------|-----------------------------------------------------|-------------------------------------------------------------|
| <code>wd.authn.authModel</code> | JWT_BEARER                                          | AUTHORIZATION_CODE                                          |
| <code>wd.authn.clientId</code>  | Client ID created in the Live Data Query ISU setup. | Client ID created in the User-As-Self Authentication setup. |



| Property                       | ISU                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | User As Self                                                                                                                |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| wd.authn.clientSecret          | N/A                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Client secret created in the User-As-Self Authentication setup.                                                             |
| wd.authn.isu                   | Your ISU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | N/A                                                                                                                         |
| wd.authn.accessTokenEndpoint   | Token endpoint created in the Live Data Query ISU setup.<br><br>Example: https://{HOST1}/ccx/oauth2/{TENANT}/token                                                                                                                                                                                                                                                                                                                                                                                                                                               | Token endpoint created in the User-As-Self Authentication setup.<br><br>Example: https://{HOST1}/ccx/oauth2/{TENANT}/token  |
| wd.authn.authorizationEndpoint | N/A                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Authorization endpoint created in the User-As-Self Authentication setup.<br><br>Example: https://{HOST2}/{TENANT}/authorize |
| wd.authn.privateKeyFilePath    | <p>The full path of the private key file.</p> <ul style="list-style-type: none"> <li>Windows Example:<br/>C:\Users<br/>  \user.name<br/>  \Documents<br/>  \&lt;privateKeyFile.txt&gt;</li> <li>Mac Example: /<br/>Users/path/<br/>to/your/<br/>private_key.pem</li> </ul> <p>Note:<br/>For a Windows machine, you must specify the wd.authn.privateKeyFilePath property.</p> <p>The wd.authn.privateKey and wd.authn.privateKeyFilePath properties are mutually exclusive. If you specified wd.authn.privateKey, don't specify wd.authn.privateKeyFilePath.</p> | N/A                                                                                                                         |
| wd.authn.privateKey            | <p>Note:<br/>The wd.authn.privateKey and wd.authn.privateKeyFilePath properties are mutually exclusive. If you specified</p>                                                                                                                                                                                                                                                                                                                                                                                                                                     | N/A                                                                                                                         |

| Property             | ISU                                                                                                                                                                                                                                                                                                                                                              | User As Self                    |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
|                      | <p>wd.authn.privateKey, don't specify wd.authn.privateKeyFilePath.</p> <p>We recommend using wd.authn.privateKeyFilePath for LDQ in DBeaver.</p> <p>Privacy-Enhanced Mail (PEM) format text. Must include the header/footer.</p> <p>For information about the private key, see <a href="#">Generate Public and Private Keys for Workday Live Data Query</a>.</p> |                                 |
| wd.authn.redirectURL | N/A                                                                                                                                                                                                                                                                                                                                                              | https://localhost:8888/callback |

- c) Note: If you use a Single Sign-On, you must log in to your Workday tenant using your default browser before you attempt to connect. This only applies to user-as-self authentication.

Click **Test Connection** to verify the configuration.

A successful test results in a Connected message. When there are errors, verify that the driver properties and connection settings are configured correctly.

- d) Save the connection.

### 3. Connect (User-As-Self Authentication Only)

- a) Double-click the new connection.

DBeaver will trigger the OAuth flow, automatically opening your default browser on the Workday login page.

- b) Log in with your Workday credentials and approve access.

The browser redirects to `https://localhost:8888/callback`.

- c) A browser security warning may appear because the callback server uses a self-signed certificate. Accept the warning to proceed (this is expected — the certificate is local to your machine).

Note: You have 5 minutes to complete the browser login. If you exceed this, the connection attempt fails and you can retry.

- d) Once the browser shows a success message, return to DBeaver. The connection completes automatically.

### 4. Execute a Query

- a) In the **Database Navigator**, expand the new connection.

- b) Right-click the connection and select **SQL Editor > New SQL Script**.

- c) Enter the query in the SQL Editor and click **Execute** to run it.

## Next Steps

See [Query Workday Data Using Live Data Query](#).

## Set Up Live Data Query Using Python Client and Driver

### Prerequisites

- Either:
  - [Set Up Live Data Query ISU Authentication in Workday](#)
  - [Set Up Live Data Query User-as-Self Authentication in Workday](#)
- Download and install Python (v3.9 or above).
- Download the Live Data Query Python wheel file. See [Live Data Query Downloads](#).

### Context

You can connect to Live Data Query using a Python script.

### Steps

1. Create the `dataservice.properties` file with the following configuration details using the `property=value` format. For sample code, see the Example section below.

We recommend that you place the properties file and the connector wheel file in the same Python project directory before installation and execution.

| Property                                    | ISU                                                                                                                                                                                                                                                     | User-As-Self                                                                                                                                                                                                                                                    |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>wd.authn.authModel</code>             | JWT_BEARER                                                                                                                                                                                                                                              | AUTHORIZATION_CODE                                                                                                                                                                                                                                              |
| <code>wd.authn.clientId</code>              | Client ID created in the Live Data Query ISU setup.                                                                                                                                                                                                     | Client ID created in the User-As-Self Authentication setup.                                                                                                                                                                                                     |
| <code>wd.authn.clientSecret</code>          | N/A                                                                                                                                                                                                                                                     | Client secret created in the User-As-Self Authentication setup.                                                                                                                                                                                                 |
| <code>wd.authn.isu</code>                   | Your ISU.                                                                                                                                                                                                                                               | N/A                                                                                                                                                                                                                                                             |
| <code>wd.authn.accessTokenEndpoint</code>   | Token endpoint created in the Live Data Query ISU setup.<br><br>Example: <code>https://{HOST1}/ccx/oauth2/{TENANT}/token</code><br><br>If the client is behind a proxy server, specify <code>wd.authn.accessTokenEndpoint=&lt;proxy-server&gt;</code> . | Token endpoint created in the User-As-Self Authentication setup.<br><br>Example: <code>https://{HOST1}/ccx/oauth2/{TENANT}/token</code><br><br>If the client is behind a proxy server, specify <code>wd.authn.accessTokenEndpoint=&lt;proxy-server&gt;</code> . |
| <code>wd.authn.authorizationEndpoint</code> | N/A                                                                                                                                                                                                                                                     | Authorization endpoint created in the User-As-Self Authentication setup.<br><br>Example: <code>https://{HOST2}/{TENANT}/authorize</code>                                                                                                                        |
| <code>wd.authn.privateKeyFilePath</code>    | The full path of the private key file.<br><br><ul style="list-style-type: none"> <li>• Windows Example: <code>C:\Users\user.name</code></li> </ul>                                                                                                      | N/A                                                                                                                                                                                                                                                             |

| Property             | ISU                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | User-As-Self                                                                                                                                                                                            |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | <p>\Documents<br/>\&lt;privateKeyFile.txt&gt;</p> <ul style="list-style-type: none"> <li>Mac Example: /<br/>Users/path/<br/>to/your/<br/>private_key.pem</li> </ul> <p>Note:</p> <p>For a Windows machine,<br/>you must specify the<br/>wd.authn.privateKeyFilePath<br/>property.</p> <p>The<br/>wd.authn.privateKey and<br/>wd.authn.privateKeyFilePath<br/>properties are mutually<br/>exclusive. If you specified<br/>wd.authn.privateKey,<br/>don't specify<br/>wd.authn.privateKeyFilePath.</p> |                                                                                                                                                                                                         |
| wd.authn.privateKey  | <p>Note:</p> <p>The<br/>wd.authn.privateKey and<br/>wd.authn.privateKeyFilePath<br/>properties are mutually<br/>exclusive. If you specified<br/>wd.authn.privateKey,<br/>don't specify<br/>wd.authn.privateKeyFilePath.</p> <p>Privacy-Enhanced Mail (PEM)<br/>format text. Must include the<br/>header/footer.</p> <p>For information about the<br/>private key, see <a href="#">Generate<br/>Public and Private Keys for<br/>Workday Live Data Query</a> .</p>                                     | N/A                                                                                                                                                                                                     |
| wd.authn.redirectURL | N/A                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | https://<br>localhost:8888/<br>callback                                                                                                                                                                 |
| wd.host              | <p>The {HOST1} value from the<br/>Token Endpoint on the API<br/>Client you created in the ISU<br/>setup.</p> <p>If the client is behind a<br/>proxy server, specify<br/>wd.host=&lt;proxy host&gt;.</p>                                                                                                                                                                                                                                                                                              | <p>The {HOST1} value from the<br/>Token Endpoint on the API<br/>Client you created in the ISU<br/>setup.</p> <p>If the client is behind a<br/>proxy server, specify<br/>wd.host=&lt;proxy host&gt;.</p> |
| wd.port              | 443                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 443                                                                                                                                                                                                     |

2. Install the Python library (wheel file) using the following command:

```
pip3 install <wheel.whl>
```

3. Create a Python script that loads the configuration from the `dataservice.properties` file, creates a connection to the service, and executes the query.

Example:

```
from workday_ldq import DataServiceConfig, create_connection

# Load configuration from current directory
config = DataServiceConfig.from_file("dataservice.properties")

# Create connection
connection = create_connection(config)

# Execute the Query
cursor = connection.cursor()
cursor.execute("SELECT * FROM {tablename} LIMIT 10")
results = cursor.fetchall()
cursor.close()
connection.close()
```

### Example

dataservice.properties Sample Code

| ISU Authentication                                                                                                                                                                                                                                   | User-As-Self Authentication                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>wd.authn.authModel=JWT_BEARER wd.authn.clientId=your_client_id_here wd.authn.isu=ISU_username_here wd.authn.accessTokenEndpoint=https://{HOST1}/ccx/oauth2/{TENANT}/token wd.authn.privateKeyFilePath=/Users/path/to/your/private_key.pem</pre> | <pre>wd.authn.authModel=AUTHORIZATION_CODE wd.authn.clientId=your_user_client_id_here wd.authn.clientSecret=your_client_secret_here wd.authn.accessTokenEndpoint=https://{HOST1}/ccx/oauth2/{TENANT}/token wd.authn.authorizationEndpoint=https://{HOST2}/{TENANT}/authorize wd.authn.redirectUrl=https://localhost:8888/callback</pre> |

### Next Steps

See [Query Workday Data Using Live Data Query](#) .

## Set Up Live Data Query in Tableau with JDBC Driver

### Prerequisites

- Either:
  - [Set Up Live Data Query ISU Authentication in Workday](#)
  - [Set Up Live Data Query User-as-Self Authentication in Workday](#)
- Download the JDBC driver .jar file and save the file to a local directory. See [Live Data Query Downloads](#).

### Context

Through a JDBC connection, Live Data Query enables SQL access from Tableau to Workday business objects.

## Steps

1. Install the JDBC Driver in the Tableau Directory.

For Tableau Desktop, copy the JDBC driver .jar file to:

- macOS : ~/Library/Tableau/Drivers/
- Windows : C:\Program Files\Tableau\Drivers\

If the Drivers directory does not exist, create it.

2. (For Linux only) Set the file permissions.

