

Streaming History

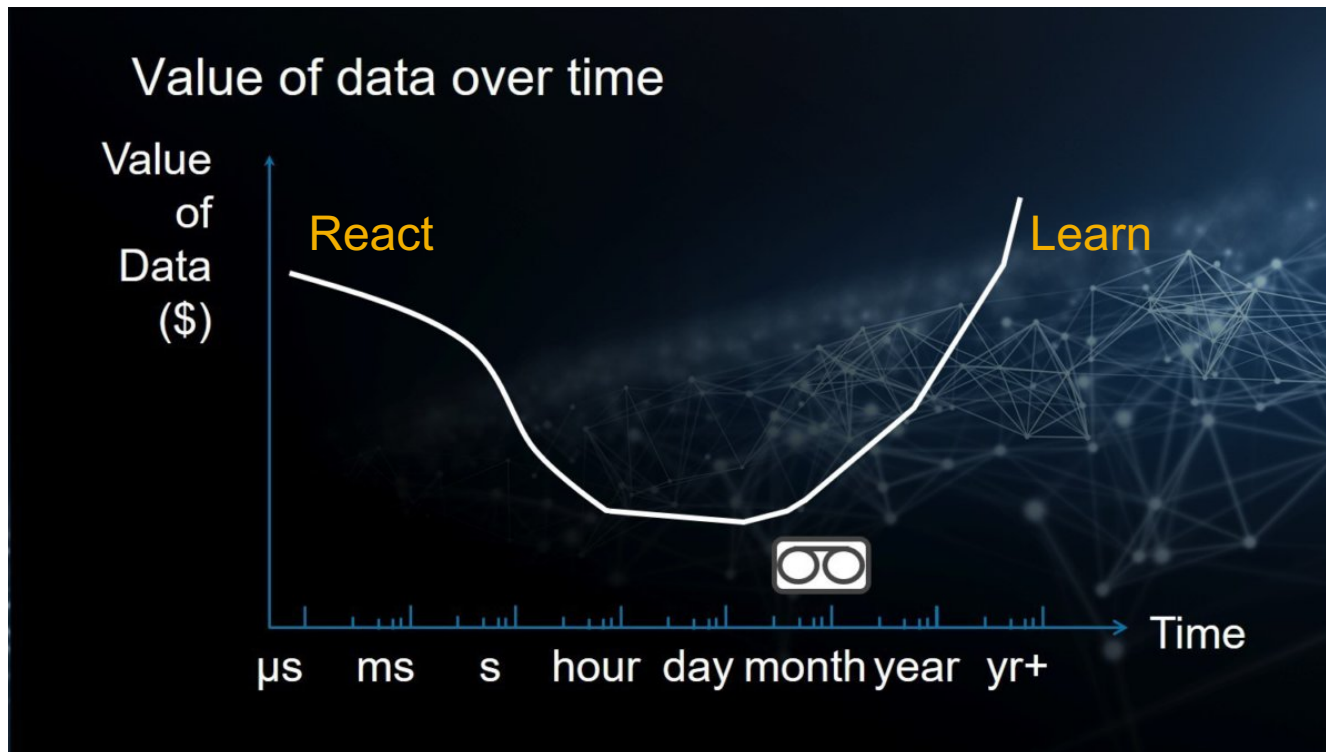
Minimizing Complexity to Maximize Value

Russ Caldwell

Dell EMC • Product Management

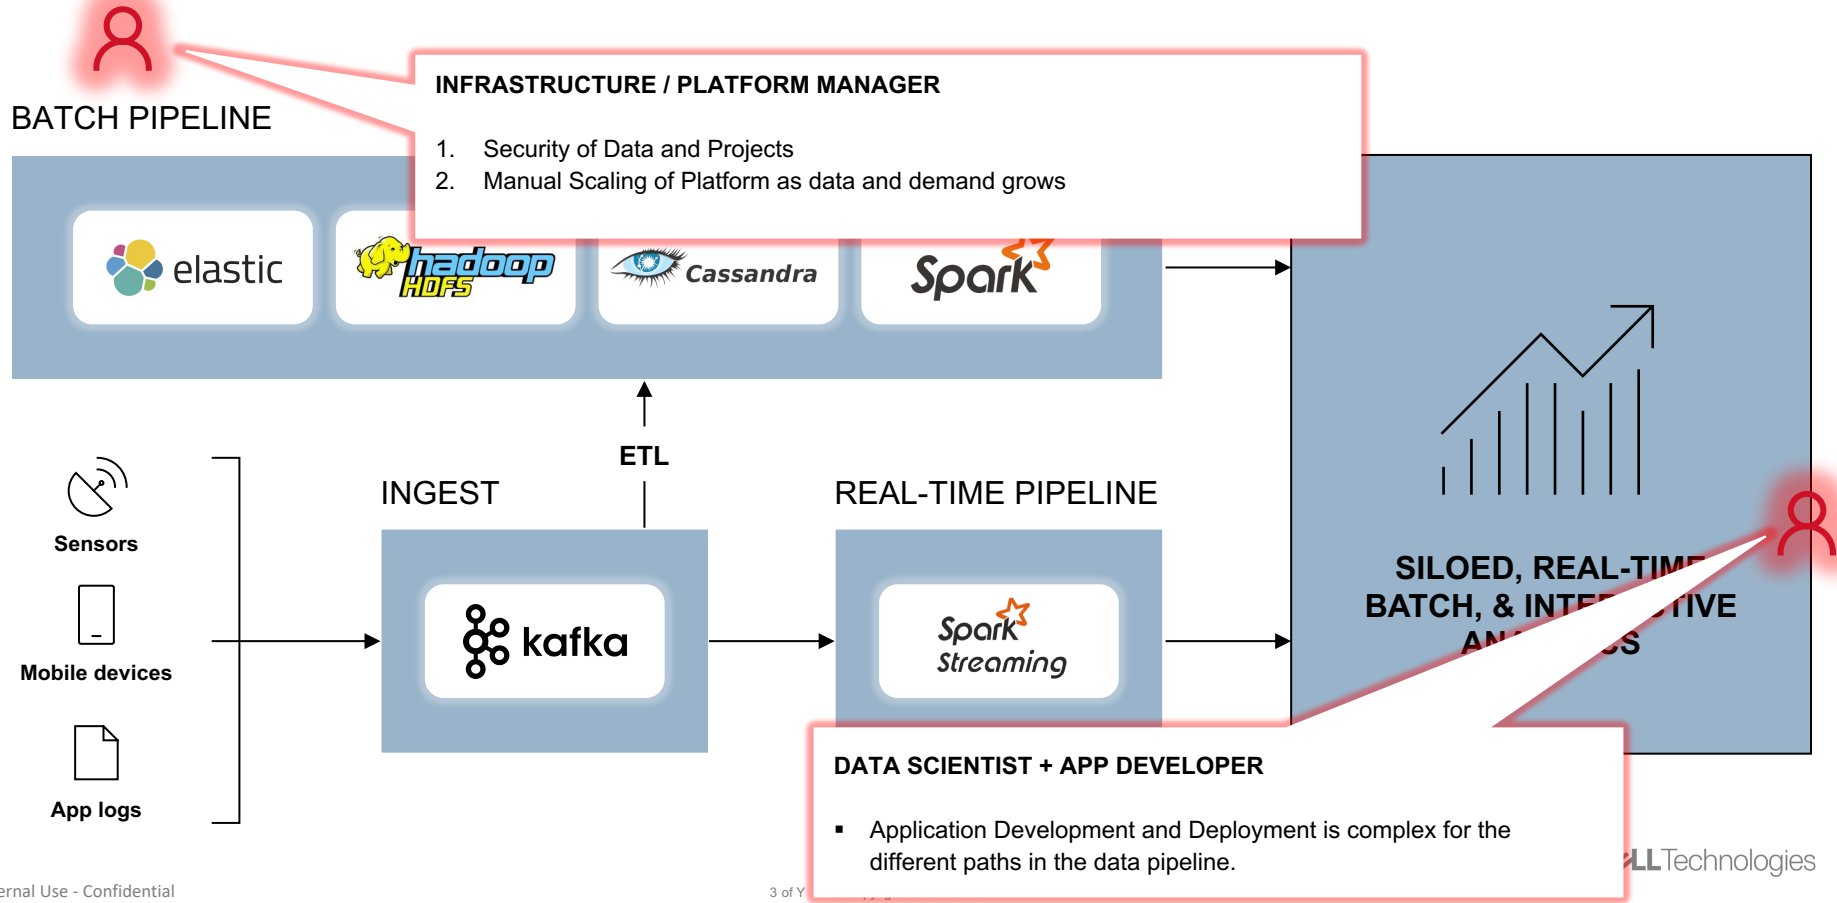
