

Data Cosmos_PULSAR

Agentic Production Support

Part of the Coforge Data Cosmos Platform | Pulsar
Portfolio

Coforge Limited | Data HBU

2026

Solution Overview

Data Cosmos_PULSAR is an enterprise-grade agentic production support system designed to autonomously monitor, triage, and resolve issues in Data and AI batch processing environments. It replaces reactive human-centric L1/L2 support models with an intelligent, always-on multi-agent system augmented by human-in-the-loop oversight. PULSAR is a core component of the Data Cosmos Platform, built by Coforge Limited to transform how production support operates for enterprise data programs, dramatically reducing Mean Time to Resolution (MTTR), improving SLA compliance, and providing complete observability into production support operations.

Enterprise data programs today run hundreds to thousands of interdependent batch jobs nightly across multiple orchestration platforms — Airflow, Control-M, Autosys, Azure Data Factory. When a critical job fails at 2 AM, the current model relies on an L1 operator to wake up, find the right runbook (if it exists), diagnose the issue, and either fix it or escalate. This reactive model is expensive, error-prone, and increasingly unsustainable as data estates grow in complexity. PULSAR fundamentally reimagines this model by deploying intelligent AI agents that never sleep, never forget a resolution pattern, and continuously learn from every incident — turning production support from a cost center into a compounding knowledge asset.

The Problem

Challenge	Impact
Reactive Support Model	Human operators discover failures after significant delays, leading to missed SLA windows, downstream cascading failures, and business-critical reporting delays.
Knowledge Silos	Tribal knowledge resides with 5–10 senior engineers. Runbooks are outdated or incomplete. New support staff face steep learning curves. Losing any key person breaks operations.
Escalation Inefficiency	L1 operators frequently escalate tickets they could resolve with proper knowledge access, overwhelming L2/L3 teams with repetitive issues and inflating support costs.
Cost Pressure	24/7 support staffing for nightly batch windows is expensive (8–12% annual cost increase). Off-hours incidents suffer from slower response times due to reduced staffing.
Limited Observability	Fragmented tooling across 3–5 orchestration engines and vendors provides no unified view of batch health, SLA compliance, or support performance metrics.
No Trend Intelligence	Historical incident data is trapped in ticketing systems with no automated trend analysis to proactively identify recurring issues or degrading job health patterns.

The PULSAR Agent Cohort

Agent	Persona	Tier	Capabilities
Mr. Glass	The Watchman	L1	Continuous 24/7 monitoring of ETL/ELT batch jobs, data mart refreshes, reporting jobs, AI model training pipelines, DQ check jobs, and file transfer monitors. Creates ITIL-compliant tickets with auto-severity classification (P1–P4). Consults runbooks and KEDB before attempting L1 resolution. Executes job resubmission for transient failures, clears stale locks, handles file arrival waits, and refreshes expired credentials.
Mr. QuickFix	The Fixer	L1.5	Pattern-based incident resolution using accumulated knowledge and learned patterns from the KEDB. Resolution strategies include: data pipeline parameter adjustments, service restarts, connection pool resets, memory/resource allocation corrections, intermediate data cleanup, and retry with modified configurations. Makes up to 2 attempts before escalating to L2.
Ms. Troubleshoot	The Investigator	L2	Deep diagnostic analysis for complex, multi-system failures. Capabilities include: database query analysis and optimization, system configuration adjustments, data integrity verification, pipeline code-level diagnostics, infrastructure-level interventions, cross-

			system dependency resolution, and environment variable/connection string corrections. Makes up to 2 attempts before escalating to L3 human.
--	--	--	---

Agents collaborate via the A2A (Agent-to-Agent) protocol. Skills are delivered via MCP (Model Context Protocol). Memory persists across incidents. Knowledge compounds with every resolution. Human-in-the-Loop L3 support is available for complex escalations following ITIL incident management protocols.