Ensure Tableau has permission to read the driver file:

```
chmod 755 ~/Library/Tableau/Drivers/<driver>.jar
```

3. Create a dataservice.properties file with the following configuration details using the property=value format. For sample code, see the Example section below.

| Property                       | ISU                                                                                                                                                                                                                                                                                             | User-As-Self                                                                                                            |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| wd.authn.authModel             | JWT_BEARER                                                                                                                                                                                                                                                                                      | AUTHORIZATION_CODE                                                                                                      |
| wd.authn.clientId              | Client ID created in the Live Data Query ISU setup.                                                                                                                                                                                                                                             | Client ID created in the User-As-Self Authentication setup.                                                             |
| wd.authn.clientSecret          | N/A                                                                                                                                                                                                                                                                                             | Client secret created in the User-As-Self Authentication setup.                                                         |
| wd.authn.isu                   | Your ISU.                                                                                                                                                                                                                                                                                       | N/A                                                                                                                     |
| wd.authn.accessTokenEndpoint   | Token endpoint created in the Live Data Query ISU setup.<br>Example: https://{HOST1}/ccx/oauth2/{TENANT}/token                                                                                                                                                                                  | Token endpoint created in the User-As-Self Authentication setup.<br>Example: https://{HOST1}/ccx/oauth2/{TENANT}/token  |
| wd.authn.authorizationEndpoint | N/A                                                                                                                                                                                                                                                                                             | Authorization endpoint created in the User-As-Self Authentication setup.<br>Example: https://{HOST2}/{TENANT}/authorize |
| wd.authn.privateKeyFilePath    | The full path of the private key file.<br><ul style="list-style-type: none"> <li>• Windows Example: C:\Users\user.name\Documents\&lt;privateKeyFile.txt&gt;</li> <li>• Mac Example: /Users/path/to/your/private_key.pem</li> </ul> <p>Note:<br/>For a Windows machine, you must specify the</p> | N/A                                                                                                                     |

| Property                          | ISU                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | User-As-Self                                 |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
|                                   | <p><code>wd.authn.privateKeyFilePath</code> property.</p> <p>The <code>wd.authn.privateKey</code> and <code>wd.authn.privateKeyFilePath</code> properties are mutually exclusive. If you specified <code>wd.authn.privateKey</code>, don't specify <code>wd.authn.privateKeyFilePath</code>.</p>                                                                                                                                                                    |                                              |
| <code>wd.authn.privateKey</code>  | <p>Note:</p> <p>The <code>wd.authn.privateKey</code> and <code>wd.authn.privateKeyFilePath</code> properties are mutually exclusive. If you specified <code>wd.authn.privateKey</code>, don't specify <code>wd.authn.privateKeyFilePath</code>.</p> <p>Privacy-Enhanced Mail (PEM) format text. Must include the header/footer.</p> <p>For information about the private key, see <a href="#">Generate Public and Private Keys for Workday Live Data Query</a>.</p> | N/A                                          |
| <code>wd.authn.redirectURL</code> | N/A                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <code>https://localhost:8888/callback</code> |

Note: If the client is behind a proxy server, add the `wd.http-proxy=<proxy-host>` property in the `dataservice.properties` file.

4. Launch Tableau Desktop.
5. [Configure the JDBC Connection in Tableau](#).
  - a. Open Tableau Desktop.
  - b. Click **Connect > To a Server > Other Databases (JDBC)**.
  - c. Enter the following information:

| Option   | Description (ISU and User-As-Self)                                                                                                                                |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL      | <p><code>jdbc:workday://{HOST}:443/workday_core?</code></p> <p>Note: Replace {HOST} with the {HOST1} value from the ISU or User-As-Self Authentication setup.</p> |
| Dialect  | Leave it default                                                                                                                                                  |
| Username | Leave it blank                                                                                                                                                    |
| Password | Leave it blank                                                                                                                                                    |

| Option          | Description (ISU and User-As-Self)                             |
|-----------------|----------------------------------------------------------------|
| Properties File | Upload the dataservice.properties file using the Browse button |

d. Click **Sign In** to open the Application.

### Example

dataservice.properties Sample Code

| ISU Authentication                                                                                                                                                                                                                                   | User-As-Self Authentication                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>wd.authn.authModel=JWT_BEARER wd.authn.clientId=your_client_id_here wd.authn.isu=ISU_username_here wd.authn.accessTokenEndpoint=https://{HOST1}/ccx/oauth2/{TENANT}/token wd.authn.privateKeyFilePath=/Users/path/to/your/private_key.pem</pre> | <pre>wd.authn.authModel=AUTHORIZATION_CODE wd.authn.clientId=your_user_client_id_here wd.authn.clientSecret=your_client_secret_here wd.authn.accessTokenEndpoint=https://{HOST1}/ccx/oauth2/{TENANT}/token wd.authn.authorizationEndpoint=https://{HOST2}/{TENANT}/authorize wd.authn.redirectUrl=https://localhost:8888/callback</pre> |

### Next Steps

See [Query Workday Data Using Live Data Query](#) .

## Set Up Live Data Query in Power BI Desktop

### Prerequisites

Note:

At this time:

- SSO-enabled tenants can't connect to LDQ from Power BI.
- ISU authentication isn't supported with Power BI.
- [Set Up Live Data Query User-as-Self Authentication in Workday](#)
- Contact Workday Support to get access to the .mez file required in this procedure.

### Steps

1. Right-click the zip containing the .mez file and select **Extract All**.  
Extract the whole folder somewhere easy to find, such as your Desktop. Don't run the files from inside the zip preview. Always extract the contents first.
2. Open config.txt from the extracted file in Notepad.
  - a) Enter these values after the = signs (don't use quotes):

| Property                     | Value                                                            |
|------------------------------|------------------------------------------------------------------|
| wd.authn.clientId            | Client ID created in the User-As-Self Authentication setup.      |
| wd.authn.clientSecret        | Client secret created in the User-As-Self Authentication setup.  |
| wd.authn.accessTokenEndpoint | Token endpoint created in the User-As-Self Authentication setup. |



| Property                                    | Value                                                                                                                                   |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
|                                             | Example: <code>https://{HOST}/ccx/oauth2/{TENANT}/token</code>                                                                          |
| <code>wd.authn.authorizationEndpoint</code> | Authorization endpoint created in the User-As-Self Authentication setup.<br><br>Example: <code>https://{HOST}/{TENANT}/authorize</code> |

- b) Save and close the file.
3. In the extracted file, double-click `Build-Connector.bat`.  
  
A window opens, builds the connector, and installs it automatically. When it's finished, close the window. The connector is installed to: `Documents\Microsoft Power BI Desktop\Custom Connectors\WorkdayLDQ.mez`.  
  
If Windows:
  - Shows a security error, run the `.bat` file anyway. This is expected behavior.
  - Can't copy the connector, close Power BI Desktop completely and run the `.bat` file again.
4. [Allow uncertified custom connectors](#) in Power BI Desktop.
  - a) In Power BI, select **File > Options**.
  - b) Select **Settings > Options > Security > Data Extensions**.
  - c) Select **Allow any extension to load without validation or warning**. This is necessary because this is an unsigned custom connector file.
  - d) Click **OK**.
  - e) Close Power BI Desktop completely and reopen it.
5. [Connect and browse the catalog](#).
  - a) In Power BI Desktop, select **Get data from other sources**.
  - b) Search for Workday LDQ and click **Connect**.
  - c) Sign in with your Workday account when prompted.
  - d) Click **Allow** when prompted.
  - e) The Navigator opens. Browse `workday_core > public`, select the tables you want, and click **Load**.

## Generate Public and Private Keys for Workday Live Data Query

Public and private keys are required when you configure Workday Live Data Query and client connector properties:

- The x509 public key is required when you set up Live Data Query in Workday, specifically when you register the API Client for interacting with Workday.
- The private key is required when you configure Workday connector properties in your client (such as DBeaver, Tableau, and Python).

Ensure that you keep the keys protected.

### Running the OpenSSL Command

To generate the public and private key, run the `openssl` command from the terminal of your choice.

To check if your machine has `openssl`, run this command and confirm it prints a version:

```
openssl --version
```

If you're using a Windows machine and you don't have `openssl`, use the Windows Package Manager (winget) to install `openssl`. Example:

```
winget install -e --id ShiningLight.OpenSSL
```

The `openssl` command is pre-installed with the macOS.

This sample command generates the key in PEM format:

```
openssl req -x509 -newkey rsa:2048 -keyout private-key.pem -out public-key.cert -sha256 -days 365 -nodes
```

## Query Workday Data Using Live Data Query

After you set up Live Data Query, you can query data in the Data Cloud Catalog using these sample queries. For more information about the catalog and column types, see [Data Cloud Catalog](#).

### Running the Queries

These sample queries access Workday data. Execute them through the driver connector. All queries in this section use `workday_core.public` as the catalog and schema. Your environment may differ. Always verify your catalog and schema before running data queries:

```
SHOW CATALOGS
SHOW SCHEMAS IN workday_core
SHOW TABLES IN workday_core.public
```

Replace `workday_core.public` in the sample queries with your actual catalog/schema, if different.

Visible tables and columns depend on the security domains granted to your ISU or user-as-self connection.

### Query Patterns

#### Join core tables to reference tables

The most common pattern is joining a core table to a reference table on a foreign key, which turns a stored ID into a readable value. A single foreign key (FK) joins on one column. Without a join, a column like `marital_status` returns the raw ID it stores. When you join it to the `marital_status` reference table, you get a real value such as `Married`.

#### Work with array columns

When you query the array column `worker.languages` directly, you get a list of numeric IDs rather than readable values. SQL has two ways to make the array usable:

- `cardinality(arr)` returns how many elements the array holds.
- `UNNEST(arr)` expands the array into one row per element, so you can join each element to its reference table.

#### Aggregate and filter

Standard SQL filters and aggregates apply to catalog columns. Use catalog column names and data types—for example, filter active workers with `currently_active = true`, or compute tenure from `hire_date` using `date_diff` and `GROUP BY`.

### Sample Queries

#### Count all workers

```
SELECT COUNT(*) AS total_workers
FROM workday_core.public.worker
```

**List active workers with their job titles**

```

SELECT
  w.id,
  w.full_name,
  w.employee_type,
  jp.job_title
FROM workday_core.public.worker w
JOIN workday_core.public.job_profile jp
ON w.job_profile_id = jp.job_profile_id
WHERE w.active = TRUE
LIMIT 50

```

**Workers by management level**

```

SELECT
  management_level,
  COUNT(*)
AS headcount
FROM workday_core.public.worker
WHERE active = TRUE
GROUP BY management_level
ORDER BY headcount DESC

```

**Workers hired in the last 90 days**

```

SELECT
  id,
  full_name, hire_date,
  business_title
FROM
  workday_core.public.worker
WHERE
  hire_date >= CURRENT_DATE - INTERVAL '90' DAY
ORDER BY
  hire_date DESC;

```

**Join worker to marital\_status**

```

SELECT
  w.display_id AS worker_name,
  m.display_id AS worker_marital_status
FROM workday_core.public.worker AS w
JOIN
  workday_core.public.marital_status m ON w.marital_status = m.id
LIMIT 10;

```

**Count languages per worker**

```

SELECT display_id AS worker, cardinality(languages) AS language_count
FROM workday_core.public.worker
WHERE cardinality(languages) > 0
LIMIT 10;

```

**Expand worker languages with UNNEST**

```

SELECT
  w.display_id AS worker,
  lang.display_id AS language
FROM workday_core.public.worker AS w
CROSS JOIN UNNEST(w.languages) AS t (language_id)

```

```
JOIN workday_core.public.language AS lang ON lang.id = t.language_id
LIMIT 10;
```

### Worker tenure by years of service

```
SELECT
    date_diff('year', hire_date, current_date) AS tenure_years,
    count(*) AS worker_count
FROM workday_core.public.worker
WHERE currently_active = true
GROUP BY 1
ORDER BY tenure_years DESC;
```

## Troubleshooting: Workday Live Data Query

### HTTP ERROR 404 Not Found

The client successfully connected to the server, but the specific address/endpoint it asked for doesn't exist on that server.

