

Streaming Analytics Package

Build reliable, real-time analytics at scale with the Unified Data Analytics Platform

Summary

Today's fast-paced, always-connected world is making streaming data a norm for core business operations. Use cases — fraud detection, ad optimization and predictive maintenance — require processing and analyzing ever-growing volumes of streaming data. This data arrives with high velocity and also requires monitoring the landing zone to automatically kick-start processing. Our offering accelerates building your streaming analytics pipeline leveraging Structured Streaming and Delta with our expert guidance so you can make the most timely decisions with your data.

Key outcomes

- Build modern multi-hop streaming pipeline reference architecture for co-selected scenarios
- Build a reference implementation for scoped streaming pipelines of your choice
- Build operational visibility tracking metrics, KPIs, trends and alert the right stakeholders
- Extended options for additional pipeline scenarios plus performance optimizations



Strategy

Build and optimize a Streaming Analytics workflow following the most up-to-date best practices, guided by experts. By leveraging structured streaming, Delta Lake and the associated architecture patterns, these packaged offerings address many of the common challenges faced when dealing with the complexities of delivering low-latency insights. The packages offer three tiers: **Foundation**, **Extended** and **Optimized**. Milestones and the outcome for each tier are produced by our prescriptive methodology, and each tier can be chained for greater impact on your streaming analytics exploration and adoption. See the **Resources and schedule** section for details.

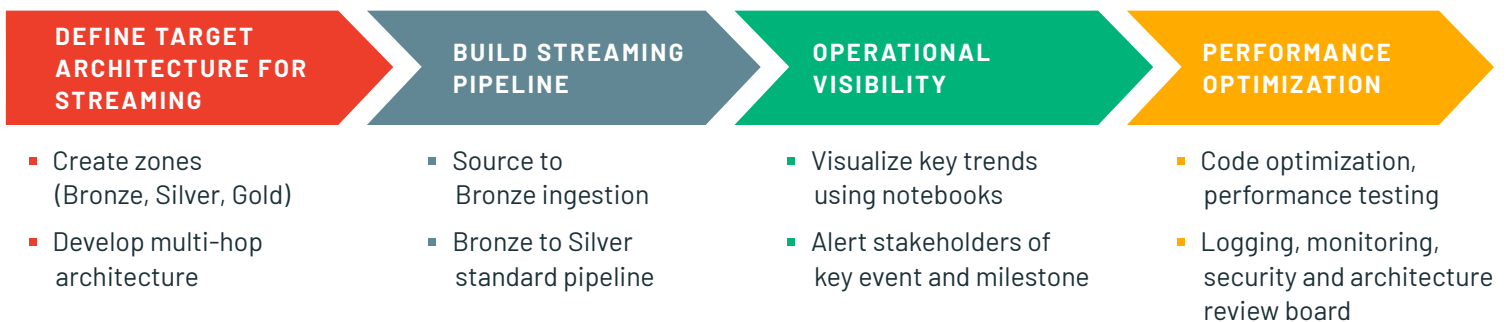
Challenges of building high-volume streaming data pipelines

- Writing data at low latency without corruption
- Managing multiple readers/writers
- Velocity vs. quality trade-offs
- Complex Lambda architecture
- Managing advanced analytics and ML on streaming data sets in real time

Key benefits

- Unified batch and streaming simplifies the data engineer's job
- Process streaming data at scale with enterprise-grade reliability
- Faster insights from your streaming data
- Secure a quick win and scale your streaming use cases

Databricks streaming analytics build process



Resources and schedule

FOUNDATION 2 WEEKS, \$42K

- Reference architecture
- Reference implementation of one streaming pipeline

EXTENDED 2 WEEKS, \$46K

- Implement additional pipeline
- Operational visibility

OPTIMIZED 2 WEEKS, \$46K

- Performance optimization
- Workflow automation
- Logging, monitoring, security

Up to 4 resources supporting the activity over a 2-week sprint
Prior to kickoff, be sure to review the readiness checklist and complete required tasks

Out of scope

- Configuration and integration of non-Databricks products and systems
- Data cleansing and solving data quality issues