DELLTechnologies

Value of data over time

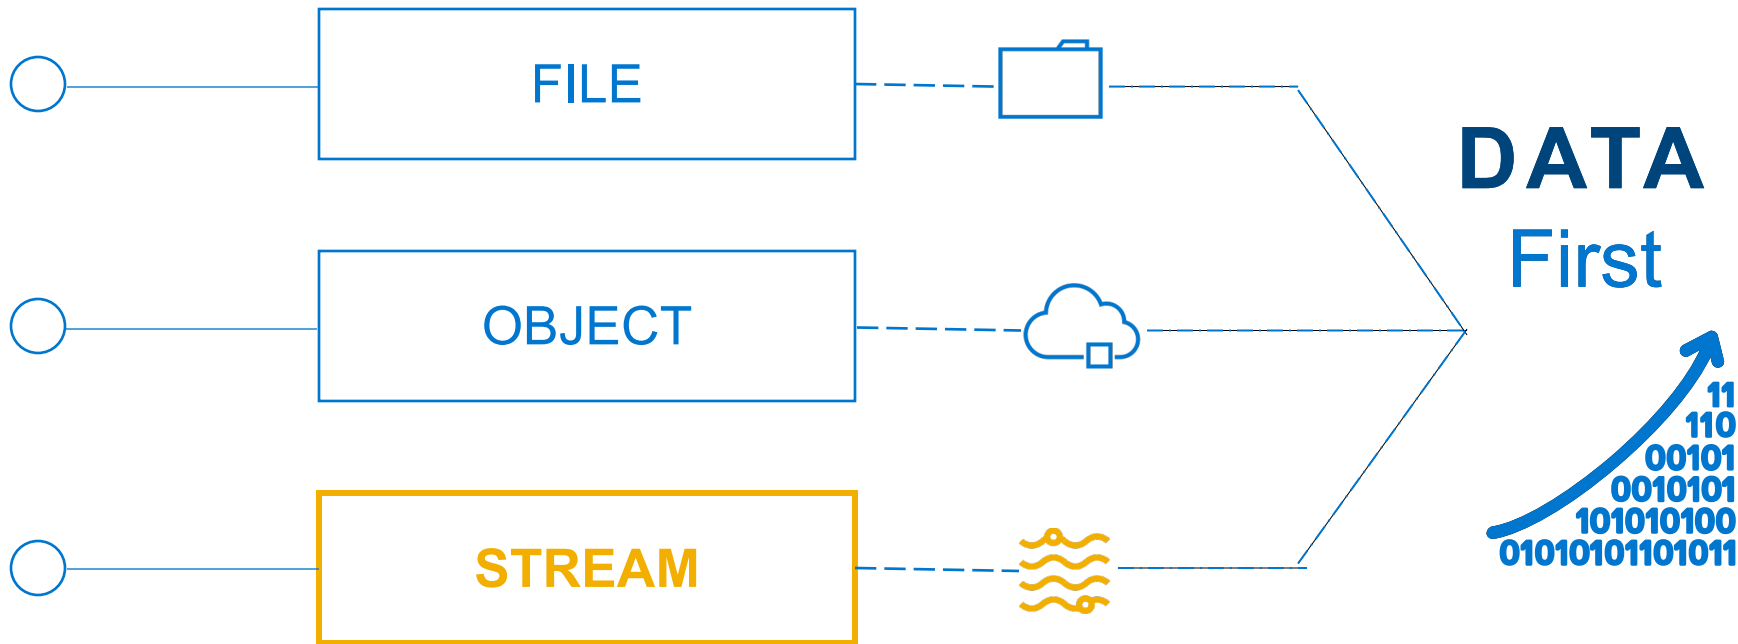


<https://www.cio.com/article/3314742/trends-in-big-data-and-artificial-intelligence-data.html>