Cause: The port is incorrect.

Solution: Modify the connection port to 443.

### HTTP 400 Bad Request

The Workday Authorization Server received your request but refused to process it because the information sent was invalid.

Cause: The request payload is syntactically invalid or contains incorrect parameters.

Solution: Confirm the ISU name matches the **View Integration System User** task.

### HTTP 401 Unauthorized

The Workday Authorization Server received the request but refused to validate the identity of the application due to invalid or unrecognized credentials.

Cause: A failure occurred in the authentication process due to invalid or unrecognized credentials.

Solution: Ensure that the Client ID specified in your configuration aligns with the one generated during the **Register API Client** task.

### Access token request failed due to network error

Application was unable to establish a stable connection to the remote server to request an OAuth2 access token.

Cause: This signifies that the application was unable to establish a TCP/IP connection with the specified authorization endpoint, preventing the authentication handshake from initiating.

Solution: Ensure that the Token Endpoint URL is syntactically correct and aligns with the specific Workday environment you are attempting to access (example: PROD, SBOX, IMPL).

### Failed to obtain access token - 404

Indicates successful connection to the server, but the specific OAuth2 Token Endpoint requested does not exist at that location.

Cause: An error occurred during the token exchange.

Solution: Ensure the Token Endpoint URL is correct.

### **Failed to obtain access token - 401**

Error during the token exchange signifies a Cryptographic Authentication Failure.

Cause: Mismatch between the local private key used for signing and the public key registered in the Workday API Client.

Solution: Ensure that the local private key utilized is correct.

### **Objects or fields are missing from the LDQ connection**

Cause: Either:

- The objects aren't available in Data Cloud. See the [Data Cloud Catalog](#) for a list of objects that are available.
- The ISU or user-as-self account you're using to authenticate the LDQ connection doesn't have access to the objects.

Solution: If the object is available in Data Cloud, add the domain for the object to one of the security groups that the ISU or user-as-self account has access to. See [Edit Domain Security Policies](#).

Note: If you make security group changes after the client connection has been established, subsequent permission changes might take up to 15 minutes to propagate before users can access newly authorized tables or columns.

## **Workday Data Lake**

### **Get Started with Workday Data Lake**

#### **Overview**

Note: You might need to take additional steps to enable this feature depending on your organization's subscription service agreement. For more information, see this [Community](#) article.

Workday Data Lake (or "Data Lake") stores Workday data as Iceberg tables. External platforms connect to Data Lake's Iceberg tables through zero-copy connectors (referred to as Workday Data Connect in partner documentation), reading data where it lives rather than exporting or replicating it in the partner platform. This ensures that your data:

- Remains within Workday's security boundary until requested.
- Has a single source of truth, reducing duplication and improving consistency across different platforms and applications.

Data Lake is best suited for large-scale enterprise analytics and data warehousing use cases where periodic data refresh is acceptable, serving as an alternative to traditional ETL pipelines.

#### **Data Lake Setup in Workday**

Before you can connect to Data Lake from an external system, you need to complete these steps:

1. [Set Up the Data Cloud Functional Area](#)
2. [Set Up Integration System Users for Data Lake](#)
3. Register API Clients. Depending on the connector options in the data platform, either:
  - [Register API Client for Data Lake \(JWT Bearer Grant\)](#)
  - [Register API Client for Data Lake \(Refresh Token Grant\)](#)

4. [Set Up Security for Data Lake Admin Console](#)
5. [Add Tables and Columns to Data Lake](#)
6. [Configure Data Lake Access Roles](#)

### Connecting to Data Lake

After you have completed the setup in Workday, you can connect to Data Lake from these external systems:

- Snowflake: Contact your Snowflake Account Executive for a connectivity guide from Snowflake to Workday Data Lake
- Google Cloud: Contact your Google Cloud Account Executive for a connectivity guide from Google Cloud to Workday Data Lake. Alternatively, you can email them at [biglake-help@google.com](mailto:biglake-help@google.com).
- Salesforce: Follow the instructions for the Workday Data Connect connector available in [Data 360 Connectors and Integrations](#).

You may need the following information to configure Data Lake in your data platform:

| Information                                                   | Details                                                                                                                    |
|---------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Token Endpoint                                                | <code>https://{HOST}/ccx/oauth2//token</code>                                                                              |
| Polaris Catalog URL                                           | <code>https://impl-catalog..myworkday.com</code>                                                                           |
| Polaris Catalog/Warehouse Name                                | <code>workday</code>                                                                                                       |
| Principal Role                                                | <code>all-reads-role</code>                                                                                                |
| Realm                                                         | <code>{TENANT}</code>                                                                                                      |
| Private key file associated with API Client (JWT Bearer Flow) | <code>{private-key}.pem</code><br>RSA 2048 PEM format associated with the API Client you configured in your Workday tenant |

### Using Workday Data Lake

After you've established a data connection from the external client, you can query Workday data. See [Query Workday Data Lake](#).

### Limitations

Only certain Workday business objects are available in Data Cloud. See [Data Cloud Catalog](#).

## Set Up the Data Cloud Functional Area

### Prerequisites

Security: These domains in the System functional area:

- *Security Administration*
- *Security Configuration*

### Context

To access Workday Data Lake, you need to set up the Data Cloud functional area. This custom functional area establishes the authorization perimeter required to register secure API Clients for Data Lake.

The Data Cloud functional area may already be configured in your tenant. To verify, follow the steps below. If the functional area doesn't exist, follow the steps to create it.

Note: Custom Functional Area framework setup must be initiated, tested, and audited within a Sandbox or Implementation (Non-Production) tenant environment. Upon verified validation, configurations can be promoted to the Production (PROD) environment.

## Steps

Check if the Data Cloud functional area is already configured in your tenant:

1. Access the **Maintain Functional Areas** task.
2. In the Functional Area column, find Data Cloud and ensure the Enabled check box is selected. If the Data Cloud functional area doesn't display in the grid, skip to the Create the Data Cloud functional area steps below.
3. Access the **Activate Pending Security Policy Changes** task.
  - a) Add an appropriate comment and click **OK**.
  - b) Review the changes, check the **Confirm** check box, and click **OK**.

Create the Data Cloud functional area:

4. Access the **Create Functional Area** task.
5. From the **Domain Group without Functional Area** prompt, select *Data Cloud*.
6. Select the **Confirm** check box.
7. Click **OK**.
8. Activate pending security policy changes. See step 3 above.

## Next Steps

[Set Up Integration System Users for Data Lake](#) .

# Set Up Integration System Users for Data Lake

## Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

## Context

External data platforms connect to Workday Data Lake as integration system users (ISUs). Each ISU represents a non-person system account that a partner platform uses to authenticate to Data Lake.

We recommend creating one or more Workday ISU per external platform and access pattern.

## Steps

1. Create integration system security groups. If you already have appropriate security groups, skip this step.  
Security: *Security Configuration* domain in the System functional area.
  - a) Access the **Create Security Group** task.
  - b) From the **Type of Tenanted Security Group** prompt, select *Integration System Security Group (Unconstrained)*.
  - c) (Optional) If you already have ISUs for Data Lake, select them from the **Integration System Users** prompt.

## 2. Create user provisioning groups for Data Lake.

Security: *Data Cloud: Manage Data Lake Access Policies* and *Manage: Datalake User Provisioning* domains in the Data Cloud functional area.

- Access the **Manage User Provisioning for Workday Products** worklet. If you don't have access to this task, see [Set Up Access to User Provisioning](#).
- Select **Set Up Security Groups for User Provisioning**.
- In the grid, add the security groups you created in step 1.
- Click **OK**.
- On the **Manage User Provisioning for Workday Products** page, select **User Provisioning Workspace**.
- On the **Data Cloud: Data Lake** tile, select **Configure**.
- On the **Configuration** tab that displays, select your security groups from the **Workday Security Groups** prompt and click **Create Provisioning Group**.
- Select **Preview and Enable Sync**. For more information on provisioning group previews, see [Manage Reports for User Provisioning](#).
- Once you generate a preview report and review any errors, select both the check box and the **Enable Sync** button to synchronize the provisioning group across Workday and Data Lake. The Sync label will change from Off to On once synchronization is complete.

## 3. Create ISUs for Data Lake.

We recommend one or more ISU per platform and authorization policy defined in Admin Console.

Security: *Integration Security* domain in the Integration functional area.

- Access the **Create Integration System User** task.
- Complete the task:

| Option                                            | Description                                                                          |
|---------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>User Name</b>                                  | Enter a unique ISU user name.                                                        |
| <b>New Password</b><br><b>New Password Verify</b> | Set a password.                                                                      |
| <b>Require New Password at Next Sign In</b>       | Don't select the check box.                                                          |
| <b>Session Timeout Minutes</b>                    | Keep the value at zero to prevent the integration system user session from expiring. |
| <b>Do Not Allow UI Sessions</b>                   | Select the check box.                                                                |

- Repeat these steps for each ISU required for your use case.

## 4. Add Data Lake ISUs to security groups.

Security: *Security Configuration* domain in the System functional area.

- Access the **Edit Security Group** task, and select a security group you created in step 1.
- Click **OK**.
- Select your ISUs from the **Integration System Users** prompt.
- Click **OK**.

## Result

The integration system users in your security groups are synchronized to Data Lake and can be used to authenticate from external clients.

Access to Data Lake tables is still controlled by access roles in the Admin Console. Configure those policies separately for each ISU or use case.



## Next Steps

Register API Clients. Depending on the connector options in the data platform, either:

- Establish connection using a [JWT Bearer Grant flow](#).
- Establish connection using a [Refresh Token Grant flow](#).

## Register API Client for Data Lake (JWT Bearer Grant)

### Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

Security:

- *Integration Security* domain in the Integration functional area.
- *Set Up: Tenant Setup - Security* and *Security Administration* domains in the System functional area.

### Context

Data Platform Clients authenticate their connection to Workday Data Lake through API clients. Depending on the connector options in the data platform, you can either establish your connection with a JWT Bearer Grant flow (this topic) or a [Refresh Token Flow](#).

