

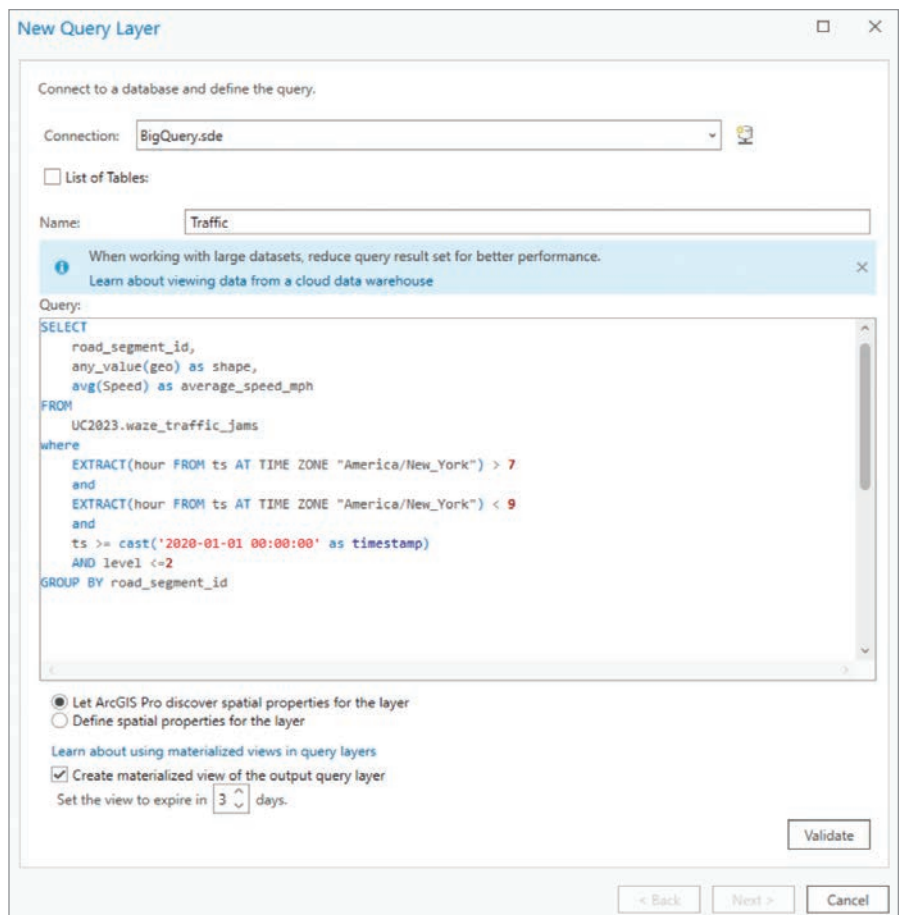
Leverage Cloud Data Warehouses with ArcGIS Pro and ArcGIS Enterprise


By Sarah Hanson

Cloud data warehouses are scalable, resilient, secure, and performant solutions for storage that are available in public clouds as managed services. Cloud elasticity means that they can scale infinitely and automatically—in terms of both storage capacity and processing power—to support demand at any given time. Cloud data warehouses offer a pay-as-you-go model so that organizations only pay for the resources they use as they use them.

For all these reasons, cloud data warehouses are becoming an integral part of many organizations, especially those that have adopted or are moving toward a cloud-first strategy. The 2021 fall release of ArcGIS Pro 2.9 and ArcGIS Enterprise 10.9.1 added functionality that allows you to directly connect to and access data stored in cloud data warehouses. Support for cloud data warehouses creates an opportunity for organizations using ArcGIS Pro and ArcGIS Enterprise to integrate yet another type of cloud-native service with their systems. Like database as a service (DBaaS) and cloud storage (blob storage), cloud data warehouses are available in public clouds as managed services.

Cloud data warehouses are not designed for transactions, or row-based editing. Instead, they are designed to support analysis. Most commonly, organizations move to the cloud and use cloud data warehouses to archive their historical records. These records could have been previously maintained in a transactional database, a relational database management system (RDBMS), or some other tabular format. Cloud data warehouses are also utilized for writing Internet of Things (IoT) sensor data,





- ↑ Esri has enhanced many chart types available in ArcGIS Pro.
- ← SQL queries are used to create a query layer definition.

which can generate tremendous volumes of data in a short amount of time.