Today's DIY stream solutions are complex



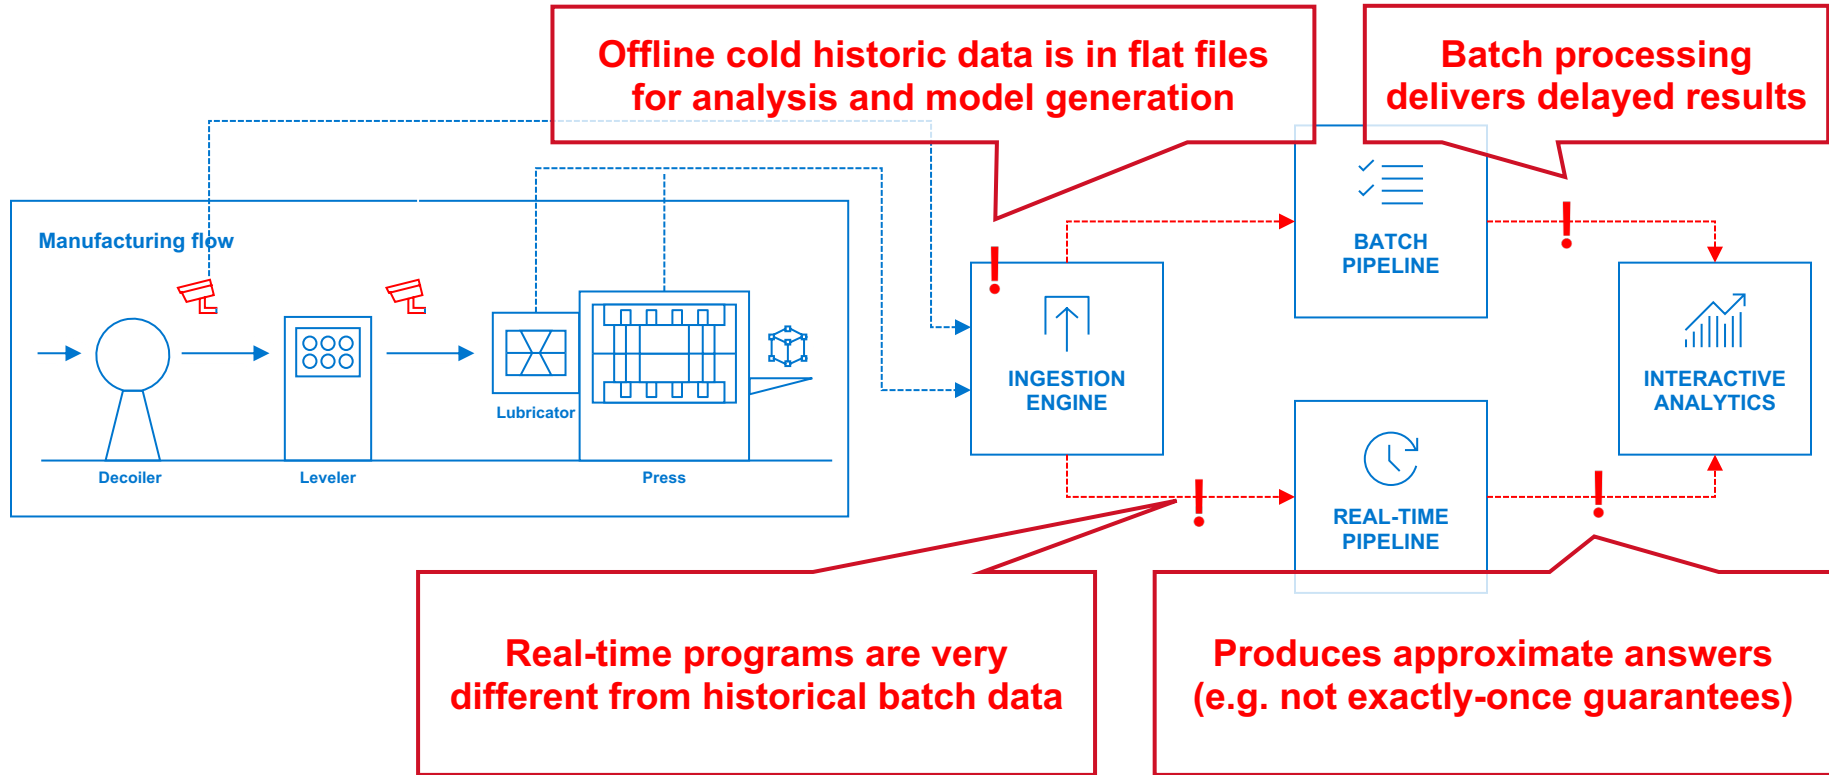
Stream: A First-Class Data Primitive





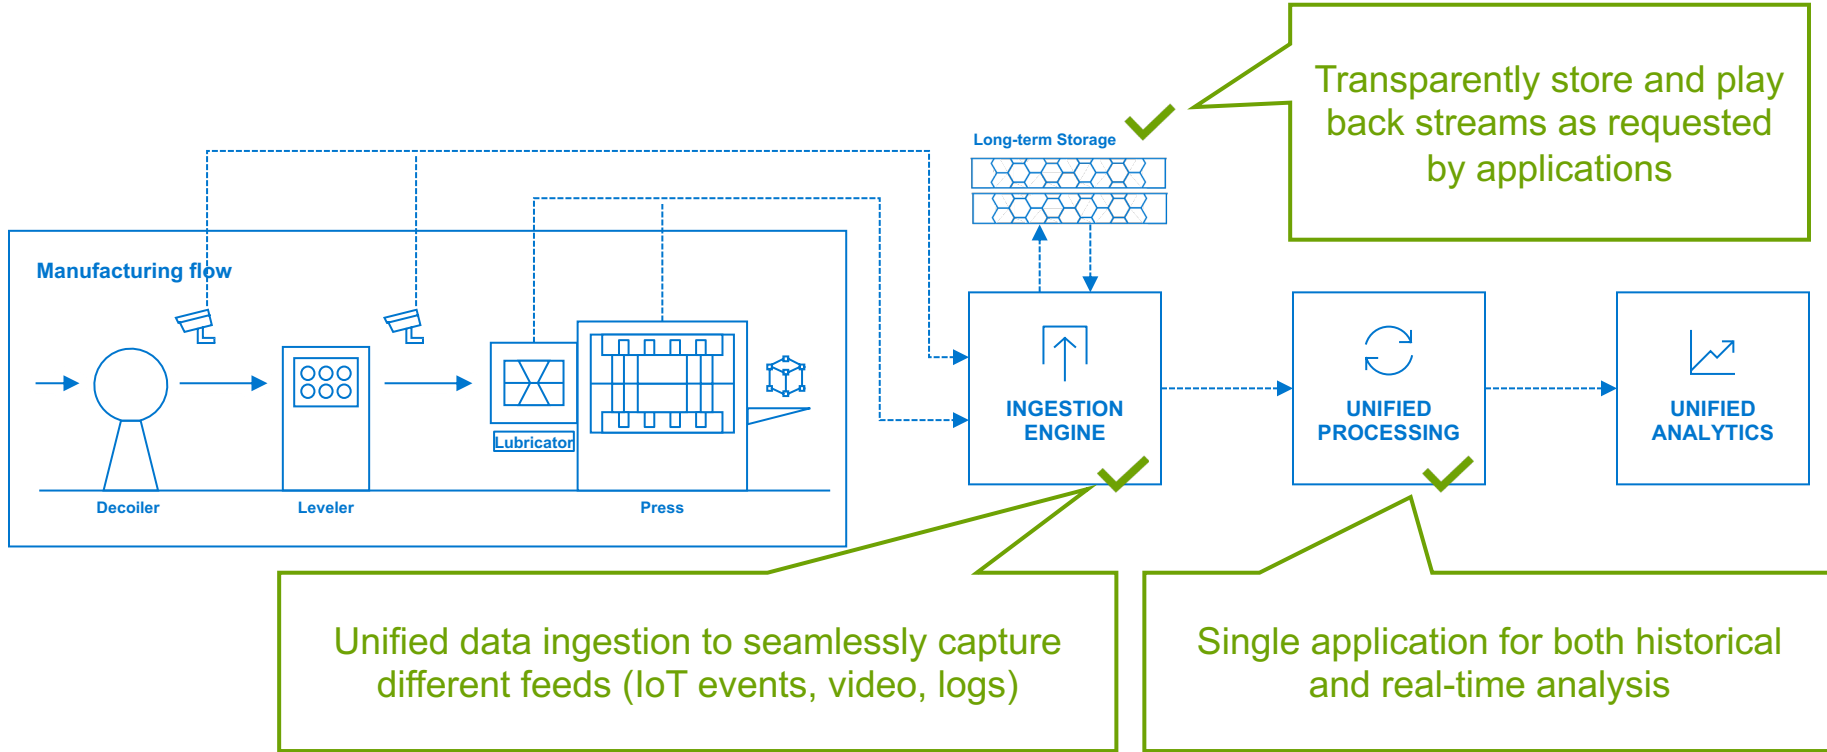
Industrial Sensor Data Anomaly Detection

DIY Lambda Architecture






Industrial Sensor Data Anomaly Detection



Streaming Storage

Architectural Streaming Goals



A. Durable Unbounded Data Streams

- Data is **replicated and persisted to disk** before being acknowledged

B. Low-Latency Response AND Long-term Storage

- **Tier 1 Cache** for fast real-time reads, < 10ms writes
- **Long-term (Tier 2) Storage** should provide transparent access to unlimited historical data

C. Strict ordering guarantees, Exactly-once semantics & Transactions

- **In-Order and exactly-once** of events guaranteed
- SQL-like **Rollback and Commit** transactions within data streams

D. Enterprise Scale, Performance and Support

- Support **Tens of Millions** of streams
- **Dynamic partitioning** of streams based on load and throughput SLO
- Size is **not bounded** by the capacity of a single node
- Can easily develop and deploy production solutions