### Steps

1. Create an Asymmetric Key Pair using your organization's standard tool.  
The public key is used in registering the API client within Workday. The private key is used to set up a connection to Data Lake from partner platforms.
2. Access the **Register API Client** task.

Note: The private key from the asymmetric key-pair is used in the partner platform to access the Data Lake.

- a) Complete the task:

| Option                  | Description                                                                                                                                                                            |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client Name             | Enter a unique API client name.                                                                                                                                                        |
| Client Grant Type       | Select <b>JWT Bearer Grant</b> .                                                                                                                                                       |
| x509 Certificate        | Select <b>Create x509 Public Key</b> .<br>Enter all details of the created certificate.<br>Example: <pre>-----BEGIN CERTIFICATE----- MIID.... ....Nl4= -----END CERTIFICATE-----</pre> |
| Integration System User | Add the ISU user name.                                                                                                                                                                 |
| Scope                   | Select <b>Data Cloud</b> . If this option doesn't display, see <a href="#">Set Up the Data Cloud Functional Area</a> .                                                                 |

- b) Click **OK**.

3. Save the API Client information from the **Register API Client** page.

You'll need this information to configure Workday connector properties in the partner platform:

| Information             | Example or Details                                                                           |
|-------------------------|----------------------------------------------------------------------------------------------|
| Token Endpoint          | https://<HOST>/ccx/oauth2/<TENANT>/token                                                     |
| Integration System User | demo-isu                                                                                     |
| Client ID               | BASE64_CLIENT_ID<br>Sample Value: OTc3MmE4ZmEt . . .<br>[TRUNCATED_FOR_EXAMPLE] . . .bUxZjIz |

### Next Steps

[Set Up Security for Data Lake Admin Console](#) .

## Register API Client for Data Lake (Refresh Token Grant)

### Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

Security:

- *Integration Security* domain in the Integration functional area.
- *Set Up: Tenant Setup - Security* and *Security Administration* domains in the System functional area.

### Context

Data Platform Clients authenticate their connection to Workday Data Lake through API clients. Depending on the connector options in the data platform, you can either establish your connection with a Refresh Token Grant flow (this topic) or a [JWT Bearer Grant flow](#).

### Steps

Access the **Register API Client for Integrations** task.

- a) Complete the task:

| Option                      | Description                                                                                                            |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------|
| Client Name                 | Enter a unique API client name.                                                                                        |
| Non-Expiring Refresh Tokens | Select this check box.                                                                                                 |
| Scope (Functional Areas)    | Select <b>Data Cloud</b> . If this option doesn't display, see <a href="#">Set Up the Data Cloud Functional Area</a> . |

- b) Click **OK**.
- c) Save the Client ID and Client Secret argument values from the **Register API Client for Integrations** page.

| Information | Example or Details                                                                           |
|-------------|----------------------------------------------------------------------------------------------|
| Client ID   | BASE64_CLIENT_ID<br>Sample Value: OTc3MmE4ZmEt . . .<br>[TRUNCATED_FOR_EXAMPLE] . . .bUxZjIz |

| Information   | Example or Details |
|---------------|--------------------|
| Client Secret | yego9m0c . . . . . |

- d) On the same page, from the related actions menu of the client, select **API Client > Manage Refresh Tokens for Integrations**.
- e) Select the ISU name from the **Workday Account** field, then click **OK**.
- f) Save the Refresh Token value from the **Successfully Regenerated Refresh Token** page.

| Information   | Example                                                       |
|---------------|---------------------------------------------------------------|
| Refresh Token | nmdjc98r2j . . .<br>[TRUNCATED_FOR_EXAMPLE] . . . 4e6fqz9swlc |

### Next Steps

[Set Up Security for Data Lake Admin Console](#) .

## Set Up Security for Data Lake Admin Console

### Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

Security: *Security Administration* and *Security Configuration* domains in the System functional area.

### Context

Note: During the Implementation-only phase, access to the **Data Lake Admin Console** is not segmented by security domain. Users who have access to either domain can use all console capabilities. Domain separation will be enforced in a future release.

Access to the **Data Lake Admin Console** report is controlled by two distinct Workday-delivered security domains in the Data Cloud functional area:

- *Data Cloud: Manage Data Lake Application Configurations*: Access the console only to configure which Workday objects and fields are made available in Data Lake.
- *Data Cloud: Manage Data Lake Access Policies*: Access the console only to manage Data Lake ISUs (Apache Polaris principals) and define configurable security on Data Lake objects.

Add users to security groups associated with these domains depending on their responsibilities.

### Steps

1. Create security groups for each domain and assign users to them (Example: SYS\_Data\_Lake\_Admins). If you already have appropriate security groups, skip this step.
  - a) Access the **Create Security Group** task.
  - b) From the **Type of Tenanted Security Group** prompt, select **User-Based Security Group**.
  - c) Access the **Assign Users to User-Based Security Group** task to add users to your security groups.

2. Update the domain security policies.
  - a) Access the **View Domain** task and select the *Data Cloud: Manage Data Lake Application Configurations* domain from the prompt.
  - b) From the related actions menu of the domain, select **Domain > Create Security Policy**.
  - c) Select the **Confirm** check box and click **OK**.
  - d) In the **Report/Task Permissions** grid, add rows and select your security groups from the prompt.
  - e) Select the **View and Modify** check boxes for each security group, according to the level of access required.
  - f) Click **OK**, and then **Done**.
  - g) Repeat these steps for the *Data Cloud: Manage Data Lake Access Policies* domain.
3. Access the **Activate Pending Security Policy Changes** task.
  - a) Add an appropriate comment and click **OK**.
  - b) Review the changes, check the **Confirm** check box, and click **OK**.

### Next Steps

[Add Tables and Columns to Data Lake](#) .

## Add Tables and Columns to Data Lake

### Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

Security: *Data Cloud: Manage Data Lake Application Configurations* domain in the Data Cloud functional area.

### Context

Workday Data Lake uses an opt-in model: nothing from the Data Cloud Catalog is materialized until you configure it. Discovery (what's available in the catalog) is separate from materialization (what is stored in the lake and refreshed on a schedule). See [Data Cloud Catalog](#) for information about available tables and columns, core vs reference tables, and table associations via foreign keys.

Configuration happens in three stages:

- **Table selection:** Choose which core tables to register for Data Lake. This syncs table metadata into the lake environment; underlying records are not materialized until you select columns.
- **Column selection:** Choose which columns to include from each table. Only selected columns are materialized for zero-copy access. ID columns are included by default; other columns are excluded until you opt in. When you include multi-instance (array) columns, you can opt in related reference tables from the **Opt-In Reference Tables** grid.
- **Frequency configuration:** Set a refresh frequency for how often materialized data is updated.

New tables and columns added to the Data Cloud Catalog in future releases are not automatically added to your data lake. You must explicitly include them if you want them materialized.

## Steps

1. Add tables to your data lake.

You can only select core tables at this step. Reference tables and core tables with foreign key associations can be added in step 2. See [Data Cloud Catalog](#) for information on core tables, reference tables, and foreign keys.

- a) Access the **Data Lake Admin Console** report.
  - b) On the **Applications** tab, select **Edit** from the actions menu of the **Workday Core** application.
  - c) On the **Edit Core Workday Catalog for Data Lake** task that displays, select core tables that you want to add to your data lake from the **Tables for Data Lake Materialization** prompt.
  - d) Click **Add**.
  - e) Click **OK**.
2. Select the columns to include from the tables selected in step 1.  
Aside from ID columns, all columns are excluded by default when you add a table.
    - a) Select **Manage Fields** for the table you added.
    - b) On the **Manage Fields for Data Lake Materialization** page that displays, select the **Include** check box for your desired fields.
    - c) Click **OK**.
    - d) When you include multi-instance fields, an **Opt-In Reference Tables** grid appears. Select the **Include** check box to add the desired tables/columns to your data lake.
  3. After adding your desired tables and columns, set a refresh frequency for your Data Lake data.
    - a) On the **View Core Workday Catalog for Data Lake** page, click **Edit**.
    - b) Select the **Schedule** tab.
    - c) Select a **Run Frequency**.

The available options depend on your Data Cloud subscription and the type of tenant you're working in:

| Data Cloud Subscription | Tenant Type                                  | Highest Supported Frequency |
|-------------------------|----------------------------------------------|-----------------------------|
| PRO                     | Production                                   | Hourly refresh.             |
| CORE                    | Production                                   | Daily refresh.              |
| PRO<br>CORE             | Sandbox<br>Sandbox Preview<br>Implementation | Weekly refresh.             |

- d) Configure the recurrence criteria.

## Result

You can view the results of your configuration in the **Applications** tab of the **Data Lake Admin Console**.

## Next Steps

[Configure Data Lake Access Roles](#) .

## Configure Data Lake Access Roles

### Prerequisites

See the Data Lake Setup in Workday section of [Get Started with Workday Data Lake](#) . This topic is part of a larger procedure.

The data refresh, according to the schedule on your Admin Console configuration, has run at least once.

Security: *Data Cloud: Manage Data Lake Access Policies* domain in the Data Cloud functional area.

## Context

After you add tables and columns to your data lake, you can configure data lake access roles to determine which users can read them. You create a data lake access role, assign one or more integration system users, and select the tables those users may read.

Workday Data Lake enforces these rules at the catalog layer. Access is enforced at the table level. Column selection is configured when you add tables and columns to the lake. Partner platforms may apply additional filters such as row or column access, but cannot grant access beyond what you assign here.

## Steps

1. Access the **Data Lake Admin Console** report.
2. On the **Access Management** tab, select **Add Data Lake Access Role**.
3. Enter an appropriate name and description for the role.
4. From the **Principal** prompt, select an ISU you set up in [Set Up Integration System Users for Data Lake](#).
5. From the **Object Access** prompt, select tables the principal needs read access to. If the prompt is empty, review the Prerequisites section of this topic.
6. Click **Save**.

## Result

After you click **Save**, the Data Lake access role is active in Polaris:

- The principals you selected can authenticate to Data Lake (using credentials configured in Workday) only within the table scope you assigned on the role.
- Principals can't read Iceberg tables that you didn't select from the **Object Access** prompt for that role.

## Next Steps

[Query Workday Data Lake](#)

## Query Workday Data Lake

After you set up Data Lake, you can query data in the Data Cloud Catalog using these sample queries. For more information about the catalog and column types, see [Data Cloud Catalog](#).

### Running the Queries

These sample queries access Workday data. Execute them through the external client. All queries in this section use `workday.workday_core` as the catalog and namespace. Your environment and setup may differ. Always verify your catalog and schema before running data queries:

```
SHOW CATALOGS
SHOW SCHEMAS IN {Catalog/Integration Name}
SHOW TABLES IN workday.workday_core
```

Replace `workday.workday_core` in the sample queries with your actual catalog/schema, if different.

## Query Patterns

### Join core tables to reference tables

The most common pattern is joining a core table to a reference table on a foreign key, which turns a stored ID into a readable value. A single foreign key (FK) joins on one column. Without a join, a column like

`marital_status` returns the raw ID it stores. When you join it to the `marital_status` reference table, you get a real value such as *Married*.