Key Platform Capabilities

Capability	Description
Batch Schedule Monitoring	Real-time monitoring across Airflow, Control-M, Autosys, ADF with <30 second failure detection, cascading failure risk identification, and proactive SLA breach alerting.
Integrated Ticket Management	Built-in JIRA-like ticket system with full lifecycle (Open → In Progress → Resolved → Closed), P1–P4 severity levels, and configurable SLA management.
Known Error Database (KEDB)	Auto-cataloging of recurring issues, root causes, and proven resolutions. Agents query KEDB as part of every diagnostic workflow, improving resolution accuracy over time.
Knowledge Library	Searchable repository of BRDs, TRDs, runbooks, architecture diagrams, ER diagrams, and SOPs with version control and integration with Confluence/SharePoint.
Real-Time Dashboards	Job run view, SLA compliance, agent activity, Kanban ticket board, system health — all with ITIL KPIs (MTTD, MTTR, FCR, SLA %).
ITIL-Compliant Reporting	Daily ops summary, weekly performance, monthly SLA, incident analysis, agent performance, capacity planning, and executive summary reports.
Trend Analysis & Predictive Intelligence	Anomaly detection for job runtime patterns, predictive alerting for likely failures, and seasonal pattern recognition (month-end spikes, quarterly cycles).
Multi-Channel Communication	Email, Slack, Telegram, WhatsApp, and Portal Chat for bidirectional agent–human interaction and critical P1 escalation notifications.
Knowledge Graph Visualization	Interactive graph of jobs, schedules, tables, systems, files, and their relationships backed by Neo4j for impact analysis.
Dynamic Agent Provisioning	On-demand agent creation and scaling based on workload metrics with configurable auto-scaling policies on Kubernetes.

Industry Use Cases where it can be leveraged

Banking & Financial Services

Use Case	Example
Agentic L2 Data Support for Retail & Corporate Banking	A Tier-1 bank runs 3,000+ nightly batch jobs across Control-M and Airflow for core banking ETL, risk calculations, and regulatory reporting. PULSAR's Mr. Glass monitors the entire batch window, auto-triaging failures across credit risk, market risk, and liquidity data pipelines. When a BCBS 239 risk aggregation job fails due to a stale database connection at 3 AM, Mr. QuickFix identifies the pattern from KEDB (seen 47 times in the past quarter), resets the connection pool, resubmits the job, and validates completion — all within 4 minutes, zero human involvement. Ms. Troubleshoot handles complex cascading failures where a single upstream data quality issue in the core banking feed impacts 15+ downstream regulatory reports.
Payment Data Support & Reprocessing	A payments processing unit handles 2M+ daily transactions with batch settlement jobs running on Autosys. PULSAR detects settlement file arrival delays, auto-waits with intelligent timeout logic, and resubmits processing jobs when files arrive. For reconciliation mismatches, Ms. Troubleshoot performs cross-system diagnostic analysis between payment gateways and core ledger systems.

Insurance

Use Case	Example
Agentic L1/L2 Support for Claims Systems	An insurance company runs nightly claims processing batches across Guidewire and Duck Creek with data flowing into a Snowflake analytics warehouse. PULSAR monitors 800+ jobs including claims adjudication, reserve calculations, subrogation scoring, and policyholder notification pipelines. When a claims data mart refresh fails due to a schema change in Guidewire, Mr. Glass detects the failure in <30 seconds, Mr. QuickFix identifies it as a known schema drift pattern, applies the DDL reconciliation fix from KEDB, and restores the pipeline — all before the actuarial team's morning SLA window. Over 6 months, PULSAR resolves 78% of claims pipeline incidents autonomously, reducing L1/L2 support costs by 55%.
Automated Claims Workflow Optimization	PULSAR's trend analysis identifies that claims batch jobs consistently degrade in performance during month-end processing due to volume spikes. Predictive intelligence proactively scales Airflow workers before the spike, preventing SLA breaches that previously occurred 3–4 times per quarter.

Travel, Transportation & Hospitality

Use Case	Example
Agentic L1/L2 Support for Booking Platforms	A global airline runs 1,500+ daily data pipeline jobs feeding booking analytics, revenue management, loyalty programs, and disruption management systems. PULSAR monitors Airflow DAGs processing Amadeus/Sabre GDS feeds and

	detects a booking data pipeline failure during peak summer travel season. Mr. Glass triages the issue — an API rate limit from the GDS provider. Mr. QuickFix applies a learned backoff-and-retry pattern, staggers API calls, and completes ingestion within the SLA window. Ms. Troubleshoot handles complex scenarios like flight disruption cascades requiring emergency priority resequencing of rebooking data pipelines.
Intelligent Booking System Health Checks	PULSAR agents perform proactive health checks on booking platform integrations during off-peak hours, detecting connection degradation to hotel PMS systems (Opera) before they impact nightly revenue consolidation jobs. This prevented 12 potential SLA breaches in the first quarter of deployment.

