Easy Data Lakehouse Management with Dremio Arctic's Automatic Data Optimization - Q&A

While cloud data lakes address the need to efficiently store large volumes of structured, semi-structured, and unstructured data, they have traditionally lacked the data management and data governance capabilities that have tied enterprise data teams to data warehouse architectures. In this episode, learn how Dremio Arctic, a lakehouse management service, delivers automatic data optimization features that simplify data management and enable high-performance analytics directly on data in the data lake. We'll cover:

- The open data lakehouse architecture, and the importance of a lakehouse management service like Dremio Arctic.
- Dremio Arctic's data optimization capabilities.
- How these features ensure high performance analytics and optimal storage footprint while reducing the management burden for data teams.

Jeremiah Morrow
Director, Product Marketing
Dremio

Alex Merced
Developer Advocate
Dremio

Show Q&A

1. **In a data mesh, each domain gets one catalog? Or one catalog for iceberg, and one for hive etc. That is each domain owner will have a catalog of his own?**
   Answer: You can either have one Arctic catalog per product or one traditional Dremio space per product

2. **I am a data product owner. I have iceberg tables. So, I get one catalog. Then I add some hive tables, or others. I cannot use my iceberg catalog. So, I have to create another one for hive. It is a bit annoying. Often I name the my catalog with my business name?**
   Answer: Currently Arctic catalogs only catalog Iceberg tables, so if all the data is in Iceberg that Arctic catalogs make sense as a unit within a data mesh, if multiple sources using the traditional Dremio spaces feature would fit the bill better.

3. **What does Arctic do differently than OS Iceberg?**
   Answer: Arctic is a catalog that tracks catalog level commit history meaning you can rollback the entire catalog, tag commits for reproducibility and using branching & merging to Isolate ETL and carry out multi-table transactions.