### Joiner tables

Array foreign keys (FK[]) can't be expressed in every external system. Workday Data Lake provides joiner tables under the `join_workday_core` schema to account for this. Each one flattens an array relationship into a two-column bridge with one row per (source, target) pair, so you can rebuild the relationship without `UNNEST`.

Joiner tables follow a fixed naming pattern:

```
join_<source_table>_to_<target_table>_via_<source_fk_column>
```

The Data Cloud Catalog contains a list of all the joiner tables.

### Work with array columns

When you query the array column `worker.languages` directly, you get a list of numeric IDs rather than readable values. SQL has two ways to make the array usable:

- `cardinality(arr)` returns how many elements the array holds.
- `UNNEST(arr)` expands the array into one row per element, so you can join each element to its reference table.

### Aggregate and filter

Standard SQL filters and aggregates apply to catalog columns. Use catalog column names and data types—for example, filter active workers with `currently_active = true`, or compute tenure from `hire_date` using `date_diff` and `GROUP BY`.

### Sample Queries

#### Join worker to marital\_status

```
SELECT
  w.display_id AS worker_name,
  m.display_id AS worker_marital_status
FROM workday.workday_core.worker AS w
JOIN
  workday.workday_core.marital_status m ON w.marital_status = m.id
LIMIT 10;
```

#### Count languages per worker

```
SELECT display_id AS worker, cardinality(languages) AS language_count
FROM workday.workday_core.worker
WHERE cardinality(languages) > 0
LIMIT 10;
```

#### Expand workers languages with UNNEST

```
SELECT
  w.display_id AS worker,
  lang.display_id AS language
FROM workday.workday_core.worker AS w
CROSS JOIN UNNEST(w.languages) AS t (language_id)
JOIN workday.workday_core.language AS lang ON lang.id = t.language_id
LIMIT 10;
```

#### Resolve an array relationship with a joiner table

Where UNNEST isn't available, the joiner table gets you the same result as ordinary joins.

```
SELECT
  sp.display_id AS succession_plan,
  spc.display_id AS candidate,
  spc.top_candidate
FROM workday.workday_core.succession_plan AS sp
JOIN
  workday.join_workday_core.join_succession_plan_to_succession_plan_candidate_via_candidate AS j ON j.source_id = sp.id
JOIN workday.workday_core.succession_plan_candidate AS spc ON j.target_id =
  spc.id;
```

The joiner exposes the source and target primary keys as columns. Confirm exact column names against the catalog's Joiner Tables sheet.

### Worker tenure by years of service

```
SELECT
  date_diff('year', hire_date, current_date) AS tenure_years,
  count(*) AS worker_count
FROM workday.workday_core.worker
WHERE currently_active = true
GROUP BY 1
ORDER BY tenure_years DESC;
```

## Data In

### Get Started with Zero-Copy Data In

#### Overview

Note: You might need to take additional steps to enable this feature depending on your organization's subscription service agreement. For more information, see this [Community](#) article.

Zero-Copy Data In leverages Workday Prism data management tools to bring zero-copy external data into Workday. Unlike traditional Prism ingestion, which copies data into Workday, Zero-Copy Data In enables customers to query external data where it resides, without physical movement and copying of data.

#### Example Use Case

Total Cost of Labor Analysis with Zero-Copy Data Blending: Leadership requires a detailed analysis of the total cost of labor across three dimensions: Base Pay, Benefits (Workday data), and Stock Compensation (external data stored in Snowflake).

The traditional method of extracting sensitive Workday data for analysis in third-party applications presents a significant security risk.

Zero-Copy Data In solves this by enabling the secure blending of Workday data with external data (Stocks from Snowflake) to create a comprehensive visualization, eliminating the need to move sensitive labor data out of Workday.

#### Data In Setup

As a prerequisite, you must [Set Up Workday Security Policies for Zero-Copy Data In](#).



| External System | Setup Instructions                                         |
|-----------------|------------------------------------------------------------|
| Databricks      | <a href="#">Set Up Zero-Copy Data In with Databricks</a>   |
| Google Cloud    | <a href="#">Set Up Zero-Copy Data In with Google Cloud</a> |
| Salesforce      | <a href="#">Set Up Zero-Copy Data In with Salesforce</a>   |
| Snowflake       | <a href="#">Set Up Zero-Copy Data In with Snowflake</a>    |

## Set Up Workday Security Policies for Zero-Copy Data In

### Prerequisites

#### Security:

- *Set Up: Assignable Roles* domain in the Organizations and Roles functional area.
- These security domains in the Prism Analytics functional area:
  - *External Catalog Create*: Controls who can create a Catalog.
  - *External Catalog Manage*: Controls who can edit a catalog, including access to view existing catalogs, edit catalog metadata, and delete catalogs.
  - *External Catalog Owner Manage*: Controls who can perform specific catalog operations, such as sharing the table with other users and syncing the catalog.
  - *External Tables Manage*: Controls who can edit a table, including viewing the table and metadata, editing table metadata, and enabling or disabling for analysis.
  - *External Tables Owner Manage*: Controls who can perform specific table operations, such as sharing the table with other users.
  - *Prism: Manage Connection*: Controls who can create a connection.

### Context

Domain security policies secure access to configurable items in Workday. Configure the assignable roles for the Data In security domains, and set up the security policy for each of the prerequisite security domains.

### Steps

1. Access the **Maintain Assignable Roles** task.
2. In the **Maintain Assignable Roles** grid, add rows for each of these role names, with their corresponding role assignment values:

| Role Name             | Workday Role                      | Enabled for          | Self-Assign             | Is Leader / Is Supporting | Role Assignees Restricted To                                                                           | Assigned/ Reviewed By Security Groups                                                                                                               |
|-----------------------|-----------------------------------|----------------------|-------------------------|---------------------------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| External Table Editor | 02. External Table Editor - Prism | Prism External Table | Leave the box unchecked | None of the Above         | <ul style="list-style-type: none"> <li>• Accountant (Unconstrained)</li> <li>• Data Analyst</li> </ul> | <ul style="list-style-type: none"> <li>• Data Administrator</li> <li>• Prism External Table Owner (Workday Owned)</li> <li>• Prism Table</li> </ul> |

| Role Name            | Workday Role                      | Enabled for          | Self-Assign             | Is Leader / Is Supporting | Role Assignees Restricted To                                                                       | Assigned/ Reviewed By Security Groups                                                                                                                               |
|----------------------|-----------------------------------|----------------------|-------------------------|---------------------------|----------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      |                                   |                      |                         |                           |                                                                                                    | Owner (Workday Owned)                                                                                                                                               |
| External Table Owner | 01. External Table Owner - Prism  | Prism External Table | Select the check box    | Is Leader                 | Leave blank                                                                                        | <ul style="list-style-type: none"> <li>Data Administrator</li> <li>Prism External Table Owner (Workday Owned)</li> <li>Prism Table Owner (Workday Owned)</li> </ul> |
| External Table View  | 03. External Table Viewer - Prism | Prism External Table | Leave the box unchecked | None of the Above         | <ul style="list-style-type: none"> <li>Accountant (Unconstrained)</li> <li>Data Analyst</li> </ul> | <ul style="list-style-type: none"> <li>Data Administrator</li> <li>Prism External Table Owner (Workday Owned)</li> <li>Prism Table Owner (Workday Owned)</li> </ul> |

3. Edit the security policy for each required security domain. Repeat the substeps below for each security domain in the table.

- In the Search box, enter domain: <security domain name>. Example: domain: *External Catalog Create*.
- On the View Domain report, select Domain > Edit Security Policy Permissions from the related actions menu of the domain. For each of the security domains in this table, add the security groups and their corresponding permissions:

| Security Domain                | Report/Task Permissions Security Groups        | Report/Task Permissions | Integration Permissions              | Integration Permissions Security Groups |
|--------------------------------|------------------------------------------------|-------------------------|--------------------------------------|-----------------------------------------|
| <i>External Catalog Create</i> | View and Modify for Prism: Tables Owner Manage | View<br>Modify          | Get and Put for Prism: Tables Manage | Get<br>Put                              |
| <i>External Catalog Manage</i> | View Only for Prism: Tables Manage             | View                    | Get and Put for Prism: Tables Manage | Get<br>Put                              |
|                                | View and Modify for Prism: Tables Manage       | View<br>Modify          | Get and Put for Prism: Tables Manage | Get<br>Put                              |

| Security Domain                      | Report/Task Permissions Security Groups        | Report/Task Permissions | Integration Permissions              | Integration Permissions Security Groups |
|--------------------------------------|------------------------------------------------|-------------------------|--------------------------------------|-----------------------------------------|
| <i>External Catalog Owner Manage</i> | View and Modify for Prism: Tables Owner Manage | View<br>Modify          | Get and Put for Prism: Tables Manage | Get<br>Put                              |
| <i>External Tables Manage</i>        | View Only for Prism: Tables Manage             | View                    | Get and Put for Prism: Tables Manage | Get<br>Put                              |
|                                      | View and Modify for Prism: Tables Manage       | View<br>Modify          | Get and Put for Prism: Tables Manage | Get<br>Put                              |
| <i>External Tables Owner Manage</i>  | View and Modify for Prism: Tables Owner Manage | View<br>Modify          | Get and Put for Prism: Tables Manage | Get<br>Put                              |

c) Click **OK**.

The **Edit Domain Security Policy Permissions** page displays the updated configuration.

d) Click **Done**.

4. Access the **Activate Pending Security Policy Changes** task.

a) Describe your changes in the **Comment** field.

b) Click **OK**.

c) Select the **Confirm** check box to activate your changes.

The **View All Security Timestamps** report displays the confirmation.

d) Click **OK**.

## Result

Your Data In security domain policies have been updated and activated.

## Set Up Zero-Copy Data In with Databricks

### Prerequisites

- [Configure Databricks on AWS](#).
  - Take note of your Databricks **Warehouse URL** and **Account ID** from your account details.
- [Set Up Workday Security Policies for Zero-Copy Data In](#).

### Summary

- [Task 1: Create a Zero-Copy Connection to Databricks](#)
- [Task 2: Create an External Catalog](#)
- [Task 3: Explore and Validate External Tables](#)

### Task 1: Create a Zero-Copy Connection to Databricks

You can access Databricks table metadata without copying data by creating a zero-copy connection from Workday to your Databricks Iceberg catalog.

1. In Workday, access the [Prism Data Catalog](#) report.
2. Select **+ Create > Connection**.
3. In the **Create Connection** wizard, under **Zero Copy Connectors**, select the **Databricks** tile.

