

# Data Cosmos\_BricksOps

*Databricks Operations, Governance & Optimization*

Part of the Coforge Data Cosmos™ Platform | Pulsar™ Portfolio

**Coforge Limited | Data HBU**

2026

# Solution Overview

Data Cosmos\_BricksOps is a Coforge Data Cosmos accelerator (under the Pulsar suite) that transforms Databricks operations, governance analysis, and platform observability into a unified and scalable workflow. It serves as a single control plane that enables platform, engineering, and governance teams to collect workspace inventory, correlate source code and CI/CD activity with Databricks job execution timelines, detect governance signals, and generate cluster optimization recommendations. Part of Coforge Data Cosmos™, the innovation backbone combining platforms, agentic accelerators, and services, it consolidates the operations intelligence lifecycle into one governed platform.

Many organizations manage Databricks operations through disconnected workflows and manual processes — leading to time-consuming metadata extraction, slow root-cause analysis, reactive cluster optimization, and high effort for governance reviews and audits. BricksOps eliminates this burden by consolidating inventory collection, operational timelines, governance analysis, and optimization planning into a single governed workflow. Whether answering “what changed, who changed it, and which jobs were impacted” for an audit, investigating an incident, or right-sizing clusters for FinOps, BricksOps compresses hours of manual effort into minutes with automated scans, repository-to-job correlation, and structured cluster advisories.

## Key Features

| Feature                                      | Description                                                                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Automated Databricks Inventory Collection    | Section-based workspace inventory collection with targeted or full scans, asynchronous execution, status tracking, retry-resilient workflows, and snapshot-based storage for faster retrieval and reporting.                                   |
| Repository, CI/CD & Job Timeline Correlation | Source code commit tracking, pull request and workflow monitoring, CI/CD execution analysis, workflow log parsing, Terraform deployment signal extraction, and Databricks job lifecycle correlation with confidence and provenance indicators. |
| Governance and Audit Analysis                | Unified job timeline views, change traceability, PR and deployment context visualization, runtime behavior correlation, and audit-ready evidence generation — answering what changed, who changed it, and which jobs were impacted.            |
| Cluster Advisory & Optimization Planning     | Cluster configuration analysis, recommendation generation, usage pattern evaluation, optimization scenario creation, adoption planning, and rollback-aware operational planning.                                                               |
| Multi-View Operations Portal                 | Executive dashboards, KPI and health summaries, inventory exploration, job analysis, run timeline analysis, governance traceability, cluster recommendation portals, and comparative analysis pages.                                           |
| Internal Roles Architecture                  | Four specialized roles — Collector (inventory metadata), Correlator (repo/CI-CD to job mapping), Advisor (optimization recommendations), and Planner (adoption and rollout strategies) — work together across the operations lifecycle.        |

# Four-Role Operations Pipeline

| Role       | Focus                        | What It Does                                                                                                                                                                              |
|------------|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Collector  | Inventory Metadata           | Maintains Databricks workspace inventory metadata through section-based scans with asynchronous execution, status tracking, and snapshot-based storage for fast retrieval and reporting.  |
| Correlator | Repo/CI-CD to Job Mapping    | Ingests repository and CI/CD evidence — commits, PRs, workflows, Terraform signals — and maps activities to Databricks job execution timelines with confidence and provenance indicators. |
| Advisor    | Optimization Recommendations | Builds cluster optimization recommendations and improvement scenarios by analyzing configurations and usage patterns for cost efficiency and performance.                                 |
| Planner    | Adoption & Rollout Strategy  | Generates operational adoption plans and rollout strategies, including rollback-aware planning for safe implementation of optimization recommendations.                                   |

## How BricksOps Works

1. **Connect** — Connect to Databricks workspaces and repository/CI-CD systems across the platform estate
2. **Collect** — Run section-based or full workspace inventory scans with async execution and snapshots
3. **Track** — Track source code commits, pull requests, workflows, and Terraform deployment signals
4. **Correlate** — Map repository and CI/CD activity to Databricks job execution timelines with provenance
5. **Analyze Governance** — Build unified job timelines and change traceability for audit-ready evidence
6. **Investigate** — Answer what changed, who changed it, which deployment, and which jobs were impacted
7. **Advise** — Analyze cluster configurations and usage to generate optimization recommendations
8. **Plan** — Create adoption plans and rollback-aware rollout strategies for optimizations
9. **Visualize** — Explore executive dashboards, KPIs, inventory, timelines, and cluster portals
10. **Optimize** — Adopt recommendations to reduce cluster costs and improve FinOps outcomes