Healthcare & Life Sciences

Use Case	Example
Agentic L1/L2 Support for Clinical Systems	A healthcare provider runs nightly ETL jobs migrating EMR/EHR data (Epic, Cerner) into a clinical data warehouse for population health analytics and care quality reporting. PULSAR monitors 600+ jobs including patient data ingestion, lab result processing, claims reconciliation, and CMS quality measure calculations. When an HL7 FHIR interface feed fails due to a certificate expiration, Mr. Glass detects it within 30 seconds, creates a P1 ticket, and Mr. QuickFix applies the certificate refresh procedure from KEDB. For HIPAA-sensitive environments, all agent actions maintain complete audit trails with PHI-aware logging.
Clinical Data Pipeline Observability	PULSAR provides a single-pane-of-glass dashboard for clinical data operations, correlating pipeline health across Epic Clarity extracts, lab system feeds, and pharmacy data — enabling the data operations team to proactively identify degrading ETL performance before it impacts clinical decision support systems.

Public Sector

Use Case	Example
Agentic Support for Government Data Platforms	A government agency manages citizen service data pipelines processing benefits, tax, and social welfare data across multiple legacy systems migrated to Azure. PULSAR monitors batch jobs with strict SLA windows tied to regulatory reporting deadlines. When a benefits calculation pipeline fails due to a data format change from a partner agency feed, Mr. Glass auto-triages, Mr. QuickFix applies the format adaptation fix, and the pipeline completes before the statutory reporting deadline — preventing compliance penalties and citizen service disruptions.

Quantified Business Outcomes

Metric	Target
5-Year TCO Reduction	50–60%

MTTR Reduction	40–60%
Knowledge Search Latency	<30 seconds (down from 15–25 minutes)
Autonomous Resolution Rate	85%+
SLA Compliance	99.5%
Production Support Cost Reduction	60%

Engagement Model

Phase	Duration	Description	Outcome
Phase 1 Knowledge Onboarding	8–12 weeks	Ingest all operational knowledge into PULSAR KMS — runbooks, system docs, BRDs/TRDs, Confluence, past tickets, schemas. Searchable, version-controlled, agent-ready.	KMS deployed, team 80% faster at finding knowledge
Phase 2 Agent Activation	4–8 weeks	Deploy PULSAR platform, train and stand up the three-agent cohort, connect orchestration engines, run agents in parallel with team for validation.	Live agents, validated performance, signed transition plan
Phase 3 Autonomous Operations	Multi-year	Agent cohort handles 70–85% of L1/L1.5 autonomously. Senior Coforge L3 humans handle complex cases. Every resolution feeds back into the KMS.	Production stability, falling cost curve, growing knowledge asset

Technology Stack

Layer	Technology
Frontend	React + TypeScript + Tailwind CSS + shadcn/ui
Backend API	FastAPI (Python)
Backend Services	Node.js
Database (Relational)	PostgreSQL
Database (Document)	MongoDB
Cache / Pub-Sub	Redis

Agent Framework	LangChain / LangGraph
Agent Communication	A2A Protocol
Skill Provisioning	MCP (Model Context Protocol)
Graph Database	Neo4j
Containerization	Docker + Kubernetes + Helm
Search Engine	Elasticsearch
Monitoring	Prometheus + Grafana

Supported Integrations

- Apache Airflow (REST API / Webhooks)
- Control-M (Automation API)
- CA Autosys (CLI / REST API)
- Azure Data Factory (Azure REST API)
- Slack (Slack API / Webhooks)
- Telegram (Bot API)
- WhatsApp (Business API)
- Confluence (REST API)
- SharePoint (Graph API)
- LDAP / Active Directory
- Prometheus / Grafana (Metrics API)
- External Agents (A2A Protocol)

Azure Compatibility

PULSAR integrates natively with Azure Data Factory, Azure Databricks, Azure Monitor, Azure Kubernetes Service, and Azure DevOps. It is deployable on Azure Kubernetes Service (AKS) for elastic scalability with Helm chart-based deployment. PULSAR is part of the Coforge Data Cosmos Pulsar portfolio, fully designed to operate alongside Azure-native monitoring, orchestration, and data services for enterprise production support.