## 4. Enter the connection details:

| Field                      | Description                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>                | The connection name must be unique in the Data Catalog. Example: <code>Databricks_ZeroCopy_&lt;first name&gt;&lt;last initial&gt;</code> . |
| <b>Description</b>         | (Optional) A brief description of the connection's purpose.                                                                                |
| <b>Catalog URL</b>         | The URL of the Databricks catalog. Format: <code>Databricks_Warehouse_URL/api/2.1/unity-catalog/iceberg-rest</code> .                      |
| <b>Catalog Name</b>        | Databricks Catalog name.                                                                                                                   |
| <b>Authentication Type</b> | Select <i>OAuth 2.0 with Token Exchange</i> .                                                                                              |
| <b>Issuer</b>              | Issuer identifier of the authorization server. Format: <code>https://databricks.com/account/&lt;account_id&gt;/oidc/v1/token</code> .      |
| <b>Principal</b>           | Service Principal from your Databricks system.                                                                                             |
| <b>Token Endpoint</b>      | OAuth2 token endpoint URL from your external system. Format: <code>Databricks_Warehouse_URL/oidc/v1/token</code> .                         |
| <b>Scopes</b>              | OAuth2 scopes from your Databricks system that define the level of access for the Principal.                                               |

5. Save the connection but do not test it yet. Before you test the connection, you must complete the certificate setup in Databricks using the JWK certificate generated by Workday.
6. On the **View Connection Details** page that displays, select **Certificate > Download > JWK**.
7. Update Databricks with the JWK Certificate.
  - a. Log in to your Databricks account console.
  - b. Locate your service principal used for the Zero-Copy Data In connection.
  - c. Update the federation policy of your service principal, adding the JWK certificate content from Workday to the "keys" list in the existing JWKS JSON.
  - d. Save your changes.
8. In Workday, on the **View Connection Details** page, select **Test Connection**.  
If the test fails, verify:
  - Catalog URL format.
  - Token endpoint URL.
  - Service principal identity and account values.
  - JWKS JSON federation policy configuration in Databricks.

**Task 2: Create an External Catalog**

An external catalog acts as a virtual window into your external data, allowing you to browse tables in real-time.

1. Access the **Data Catalog** report.
2. Select **+ Create > External Catalog**.
3. Enter the external catalog details:

| Field                        | Description                                                                                                                                                                            |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>External Catalog Name</b> | Enter a descriptive name for your external catalog. Example: <code>EA_Procurement_Catalog</code> or <code>EA_HR_External_Data</code> .                                                 |
| <b>Connection</b>            | Select the connection you just created.                                                                                                                                                |
| <b>Namespace</b>             | The exact namespace where the Iceberg connector tables reside. This information must match the Databricks system value exactly. Example: <code>enterprise.finance.procurement</code> . |
| <b>Description</b>           | (Optional) The description of the external catalog.                                                                                                                                    |

4. Click **Save**.

5. Wait for the **Success** notification that the External Catalog was created.
6. Select **Back to Data Catalog**.
7. Select **External Catalog**, and validate that schema sync is successful on your External Catalog.

After you save the External Catalog, Workday automatically runs a **Schema Sync** to:

- Read the metadata of Iceberg tables in the specified Databricks namespace.
- Register those tables as **External Tables** in Workday.
- Make table structures visible in the Data Catalog without moving or copying data.

The duration of this sync depends on the number and complexity of tables in the namespace.

### Task 3: Explore and Validate External Tables

Explore the external tables registered from the external partner and perform basic schema validation inside Workday.

1. Locate the Databricks external tables.
  - a. Access the **Data Catalog** report.
  - b. In the left-hand menu, select **External Tables**.
  - c. Click on the **Refresh** button.
  - d. In the External Tables view, confirm that all Iceberg tables from the external partner are listed under the schema/namespace you selected during external catalog creation.
2. Inspect an external table.
  - a. Right-click or double-click the table name to open **View External Table Details**.
  - b. On the details page, review:
    - Column names
    - Data types
    - Any available column statistics
    - Sample data preview (if present)
3. Verify that:
  - Data types (Example: Numeric vs. string) align with the definitions in the external system.
  - Values appear as expected.

### Next Steps

- [Enable External Data for Analysis](#)
- [Blend External Data with Workday Data](#)
- [Create Visualizations](#)

## Set Up Zero-Copy Data In with Google Cloud

### Prerequisites

- Configure Google Cloud BigLake for zero-copy access to an Iceberg REST catalog.
  - Take note of the GCS Warehouse Name and Project ID from your Google Cloud configuration.
- [Set Up Workday Security Policies for Zero-Copy Data In](#) .

### Summary

- [Task 1: Configure Google Cloud Workload Identity Federation](#)
- [Task 2: Create a Zero-Copy Connection to Google Cloud](#)

- [Task 3: Create an External Catalog](#)
- [Task 4: Explore and Validate External Tables](#)

### Task 1: Configure Google Cloud Workload Identity Federation

Before creating the zero-copy connection in Workday, you need to configure Google Cloud Workload Identity Federation so that the external identity can impersonate the Google Cloud service account used to access BigLake.

1. In the Google Cloud console, create a Workload Identity Pool. Use a descriptive name, such as Prism BigLake pool.
2. Add a provider to the pool.
3. Select OIDC as the provider type.  
Don't select SAML or AWS.
4. Enter the issuer URL for the identity provider.
5. Configure the audience for the provider.
6. Create or identify the Google Cloud service account that will access the BigLake catalog.
7. Take note of the service account email address. You'll enter it as the **Principal** when you create the connection in Task 2.
8. Grant the service account the required read access to the catalog and underlying data.
9. Link the service account to the Workload Identity Pool provider by selecting **Grant Access** and configuring the required principal and attribute conditions.
10. Download the client library config file associated with the service account with access to the Workload Identity Pool to obtain the connection values.
  - a. Select the Workload Identity Pool and go to the **Connection Service Accounts** tab.
  - b. Download the **Client library config** associated with the service account.
  - c. Select the identity provider.
  - d. Enter an OIDC ID token path. Example: /tmp/oidc-token.
  - e. Select JSON as the format type.
  - f. Download the configuration file.
11. Record the audience and token\_url values exactly as they appear in the downloaded configuration file. You'll need these values when you create the connection in Workday.

### Task 2: Create a Zero-Copy Connection to Google Cloud

You can access Google Cloud table metadata without copying data by creating a zero-copy connection from Workday to your Google Cloud Iceberg catalog.

1. In Workday, access the [Prism Data Catalog](#) report.
2. Select **+ Create > Connection**.
3. In the **Create Connection** wizard, under **Zero Copy Connectors**, select the **Google Cloud** tile.
4. Enter the connection details:

| Field              | Description                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>        | The connection name must be unique in the Data Catalog.<br>Example: GoogleCloud_ZeroCopy_<first name><last initial>. |
| <b>Description</b> | (Optional) A brief description of the connection's purpose.                                                          |

| Field                      | Description                                                                                                                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Catalog URL</b>         | The Google Cloud BigLake Iceberg REST catalog URL.<br><br>In the Lakehouse Runtime Catalog, select the catalog that you created and copy the REST catalog URL property value exactly from the Catalog Detail page.<br><br>Example: <code>https://biglake.googleapis.com/iceberg/v1/restcatalog</code> |
| <b>Catalog Name</b>        | The GCS Warehouse Name configured for the BigLake catalog.<br><br>In the Lakehouse Runtime Catalog, select the catalog that you created and copy the GCS warehouse property value exactly from the Catalog Detail page.<br><br>Example: <code>gs://some-warehouse</code>                              |
| <b>Use Sandbox</b>         | Select this option only when connecting to the applicable sandbox environment.                                                                                                                                                                                                                        |
| <b>Project ID</b>          | The Google Cloud project ID that hosts the catalog.<br><br>In the Lakehouse Runtime Catalog, select the catalog that you created and copy the project property value exactly from the Catalog Detail page.                                                                                            |
| <b>Authentication Type</b> | Select <i>OAuth 2.0 with Token Exchange</i> .                                                                                                                                                                                                                                                         |
| <b>Issuer</b>              | The issuer URL configured on the OIDC Workload Identity Pool provider. The value must match the provider configuration exactly.                                                                                                                                                                       |
| <b>Principal</b>           | The Google Cloud service account principal used by the connection.<br><br>Example: <code>biglake-client@gcp-eng.iam.gserviceaccount.com</code>                                                                                                                                                        |
| <b>Token Endpoint</b>      | The <code>token_url</code> value from the downloaded client library config file downloaded for the above service account in task 1.                                                                                                                                                                   |
| <b>Scopes</b>              | Enter <code>https://www.googleapis.com/auth/cloud-platform</code> .<br><br>This is the standard cloud-platform scope used for service account impersonation.                                                                                                                                          |
| <b>Audience</b>            | The audience value from the downloaded client library config file downloaded for the above service account in task 1.                                                                                                                                                                                 |

5. Save the connection but do not test it yet. Before you test the connection, complete the following steps.
6. On the **View Connection Details** page that displays, select **Certificate > Download > JWK**.
7. Update the Google Cloud workload identity provider with the JWK Certificate.
  - a. In the Google Cloud console, edit the Workload Identity Pool provider you created in task 1.  
**View JWK File** shows all the existing JWK keys associated with the identity provider.
  - b. Upload a new file containing all the existing keys you want to preserve and the new JWK key downloaded from the configured connection in Workday.
  - c. Save your changes.

In Workday, you can now test your connection and validate that it is successful.

8. In Workday, on the **View Connection Details** page, select **Test Connection**.

If the test fails, verify:

- Catalog URL format.
- Token endpoint URL.
- Google Cloud project ID and service account principal values.
- Workload Identity Pool provider audience value.
- JWK certificate and Workload Identity configuration in Google Cloud.
- Service account permissions for the catalog and underlying data.

### Task 3: Create an External Catalog

An external catalog acts as a virtual window into your external data, allowing you to browse tables in real-time.

1. Access the **Data Catalog** report.
2. Select **+ Create > External Catalog**.
3. Enter the external catalog details:

| Field                        | Description                                                                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>External Catalog Name</b> | Enter a name for your external catalog.                                                                                                                              |
| <b>Connection</b>            | Select the connection you just created.                                                                                                                              |
| <b>Namespace</b>             | The exact namespace where the Iceberg connector tables reside. This information is provided by Google Cloud.<br><br>Example: <i>enterprise.finance.procurement</i> . |
| <b>Description</b>           | (Optional) The description of the external catalog.                                                                                                                  |

4. Click **Save**.
5. Wait for the **Success** notification that the External Catalog was created.
6. Select **Back to Data Catalog**.
7. Select **External Catalog**, and validate that schema sync is successful on your External Catalog.

After you save the External Catalog, Workday automatically runs a **Schema Sync** to:

- Read the metadata of Iceberg tables in the specified Google Cloud namespace.
- Register those tables as **External Tables** in Workday.
- Make table structures visible in the Data Catalog without moving or copying data.

The duration of this sync depends on the number and complexity of tables in the namespace.

### Task 4: Explore and Validate External Tables

Explore the external tables registered from Google Cloud and perform basic schema validation inside Workday.

1. Locate the Google Cloud external tables.
  - a. Access the **Data Catalog** report.
  - b. In the left-hand menu, select **External Tables**.
  - c. Click on the **Refresh** button.
  - d. In the **External Tables** view, confirm that all Iceberg tables from Google Cloud are listed under the schema/namespace you selected during external catalog creation.



2. Inspect an external table.
  - a. Right-click or double-click the table name to open **View External Table Details**.
  - b. On the details page, review:
    - Column names
    - Data types
    - Any available column statistics
    - Sample data preview (if present)
3. Verify that:
  - Data types (Example: Numeric vs. string) align with the definitions in the external system.
  - Values appear as expected.