# Industry Trends and where it can be leveraged

## Banking & Financial Services

| Use Case                          | Example                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Audit-Ready Governance            | A global bank uses BricksOps to correlate repository and CI/CD activity with Databricks job timelines across risk and regulatory workloads. Audit-ready evidence generation instantly answers what changed and which jobs were impacted — compressing multi-day audit preparation into immediate dashboard visibility for OCC and Fed reviews. |
| Incident Root-Cause Investigation | A retail bank uses timeline correlation to trace a data incident back to the specific deployment and pull request that introduced it, reducing multi-hour investigations across disparate tools to minutes.                                                                                                                                    |

## Insurance

| Use Case                        | Example                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Post-Release Validation         | A multi-line insurer uses BricksOps for post-release validation across claims and actuarial workloads, correlating deployments with job runtime behavior to confirm releases performed as expected and maintain change traceability for Solvency II evidence. |
| Governance Reporting Automation | An insurer replaces manual governance reporting with BricksOps dashboards, giving governance teams immediate visibility into change history and job impact across policy admin and underwriting pipelines.                                                    |

## Travel, Transportation & Hospitality

| Use Case                           | Example                                                                                                                                                                                                                                      |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FinOps & Cluster Cost Optimization | A global airline uses BricksOps' Advisor and Planner roles to right-size Databricks clusters for peak-season booking and revenue workloads, generating structured recommendations and rollback-aware adoption plans that reduce cloud spend. |
| Operations Observability at Scale  | A hotel chain consolidates Databricks operations visibility across guest and loyalty workloads into a single portal, replacing fragmented tools with unified inventory and timeline analysis.                                                |

## Healthcare & Life Sciences

| Use Case                      | Example                                                                                                                                                                                                                                      |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HIPAA Governance Traceability | A healthcare provider uses BricksOps to maintain governance traceability and audit evidence across HIPAA-regulated Databricks pipelines, correlating changes to job execution for PHI-handling workloads and simplifying compliance reviews. |
| Research Pipeline Operations  | A pharma company uses BricksOps to track repository-to-job correlation across clinical research pipelines, accelerating knowledge transfer and post-release validation for trial analytics.                                                  |

## Cross-Industry Use Cases

| Use Case                           | Description                                                                                                                                       |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Databricks Operations Intelligence | Consolidate inventory collection, operational timelines, and observability into a single governed control plane across all Databricks workspaces. |
| Compliance & Audit Preparation     | Generate audit-ready evidence with change traceability and timeline correlation, turning multi-day prep into immediate dashboard visibility.      |
| Incident Investigation             | Trace incidents to the exact commit, PR, or deployment that introduced them, cutting root-cause analysis from hours to minutes.                   |

| Use Case                     | Description                                                                                                          |
|------------------------------|----------------------------------------------------------------------------------------------------------------------|
| FinOps & Cost Optimization   | Generate structured cluster recommendations and rollback-aware adoption plans to reduce Databricks compute spend.    |
| Post-Release Validation & KT | Validate releases against job runtime behavior and accelerate knowledge transfer with unified operational timelines. |

## Target Audience

- Databricks Platform & Operations Teams
- Data Engineering & Platform Teams
- FinOps & Cost Optimization Teams
- Governance, Risk & Audit Teams
- Data Architects and Cloud Migration Leads

## Azure Compatibility

Data Cosmos\_BricksOps integrates natively with Azure Databricks for workspace connectivity, inventory collection, and job timeline correlation. It is deployed on Azure Kubernetes Service (AKS) with a containerized backend (Python, Flask) and frontend (React, Vite), uses Azure Data Lake Storage / Azure Blob Storage for artifact storage, and integrates with Azure DevOps Pipelines and GitHub Actions for CI/CD correlation. Azure Key Vault provides secure secrets management and Azure Monitor delivers observability. It is part of the Coforge Data Cosmos Pulsar portfolio for data-ops solution on Databricks, fully designed to leverage Azure-native services for scalable, secure, and governed enterprise Databricks operations, governance, and optimization.