Storage Types

Traditional Apps/Middleware

Streaming Apps/Middleware


BLOCKS

- Structured Data
- Relational DBs



FILES

- Unstructured Data
- Pub/Sub
- NoSQL DBs




OBJECTS

- Unstructured Data
- Internet Friendly (REST)
- Scale over Semantics
- Geo



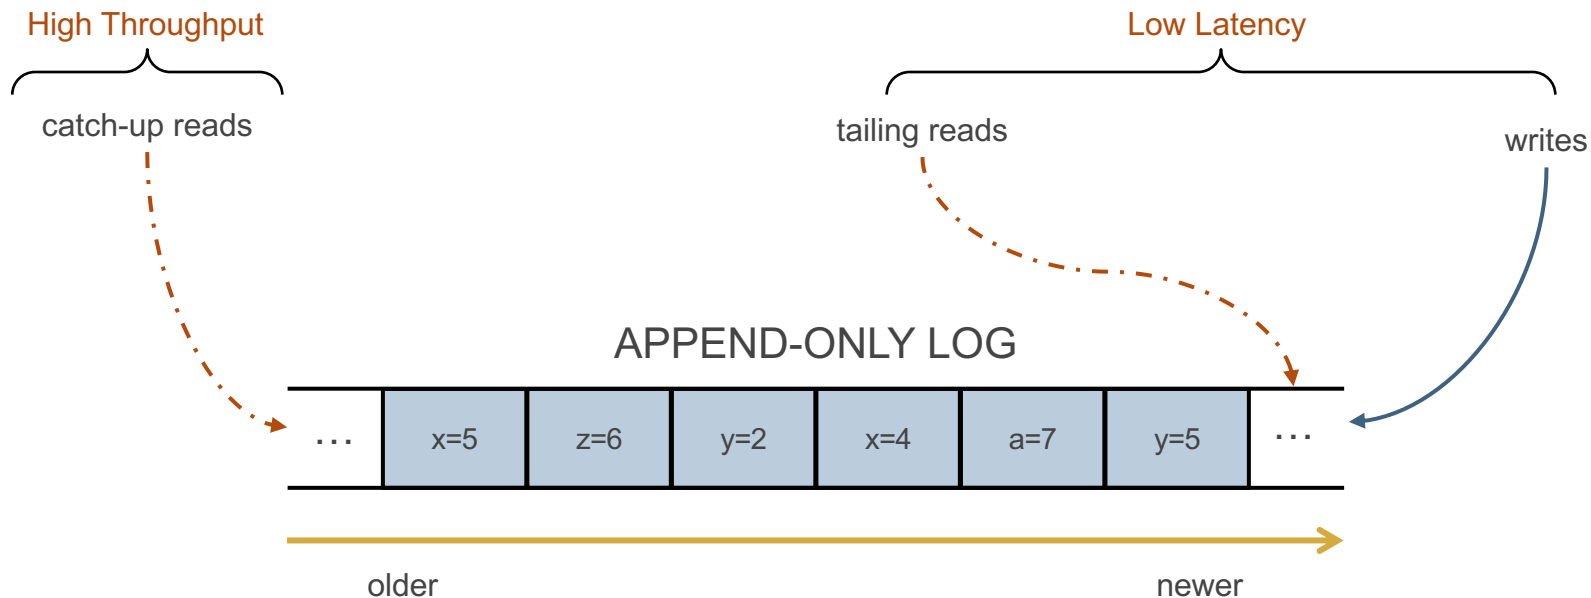