### Next Steps

- [Enable External Data for Analysis](#)
- [Blend External Data with Workday Data](#)
- [Create Visualizations](#)

## Set Up Zero-Copy Data In with Salesforce

### Prerequisites

- [Configure Salesforce](#).
  - Create a Data Stream, Data Share, and Data Share Target.
  - Take note of the **Issuer URL**, **Service Account ID**, **Tenant Endpoint**, and **Core Tenant ID** of your Data Share Target.
- [Set Up Workday Security Policies for Zero-Copy Data In](#)

### Summary

- [Task 1: Create a Zero-Copy Connection to Salesforce](#)
- [Task 2: Create an External Catalog](#)
- [Task 3: Explore and Validate External Tables](#)

### Task 1: Create a Zero-Copy Connection to Salesforce

You can access Salesforce table metadata without copying data by creating a zero-copy connection from Workday to your Salesforce Iceberg catalog.

1. In Workday, access the [Prism Data Catalog](#) report.
2. Select **+ Create > Connection**.
3. In the **Create Connection** wizard, under **Zero Copy Connectors**, select the **Salesforce** tile.
4. Enter the connection details:

| Field              | Description                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>        | The connection name must be unique in the Data Catalog.<br>Example: Salesforce_ZeroCopy_<first name><last initial>.              |
| <b>Description</b> | (Optional) A brief description of the connection's purpose.                                                                      |
| <b>Catalog URL</b> | Add https:// and enter the <b>Tenant Endpoint</b> from your Data Share Target in Salesforce. Example: https://xyz.salesforce.com |

| Field                 | Description                                                                                       |
|-----------------------|---------------------------------------------------------------------------------------------------|
| <b>Cloud Region</b>   | Example: Canada (Central) - ca-central-1 (this must match the region of the Salesforce instance). |
| <b>Core Tenant ID</b> | Paste the <b>Core Tenant ID</b> of your Salesforce Data Share Target.                             |

5. Save the connection but do not test it yet.
6. Gather keys.
  - a. On the **View Connection Details** page, locate and keep open the **Service Principal**, and **Issuer**.
  - b. Select **Certificate > Download > PKCS#8**, to download the PKCS#8 certificate.
7. Finalize the Data Share Target in Salesforce:

| Field                     | Description                                                                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Service Account ID</b> | Paste the <b>Service Principal</b> from Workday.                                                                                                                                      |
| <b>Issuer URL</b>         | Paste the <b>Issuer</b> from Workday.                                                                                                                                                 |
| <b>Public Key</b>         | Copy the contents of the PKCS#8 certificate that you downloaded from Workday. Note: Remove the -----BEGIN PUBLIC KEY----- and -----END PUBLIC KEY----- headers/footers before saving. |

8. Save your changes.
9. Link the Data Share Target to the Data Share in Salesforce.
10. In Workday, on the **Connection Details** page, select **Test Connection**.

## Task 2: Create an External Catalog

An external catalog acts as a virtual window into your external data, allowing you to browse tables in real-time.

1. Access the **Data Catalog** report.
2. Click **+ Create** and select **External Catalog**.
3. Enter the external catalog details:

| Field                        | Description                                                                                                                                                 |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>External Catalog Name</b> | Enter a name for your external catalog.                                                                                                                     |
| <b>Connection</b>            | Select the connection you just created.                                                                                                                     |
| <b>Namespace</b>             | The exact namespace where the Iceberg connector tables reside. This information is provided by Salesforce. Example: <i>enterprise.finance.procurement</i> . |
| <b>Description</b>           | (Optional) The description of the external catalog.                                                                                                         |

4. Click **Save**.
5. Wait for the **Success** notification that the External Catalog was created.
6. Select **Back to Data Catalog**.
7. Select **External Catalog**, and validate that schema sync is successful on your External Catalog.

After you save the External Catalog, Workday automatically runs a **Schema Sync** to:

- Read the metadata of Iceberg tables in the specified Salesforce namespace.
- Register those tables as **External Tables** in Workday.
- Make table structures visible in the Data Catalog without moving or copying data.

The duration of this sync depends on the number and complexity of tables in the namespace.

## Task 3: Explore and Validate External Tables

Explore the external tables registered from the external partner and perform basic schema validation inside Workday.

1. Locate the external partner's external tables.
  - a. Access the **Data Catalog** report.
  - b. In the left-hand menu, select **External Tables**.
  - c. Click on the **Refresh** button.
  - d. In the **External Tables** view, confirm that all Iceberg tables from the external partner are listed under the schema/namespace you selected during external catalog creation.
2. Inspect an external table.
  - a. Right-click or double-click the table name to open **View External Table Details**.
  - b. On the details page, review:
    - Column names
    - Data types
    - Any available column statistics
    - Sample data preview (if present)
3. Verify that:
  - Data types (Example: Numeric vs. string) align with the definitions in the external system.
  - Values appear as expected.

### Next Steps

- [Enable External Data for Analysis](#)
- [Blend External Data with Workday Data](#)
- [Create Visualizations](#)

## Set Up Zero-Copy Data In with Snowflake

### Prerequisites

- [Configure Snowflake](#).
  - Take note of your **Account Locator** value from your account details.
- [Set Up Workday Security Policies for Zero-Copy Data In](#)

### Summary

- [Task 1: Create a Zero-Copy Connection to Snowflake](#)
- [Task 2: Create an External Catalog](#)
- [Task 3: Explore and Validate External Tables](#)

### Task 1: Create a Zero-Copy Connection to Snowflake

You can access Snowflake table metadata without copying data by creating a zero-copy connection from Workday to your Snowflake Iceberg catalog.

1. In Workday, access the [Prism Data Catalog](#) report.
2. Select **+ Create > Connection**.
3. In the **Create Connection** wizard, under **Zero Copy Connectors**, select the **Snowflake** tile.
4. Enter the connection details:

| Field       | Description                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Name</b> | The connection name must be unique in the Data Catalog.<br>Example: Snowflake_ZeroCopy_<first name><last initial>. |

| Field                      | Description                                                                                                                                                                                                                   |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>         | (Optional) A brief description of the connection's purpose.                                                                                                                                                                   |
| <b>Catalog URL</b>         | The URL of the Snowflake data catalog. The URL must use only lowercase characters. Format: <code>https://&lt;account_locator&gt;.snowflakecomputing.com/polaris/api/catalog</code>                                            |
| <b>Catalog Name</b>        | The name of the Snowflake catalog that contains your Iceberg tables.                                                                                                                                                          |
| <b>Authentication Type</b> | Select <b>OAuth 2.0 with Key Pair</b> .                                                                                                                                                                                       |
| <b>Issuer</b>              | Enter initial values. You'll retrieve the <i>PUBLIC_KEY_FP</i> value later in this procedure. Format: <i>ACCOUNT_LOCATOR.SERVICE_USER.PUBLIC_KEY_FP</i> Example: <i>XTQ04827 . ICEFLOW_PRISM_USER . &lt;PUBLIC_KEY_FP&gt;</i> |
| <b>Principal</b>           | The Service Principal from Snowflake. This user will be used to authenticate the Snowflake connection. Format: <i>ACCOUNT_LOCATOR.SERVICE_USER</i> Example: <i>XTQ04827 . ICEFLOW_PRISM_USER</i>                              |
| <b>Token Endpoint</b>      | The OAuth2 token endpoint from Snowflake. Format: <i>Catalog_URL/v1/oauth/tokens</i> Example: <code>https://&lt;account_locator&gt;.snowflakecomputing.com/polaris/api/catalog/v1/oauth/tokens</code>                         |
| <b>Scopes</b>              | The OAuth2 scopes from Snowflake. The scopes define what the Principal has access to when connected to Snowflake.                                                                                                             |

5. Save the connection but do not test it yet. Before you test the connection, you must complete the certificate/key setup in the subsequent steps.

After the save is complete, the **Authentication Certificate** window displays.

6. From the **Authentication Certificate** window, download the certificate in **PKCS#8** format.

If you skipped downloading the certificate from the **Authentication Certificate** window, follow these steps:

- a. Access the **Data Catalog** report.
- b. From the **Connections** list, open the connection you just created. Example: *Snowflake\_ZeroCopy\_<first name><last initial>*
- c. On the **View Connection Details** page, locate the **Authentication Certificate** section.
- d. Click **Certificate > Download**.
- e. Select **PKCS#8**.

This file contains the RSA public key that you'll register with the Snowflake service user.

7. Set the RSA public key on the Snowflake service user.

- a. Sign in to your Snowflake account.
- b. Create a **SQL Worksheet** with sufficient privileges to alter the service user.
- c. Type and run an ALTER USER command to set the RSA public key for your service user.

Example:

```
ALTER USER ICEFLOW_PRISM_USER SET RSA_PUBLIC_KEY = '<Content of
Downloaded PKCS#8>';
```

If your environment uses a second RSA key, use RSA\_PUBLIC\_KEY\_2 instead. Example:

```
ALTER USER ICEFLOW_PRISM_USER SET RSA_PUBLIC_KEY_2 = '<Content of
Downloaded PKCS#8>';
```

- d. Run the following command to inspect the user and retrieve the public key fingerprint:

```
DESCRIBE USER ICEFLOW_PRISM_USER;
```

- e. In the result set, locate the value of either:
  - RSA\_PUBLIC\_KEY\_FP
  - RSA\_PUBLIC\_KEY\_2\_FP (if you used the second key)
- f. Copy this fingerprint value.

8. Update the Workday connection Issuer with the fingerprint value.

- a. In Workday, open your Snowflake zero-copy connection.
- b. From the related actions menu of the connection, select **Prism Connection > Edit Connection**.
- c. Update the **Issuer** field to append the RSA public key fingerprint.

Use this final Issuer format:

```
<ACCOUNT_LOCATOR> . <SERVICE_USER> . <PUBLIC_KEY_FP>
```

Where:

- *ACCOUNT\_LOCATOR* is your Snowflake account locator.
- *SERVICE\_USER* is the service user you altered above.
- *PUBLIC\_KEY\_FP* is the value copied from RSA\_PUBLIC\_KEY\_FP (or RSA\_PUBLIC\_KEY\_2\_FP).

- d. Save the updated connection.

9. Test the connection.

- a. On the **Connection Details** page, select **Test Connection**.
- b. If the test fails, verify:
  - Catalog URL format.
  - Token endpoint URL.
  - Service user and account locator values.
  - RSA public key is correctly set on the user.
  - Issuer format, including the public key fingerprint.

## Task 2: Create an External Catalog

An external catalog acts as a virtual window into your external data, allowing you to browse tables in real-time.