Get Started by Creating a Query Layer

You can connect from ArcGIS Pro to three of the most widely adopted cloud data warehouses: Amazon Redshift, Google BigQuery, and Snowflake. Once a connection is established, you can create a new query layer (spatial) or a query table (nonspatial) to extract the features and attributes you want. For example, an analyst for a retail chain could create a query layer that aggregates sales information by product category and store location. Or a transportation organization could visualize where traffic congestion occurs during the morning rush hour.

To create query layers from any database, whether a relational database or a cloud data warehouse, you need to be familiar with Structured Query Language

(SQL). SQL queries are used to create query layer definitions.

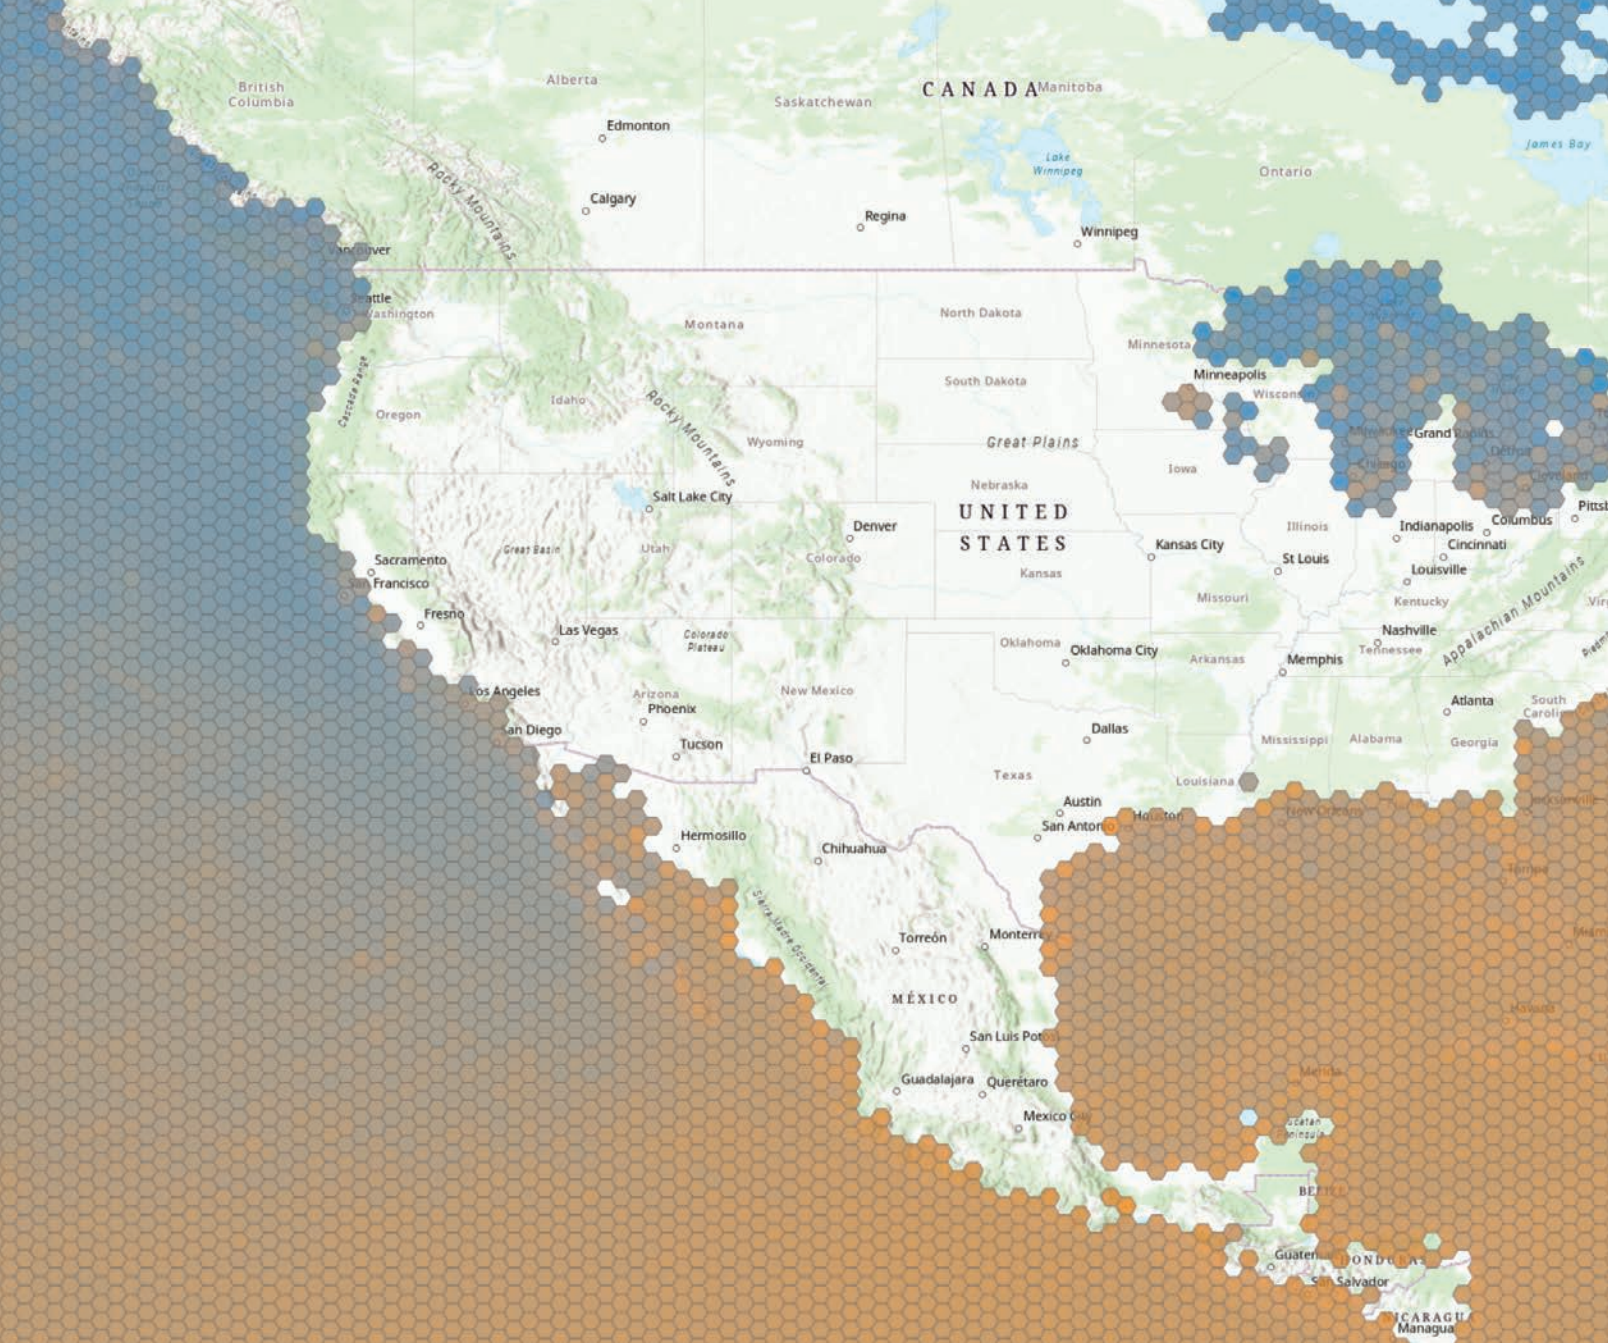
Revealing Patterns and Trends in Your Data

ArcGIS Pro also includes tools, such as charting and feature binning, that can be used to detect patterns and trends in data. Imagine a bar chart showing annual average precipitation amounts by week. With the map extent set to the entire United States, the chart may not be very revealing. However, as you zoom into a more localized area, the chart updates, showing the average precipitation for a specific area or region throughout the year. Suddenly, you have a charting tool that is exploratory and interactive and allows you to gain new insight. Chart results can be updated by simply panning around the map.