LOGS

- **Append-only**
- **Low-latency**
- **Tail Read/Write**




The Importance of Log Storage


The Fundamental Data Structure for Scale-out Distributed Systems



Historical File/Object vs Stream Data Primitive




Transparent Auto-Scaling



- Still guarantees in-order events
- Scales up & down based on ingestion rates

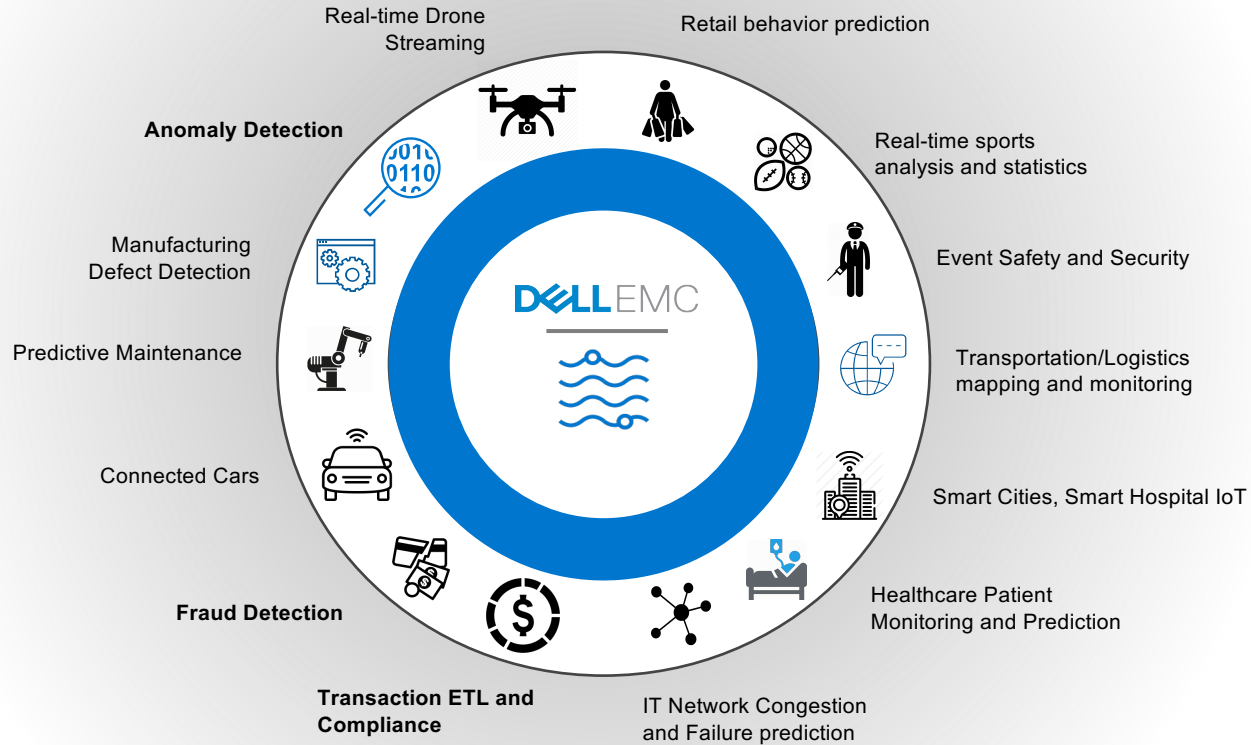


Developing with Streams



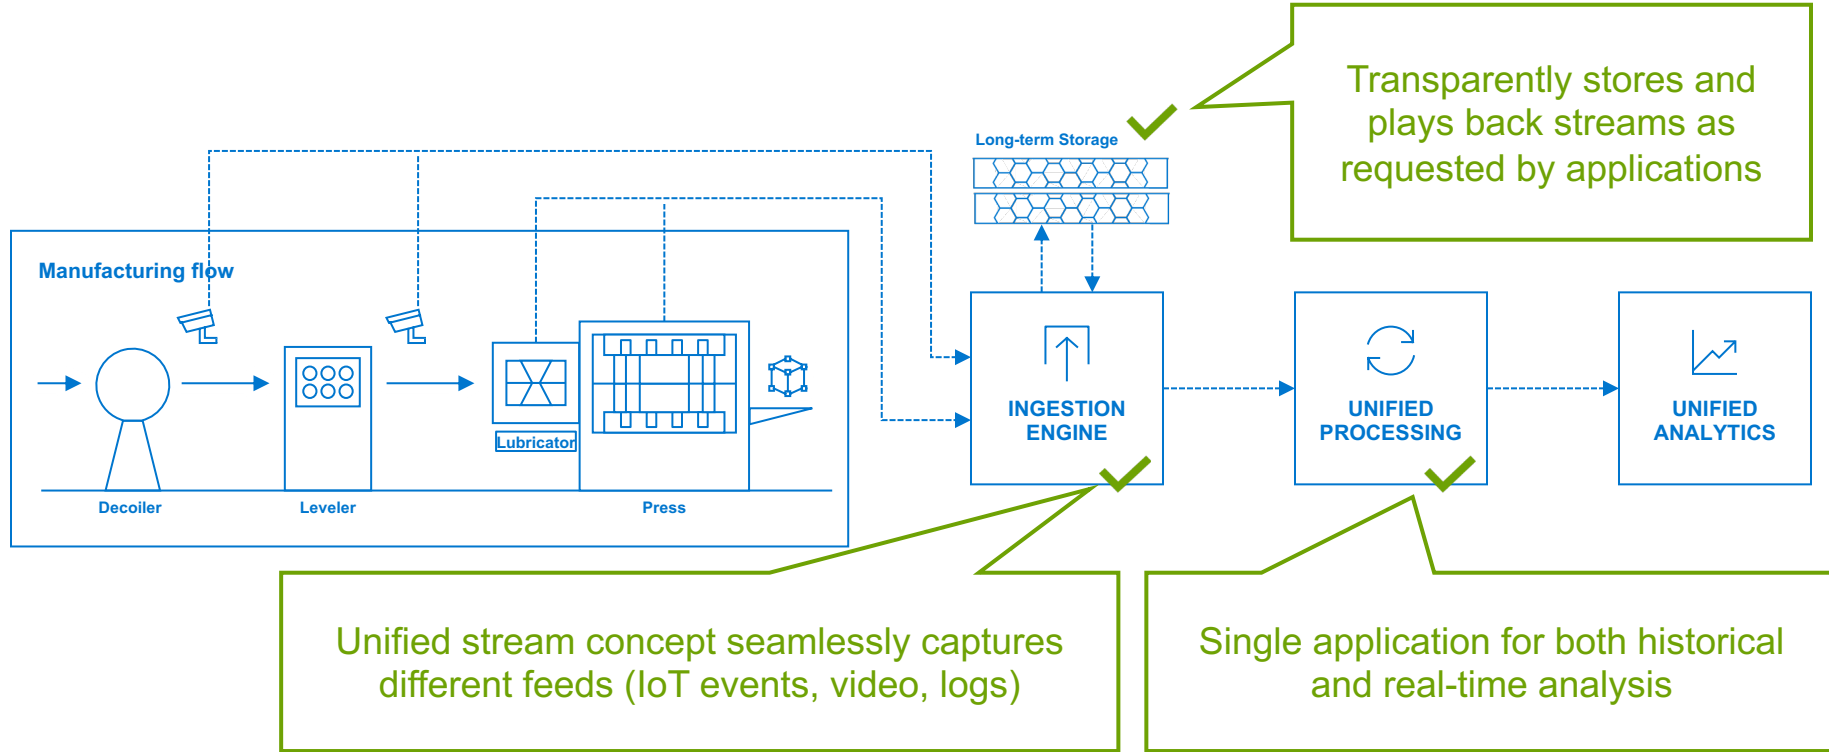
Use-Cases

Streaming Data Platform Uses Cases

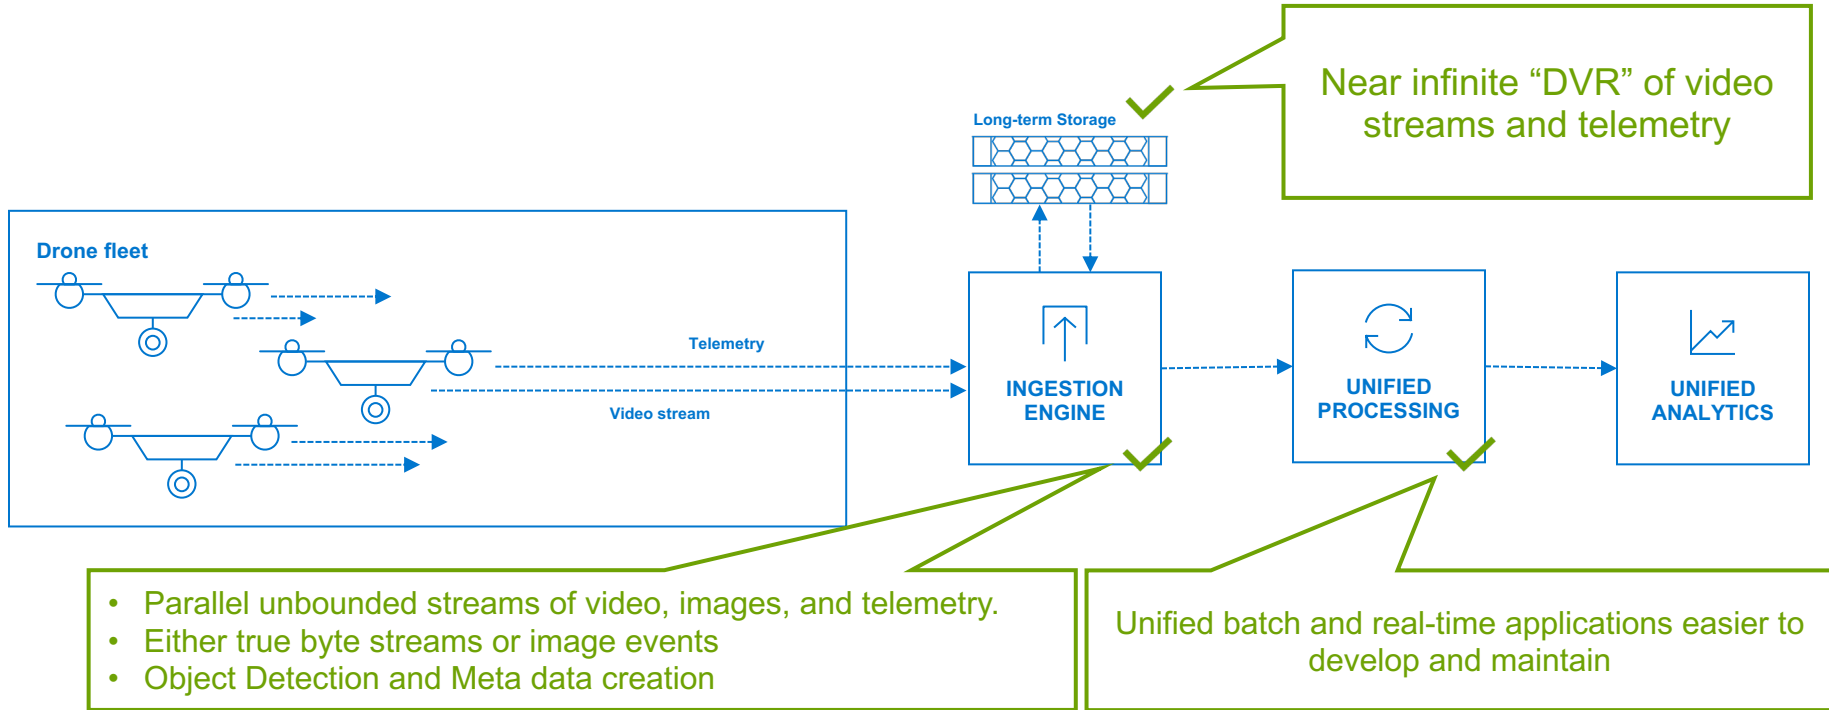




Industrial Sensor Data Anomaly Detection