1. Access the **Data Catalog** report.
2. Select **+ Create > External Catalog**.

3. Enter the external catalog details:

| Field                        | Description                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>External Catalog Name</b> | Enter a name for your external catalog.                                                                                                                    |
| <b>Connection</b>            | Select the external partner zero-copy connection you just created.                                                                                         |
| <b>Namespace</b>             | The exact namespace where the Iceberg connector tables reside. This information is provided by Snowflake. Example: <i>enterprise.finance.procurement</i> . |
| <b>Description</b>           | (Optional) The description of the external catalog.                                                                                                        |

4. Click **Save**.
5. Wait for the **Success** notification that the External Catalog was created.
6. Select **Back to Data Catalog**.
7. Select **External Catalog**, and validate that schema sync is successful on your External Catalog.

After you save the External Catalog, Workday automatically runs a **Schema Sync** to:

- Read the metadata of Iceberg tables in the specified Snowflake namespace.
- Register those tables as **External Tables** in Workday.
- Make table structures visible in the Data Catalog without moving or copying data.

The duration of this sync depends on the number and complexity of tables in the namespace.

### Task 3: Explore and Validate External Tables

Explore the external tables registered from the external partner and perform basic schema validation inside Workday.

1. Locate the external partner's external tables.
  - a. Access the **Data Catalog** report.
  - b. In the left-hand menu, select **External Tables**.
  - c. Click on the **Refresh** button.
  - d. In the **External Tables** view, confirm that all Iceberg tables from the external partner are listed under the schema/namespace you selected during external catalog creation.
2. Inspect an external table.
  - a. Right-click or double-click the table name to open **View External Table Details**.
  - b. On the details page, review:
    - Column names
    - Data types
    - Any available column statistics
    - Sample data preview (if present)
3. Verify that:
  - Data types (Example: Numeric vs. string) align with the definitions in the external system.
  - Values appear as expected.

### Next Steps

- [Enable External Data for Analysis](#)
- [Blend External Data with Workday Data](#)
- [Create Visualizations](#)

## Enable External Data for Analysis

### Prerequisites

- Set up Zero-Copy Data In with your external system:
  - [Databricks](#)
  - [Google Cloud](#)
  - [Salesforce](#)
  - [Snowflake](#)
- These security domains in the Prism Analytics functional area:
  - *Prism: External Tables Manage*
  - *Prism: External Tables Owner Manage*

### Context

You can make an external table available as a report data source inside Workday. When the schema is synced, activate the tables you wish to use for reporting.

### Steps

1. Access the [Prism Data Catalog](#) report.
2. In the left navigation menu of the Data Catalog, click **External Tables**.
3. Locate the table you need. Right-click or double-click the table name to open the **View External Table Details** page.

Note: Tip: Verify data integrity before proceeding.

4. Click the **Quick Action** button on the top right and select **Edit External Table**.
5. On the **Edit External Table** page, select the **Enable for analysis** check box.
6. Click **Save**.

### Result

Enabling the external table for analysis transforms the external table into a reporting-ready data source.

### Next Steps

You can use the external table as a data source in Prism, Discovery Boards, or other analytical experiences according to your own scenario.

## Blend External Data with Workday Data

### Context

Use Prism Data Management to transform and aggregate data from your external tables (and optionally Workday tables) into an analysis-ready dataset.

## Steps

1. Create a Derived Dataset.
  - a) Access the [Prism Data Catalog](#) report.
  - b) Select **+ Create > Derived Dataset**.
  - c) Select one or more source tables or datasets that contain the measures and dimensions you want to analyze.  
Example: A Workday dataset and an external table.
  - d) Enter a meaningful **Dataset Name** that clearly describes the content and purpose of the dataset.
  - e) Select **Save** to open the dataset in the **Edit Dataset Transformations** page.
2. Join Workday and external data.
  - a) On the **Edit Dataset Transformations** page, in the top panel, select **Add Stage > Join**.
  - b) In the **Primary Pipeline**, keep your main table or dataset (Example: Workday data).
  - c) In the **Select a Pipeline** dropdown list, click **Add Another Pipeline** and select the external table from your **External Catalog**.
  - d) Use business keys (Examples: Worker IDs, transaction IDs, cost center IDs) to define the join.
  - e) (Optional) Use Copilot suggestions to identify high-confidence field pairs and confirm the ones you want to use.
  - f) Select an appropriate **Join Type** (Example: **Inner Join** when you only want records that exist in both sources).
  - g) In **Select Fields**, include only the fields that are relevant to your analysis.
3. Create calculated fields.
 

To prepare your measures for analysis:

  - a) On the **Edit Dataset Transformations** page, select **Add Field**.
  - b) Use the **Prism Copilot Expression Builder** to:
    - Standardize numeric measures (Examples: Convert text to numeric, or normalize to a common currency or unit).
    - Derive business metrics (such as totals, ratios, flags, or categories).
    - Clean or format text and date fields into more report-friendly values.
  - c) Give each calculated field a clear, business-friendly name.
  - d) Select **Save** in the upper-right corner to save the updated dataset definition.
4. Reshape and aggregate the data.
 

Depending on your reporting needs, you may want to reshape or pre-aggregate the data:

  - Use stages such as **Unpivot** or **Pivot** to restructure multiple related measures into a format that is easier to slice and compare.
  - Use **Aggregate** or Grouping-style transformations to pre-roll up data at a useful grain (Examples: By worker, cost center, time period, or product).

After making changes, use the preview pane to confirm that:

  - The key dimensions (Examples: Organizations, dates, entities) appear as expected.
  - The measures are in the correct format and at the correct level of detail.
5. Publish the dataset.
 

When you are satisfied with the structure and contents of the dataset:

  - a) From the **View Dataset Details** page, select **Quick Actions > Publish**.
  - b) Select **Submit** to start the publish operation.
  - c) Use the **Refresh** icon to update the status until the publishing operation is complete and shows as Published.



## Result

The published dataset is now ready to be used as a data source in:

- Custom Reports
- Dashboards
- Discovery Board Visualizations
- WQL Queries

## Create Visualizations

### Context

You can build Discovery Board visualizations that use your external tables and the derived datasets.

### Steps

#### 1. Create a Discovery Board.

- a. From the Workday header, select your profile picture and then select **Drive**.
- b. Navigate to a folder.
- c. Select **+ New > Discovery Board**.
- d. Enter a name such as ZeroCopy\_Analysis\_.
- e. Click **Create**.

A new Discovery Board opens with a single blank visualization.

#### 2. Add a visualization using only external data.

- a. In the Discovery Board, open the Data Sources panel.
- b. In the Search Data Sources box, search for the external data source associated with your external catalog.
- c. Select this external data source.
- d. Open the Builder Panel to view available fields.
- e. To the right of the Search Fields box, select the down arrow next to the default visualization type and choose a chart that fits your data (for example, a bar chart, column chart, or donut chart).
- f. Drag one or more numeric measures from the external source into the appropriate blocks (for example, Y-Axis or Angle).
- g. Drag a dimension field (such as a category, organization, or level) into the X-Axis, Color, or Filter blocks to segment the measure.

This creates a visualization that uses only the external data queried via Zero Copy.

Click **Save**.

### 3. Add a visualization using the blended dataset.

- a. On the Discovery Board, select **Add Viz** to add another visualization.  
In the Data Sources panel, select the Back arrow and search for the published dataset you created.
- b. Select that dataset (for example, a dataset that blends Workday and external data).
- c. Open the Builder Panel to display the dataset fields.
- d. Choose a chart type that helps compare your key measures across categories or over time (for example, a column chart).
- e. Drag your primary measure (for example, a total or sum) into the Y-Axis block.
- f. Drag a relevant dimension (for example, a category, period, or organization field) into the X-Axis block.
- g. Optionally add more dimensions to Color or Filter to allow users to segment the results.

This creates a visualization that showcases the blended dataset.

Click **Save**.

### 4. Review and refine the board.

Review the Discovery Board:

Confirm that the first visualization is based only on external data.

Confirm that the second visualization uses your published, blended dataset.

Check that labels, titles, and legends make it clear what each visualization represents.

### Next Steps

You can:

- Add filters (such as time period, organization, or population).
- Adjust formatting and layout so the board is easy to interpret.
- Share the board according to your needs.

## Troubleshooting: Zero-Copy Data In

### Unable to Create Connection

Customers cannot create a new connection from the Prism Data Catalog.

Cause: Missing required Prism security permissions.

Solution: Verify and assign the required security domain.

Ensure the user has *Prism: Manage Connection* domain in the Prism Analytics functional area.

After the domain is assigned, log out and log back in before retrying the connection creation.

### Test Connection Fails

Connection test fails when validating credentials during setup.

Cause: Incorrect OAuth credentials or misconfigured token endpoint.

Solution: Validate credentials and authentication configuration.

Confirm the Client ID, Principal, and Token Endpoint values match those provided by the external platform's documentation.

### Connection Created but No External Catalog Can Be Created

Users can create a connection but cannot create an external catalog.

Cause: Missing External Catalog Create permission.

Solution: Assign the required External Catalog security domain.

Verify the user has *Prism: External Catalog Create* domain.

Once granted, retry creating the external catalog from the Data Catalog UI.

### External Catalog Creation Fails with Validation Error

Catalog creation fails with a generic validation or save error.

Cause: Invalid or unsupported namespace format.

Solution: Validate the namespace path.

Confirm the namespace exactly matches the external Iceberg catalog namespace.

### External Catalog Created but Schema Sync Never Starts

Catalog is created successfully but schema sync does not initiate.

Cause: Connection misconfiguration or insufficient catalog-level permissions in the external system.

Solution: Validate external catalog access.

Ensure the identity used by the connection has sufficient permissions to list schemas and tables in the external catalog.

Once access is corrected, trigger a catalog re-sync to refresh the metadata.

### External Catalog Created but Tables Are Not Visible

External catalog is created successfully, but no tables appear after schema sync.

Cause: Incorrect namespace or schema sync did not complete successfully.

Solution: Validate namespace and re-sync the catalog.

Double-check that the namespace matches the exact Iceberg namespace. If correct, manually trigger a catalog sync using the External Catalog Manage actions.

### Schema Sync Fails or Times Out

Schema sync does not complete or fails intermittently.

Cause: Temporary connectivity issues or large schema size during EA.

Solution: Retry sync and limit initial scope.

Retry the schema sync after a short wait. For large catalogs, consider starting with a smaller namespace during EA to validate end-to-end functionality before expanding scope.

### External Tables Visible but Cannot Be Enabled for Analysis

Tables appear in the External Tables list but the **Enable for analysis** option is unavailable.

Cause: Missing External Tables Manage permission.

Solution: Assign the External Tables Manage domain.

Ensure the user has *Prism: External Tables Manage* domain. This permission is required to edit table metadata and enable tables for analysis.

**Join Fails When Creating a Derived Dataset**

Join between external Iceberg tables and Workday data fails.

Cause: Data type mismatch or incompatible join keys.

Solution: Align join key data types and validate data.

Verify the join columns have compatible data types (for example, string vs numeric).

Use the External Table preview to inspect sample data and confirm key consistency before joining.

**Reports Fail After Publishing Derived Dataset**

Reports fail to run or return errors after dataset publication.

Cause: Dataset not fully published or upstream changes in external schema.

Solution: Republish dataset and validate schema.

Confirm the derived dataset is successfully published. If the external schema changed, re-sync the external catalog and republish the dataset before running reports again.