Feature binning is another great way to detect patterns and trends, especially when you are working with datasets that

have thousands or even millions of point features. To try this technique, run the Enable Feature Binning geoprocessing tool to use database-driven feature binning or set the display properties on a layer to use application-driven feature binning.

Most commonly, organizations move to the cloud and use cloud data warehouses to archive their historical records.



Esri recommends using database-driven feature binning for datasets that contain a large number of features because bin information is stored in the database, so this results in better performance. In addition, you can create summary statistic fields such as mean, min, max, and standard deviation for the numeric fields in your dataset and use those fields to symbolize feature bins.

Sharing Web Layers to ArcGIS Enterprise

Esri has added support for BigQuery, Redshift, and Snowflake in ArcGIS Enterprise and ArcGIS Pro. Users with permissions to publish to ArcGIS Enterprise

may share their query layers to ArcGIS Enterprise. Like sharing other datasets by reference, your cloud data warehouse must be registered as a data store with ArcGIS Enterprise. When sharing, you will need to determine two things: the type of web layer that you want to share and how ArcGIS Server should access the data.



Considerations for Sharing

Query layers that reference data in cloud data warehouses can be shared as map image layers to ArcGIS Enterprise 10.9.1 or later. Support for sharing web feature layers was added with ArcGIS Enterprise 11.0. This release supports

sharing a read-only feature layer for Google BigQuery. In ArcGIS Enterprise 11.1, support for sharing a read-only feature layer was added for Snowflake. Feature layer support for Amazon Redshift is planned for a future release of ArcGIS Enterprise.

While neither the map image layer nor the features layer supports editing in the context of cloud data warehouses, there are some differences between these layer types.

If you're using complex symbology or dynamic joins or relates, you will want to publish a map image layer. If you want the web layer to be able to be used as input to analysis tools, you will want to publish a



- ↑ Data that references cloud data warehouses can be displayed in ArcGIS Dashboards.
- ← Feature binning can be used to detect patterns and trends in data. These bins were aggregated from more than 350 million individual point observations.

cloud data warehouse each time the layer is accessed. Both materialized views and snapshots store the SQL query result set as a persisted table (either in the cloud data warehouse or the relational data store). While materialized views within the cloud data warehouse will update automatically if the underlying data changes, snapshots are static and must be explicitly refreshed from the web layer's item page in the ArcGIS Enterprise portal.

Refreshing the snapshot will overwrite the data stored in the relational data store with the latest results set. Updates to the snapshot require a manual refresh, but you may see the best performance using this option because the web layer is referencing data in the relational data store and does not need to reach out to the cloud data warehouse when a request is made.

The option you choose for accessing data in your cloud data warehouse will depend on considerations related to cost, performance, and the frequency of updates to the underlying data. In general, sharing a materialized view or a snapshot will result in better performance and incur less cost.

Summary

If you have a cloud data warehouse and want to get started integrating it with ArcGIS Pro and ArcGIS Enterprise, you can start immediately. ArcGIS is a powerful tool for enhancing and unveiling patterns and trends in your data to make better, more

informed decisions, and cloud data warehouses are a great way to get the most out of your ArcGIS system.

About the Author

Sarah Hanson is a product manager for ArcGIS Enterprise, with a focus on cloud data management. In her free time, she loves to travel and spend time outdoors.

feature layer. Once published, your cloud data warehouse-backed web layers can be used in your maps and apps or shared across the web just like any other map image layer. Esri has noticed that many users want to share data that references their cloud data warehouse for use in ArcGIS Dashboards, a web-based application that presents analytics using interactive maps, charts, and tables.

You must also consider how you want the data in your cloud data warehouse to be accessed. There are three choices: directly, by creating a materialized view, or by creating a snapshot. When data is accessed directly, the SQL query will be run in the

Reach GIS Professionals, Managers, & Developers

Advertise in *ArcUser* and reach more than 450,000 people who are actively using GIS and Esri technology in their organizations.

ads@esri.com

Copyright © 2023 Esri. All rights reserved.