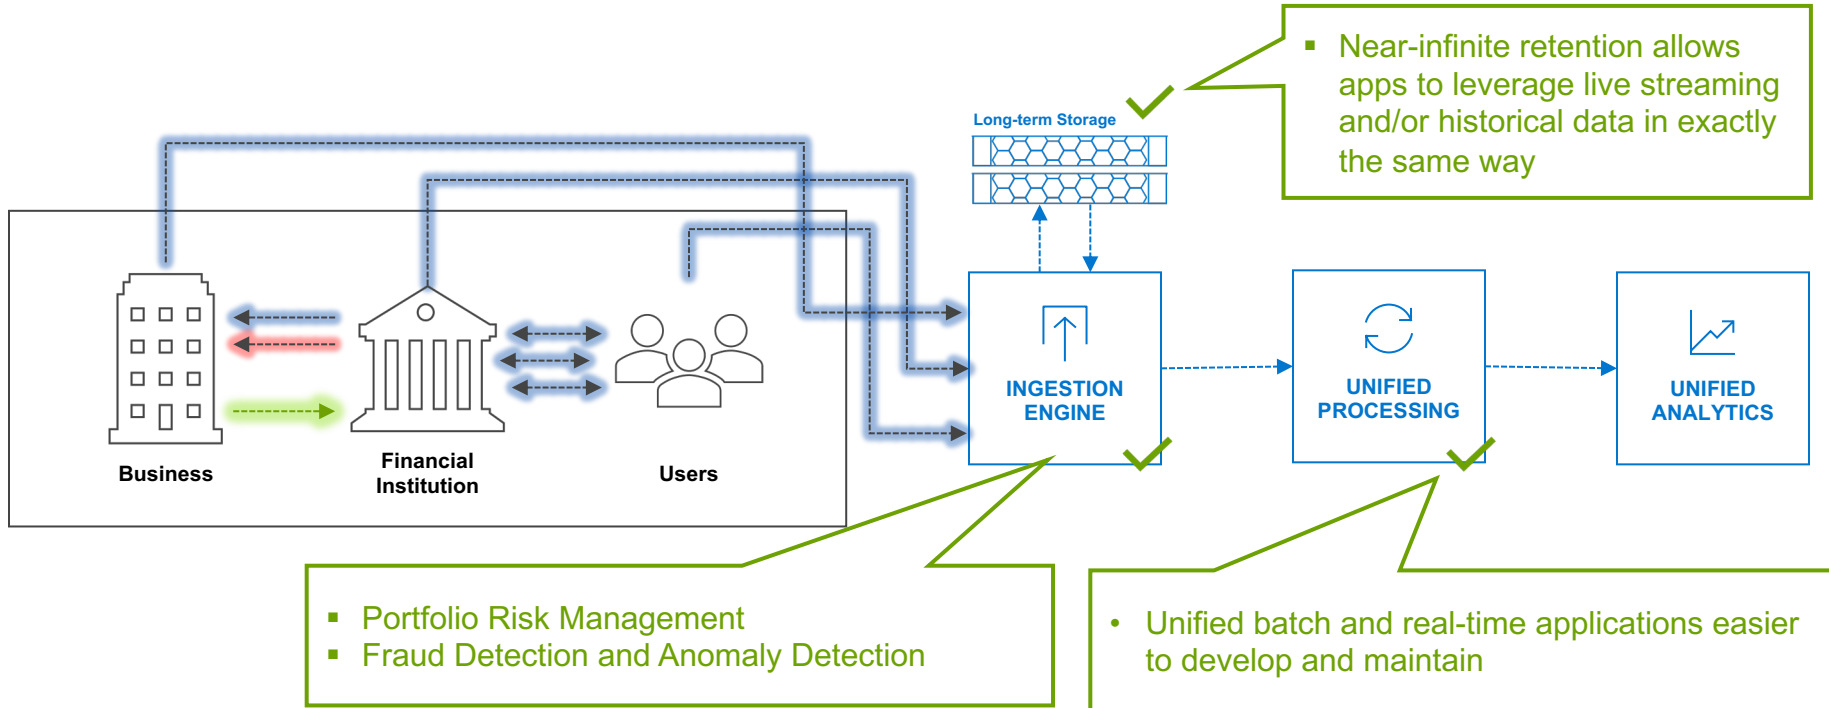


Unbounded Data Streams

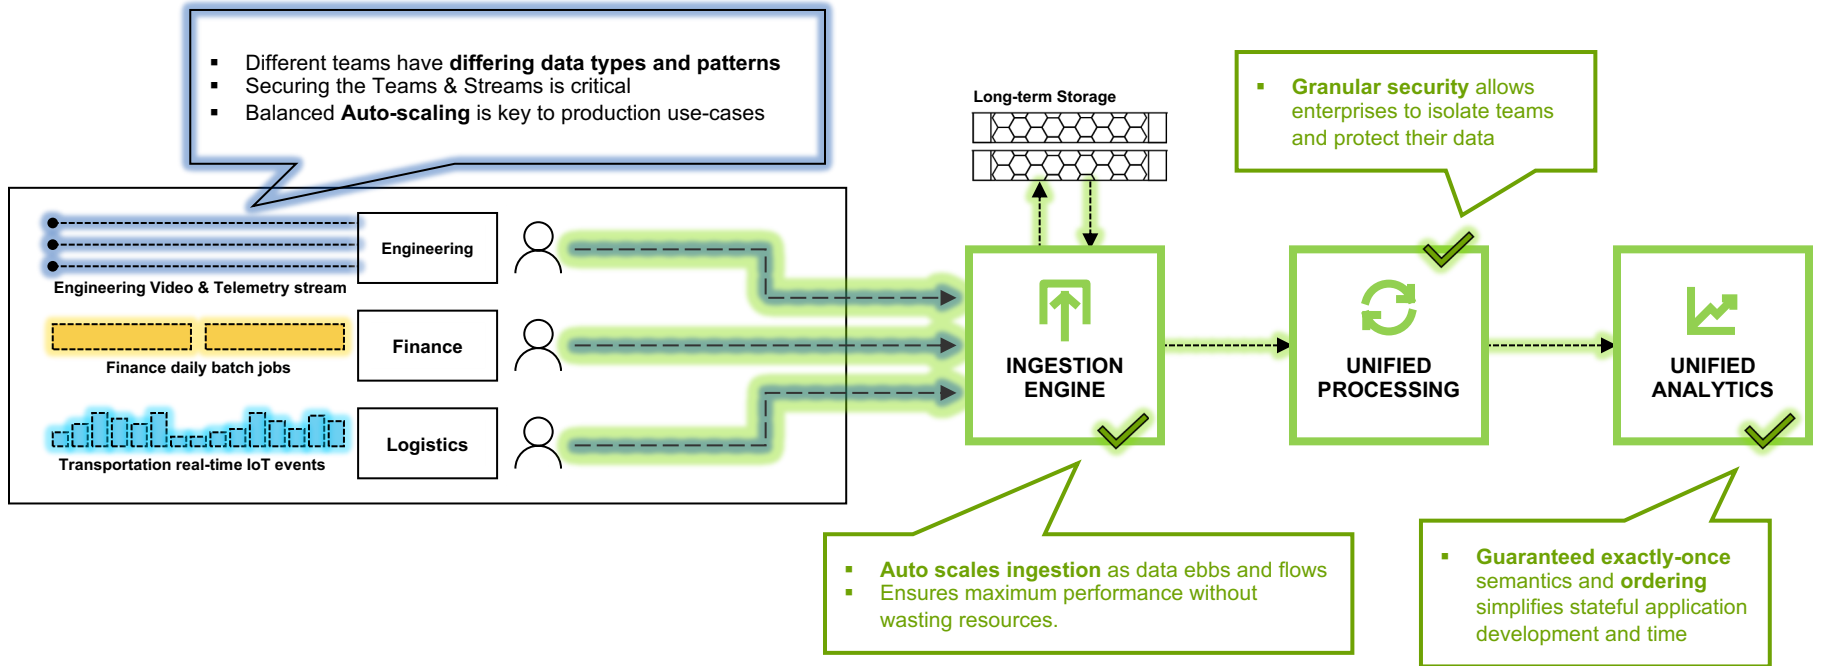




Exactly once, high consistency transactions

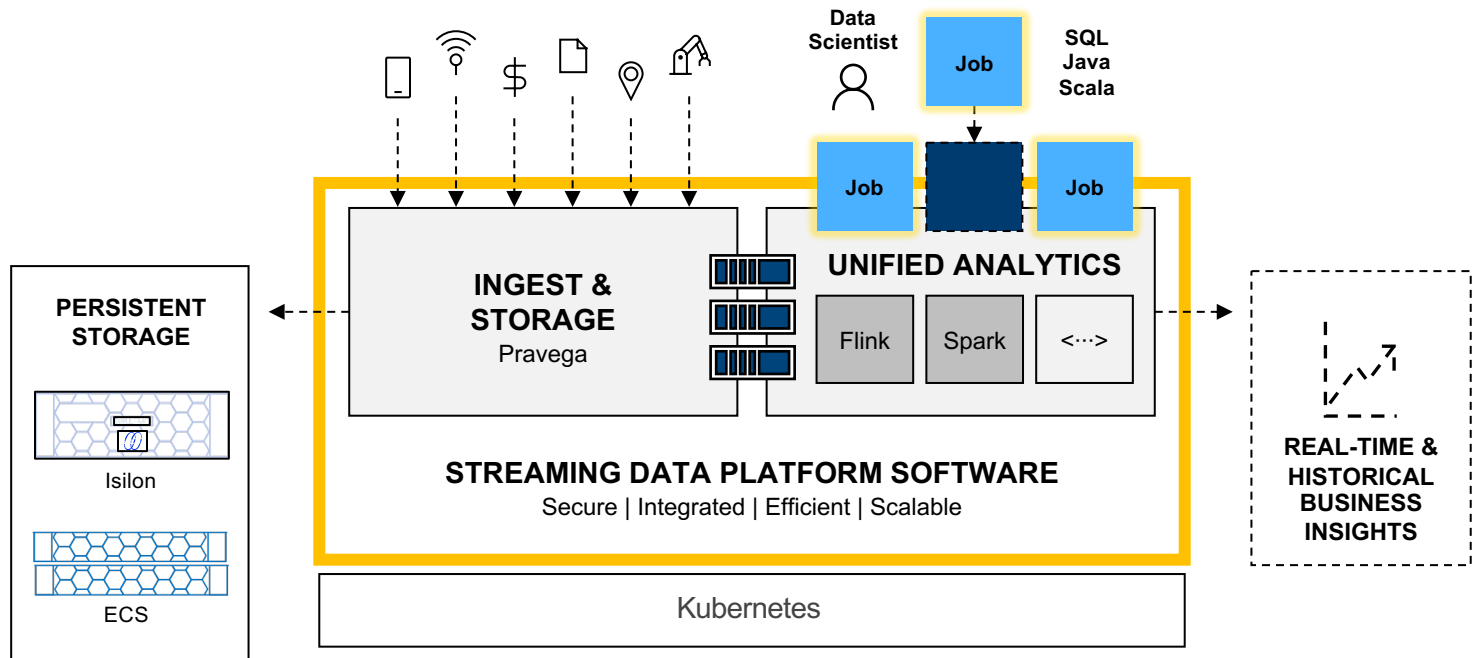


Multiple Teams & Use-Cases



Streaming Data Platform

Architectural deep dive



DELLTechnologies