



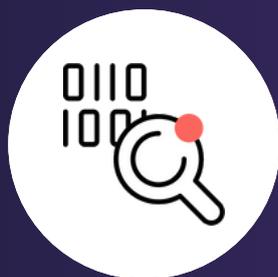
Comprehensive Databricks Health Check

Optimise Performance, Reduce Costs and Prepare for Serverless

In their 2024 Data & AI Summit, Databricks announced the general availability of Serverless Compute for Notebooks, Workflows and Delta Live Tables for Unity Catalog enabled workspaces. Serverless compute offers speed, simplicity, and reliability as well as the potential for cost optimisation. Coeo's Databricks Health Check assesses your organisation's readiness to transition to serverless.

Our evaluation covers your current setup, the roadmap to serverless, and the immediate benefits of switching. The Health Check also offers a full workspace analysis to ensure your Databricks setup supports your long-term data objectives.

Coeo's Databricks Health Check includes a comprehensive review covering 3 key areas:



Full Workspace Analysis

A detailed review of recent compute usage, performance metrics, security settings, and Unity Catalog readiness.



Cost Optimisation

Gain a clear view of current costs versus potential savings, including implementation considerations.



Change Impact Assessment

Understand the effects on your codebase, managed resource groups, and reserved instances, to meet your long-term data objectives.

The Health Check ensures that your Databricks environment is optimised for performance, reliability, and security, with the serverless upgrade in mind. We'll provide a full report and a 2-hour workshop discussing potential areas for optimisation.



Databricks Serverless Benefits & Considerations

Benefits

- **Speed:** Instant compute via Databricks pools means no more waiting for clusters.
- **Simplicity:** Automatic scaling to workloads and spark versioning means less complexity in choosing cluster configs.
- **Reliability:** Automatic instance type failover and a 'warm pool' of instances offer shields from availability shortages and cloud outages.
- **Elastic Costs:** Elastic billing modal and autoscaling to required capacity means no costs for idle time or over-provisioning.

Considerations

- **Unity Catalog:** Serverless requires Unity Catalog (UC) which may require upgrades if not currently enabled. UC has many benefits of its own.
- **Costs:** Higher DBU-Cost per hour and potential for future networking charges - Databricks have expressed that lower serverless costs are a future goal.
- **Impact on Codebase:** Some codebase will require changes given a few limitations of serverless.

Future-Proof Your Data Platform, We'll help you:



Improve your Data Culture

Fully optimise your Databricks workspace to enhance your ability to make data-driven business decisions



Optimise your Data Platform Costs

Streamline platform costs while maintaining data processing efficiency and performance



Ensure Reliable & Secure Data

Evaluate your data governance strategy, including Unity Catalog implementation readiness



Evaluate Findings and Next Steps

Align your current Databricks workspace capabilities with your organisation's long-term data objectives

Improve your data strategy and get more out of your Databricks investment. Contact our team today to schedule your three-step Databricks Health Check.

Step 1

Databricks Usage Workshop

Step 2

Full Workspace Analysis

Step 3

Findings Report & Strategy Workshop