

Whitepaper

Driving Innovation with AI-Driven Analytics and Cloud-Native Platform

APRIL 2025

Table of Contents

Rakuten Cloud-Native Platform coupled with Rakuten AI for Telecom.....	02
Analytics-as-a-Service in Orange Africa Case Study.....	03
Rakuten Cloud-Native Platform.....	04
Rearchitcting AI Cloud Platforms: A Data-First Approach.....	06
Network Monitoring Case Study for a Tier-1 Operator.....	08
Conclusion.....	10

AI and analytics are becoming critical components for how enterprises and telecom operators make decisions and automate operations. To build this capability, enterprises, system integrators and telecom operators worldwide are adopting intelligent platforms and automation technologies as they prepare to deploy the next generation of cloud-native analytics applications. This shift is driving technology investments and innovation in cloud-native platforms that are purpose-built for ingesting, processing, analyzing and presenting data generated by the organization's systems, processes and infrastructure. At Rakuten Symphony, we offer a complete end-to-end solution for your analytics initiatives, enabling you to transform not just your AI and analytics platform, but your entire business.

Rakuten Cloud-Native Platform coupled with Rakuten AI for Telecom

AI and analytics are revolutionizing business through data driven decisions and intelligent automation. In September 2012 (“ImageNet Classification with Deep Convolutional Neural Networks”), Alex Krizhevsky demonstrated how, using two \$500 GPUs, he could classify images into roughly 1,000 categories in 10–20 milliseconds. His groundbreaking speed and accuracy revolutionized AI, transforming gaming technology into an economic powerhouse. This innovation ushered in a new age of AI that is influencing every aspect of operations in data-driven enterprises worldwide.

Managing AI Platforms Remains a Challenge

Enterprises and telecom operators looking to use AI to drive automation and data-driven business operations face the challenge of deploying and managing a complex pipeline of applications. These often require expensive hardware, including GPUs, which must be provisioned efficiently and managed to provide data scientists with a simple, on-demand platform that is efficient, highly available, and extensible. Regulatory and data sovereignty issues, along with rising costs of public cloud alternatives, are driving enterprises toward the so-called “repatriation” of data workloads to private, on-premise clouds. These environments better support strictly secured and air-gapped deployments, in-country data residency, ease of customization, and cost control.

Rakuten Cloud-Native Platform (CNP), coupled with Rakuten AI for Telecom, provides automated provisioning and lifecycle management of sophisticated analytics pipelines on demand. This solution reduces platform complexity and simplifies the lifecycle management

of data-intensive applications at scale. It enables efficient data analysis, real-time event processing, and integration with existing data workflows, offering a comprehensive Analytics-as-a-Service solution. The solution provides data integration to unify data from disparate sources, creating a single source of truth for analysis. It also offers tools to enforce data governance policies, monitor usage, and track data lineage and processing—including machine learning, data visualization, and other analytics capabilities—to help organizations derive actionable insights.

AI Applications as a Service

Rakuten Symphony’s offering is designed to deliver advanced AI/ML workloads on demand, enabling telecom operators and enterprises to easily harness the power of AI for network optimization, predictive analytics, and enhanced customer experiences.

The offering provides a comprehensive environment for managing the entire data lifecycle, including data governance, querying, transformation, and discovery. It ensures secure interactions and efficient orchestration of data operations through seamless integration with cloud-native platform including storage, object storage as well as CPU and GPU resources.

The platform is designed to support advanced analytics and AI workloads, enabling organizations to leverage data-driven insights for improved decision-making. By integrating machine learning capabilities, the platform empowers teams to automate processes, detect anomalies, and forecast trends, ensuring operational efficiency and business continuity.




Figure 1 – Rakuten Symphony's offering for AI and analytics, powered by Rakuten Cloud-Native Platform

Analytics-as-a-Service in Orange Africa Case Study

Orange Africa was looking for a standardized solution for data analytics application delivery and lifecycle management that would simplify the complex tasks they faced as part of the deployment and operation of their data analytics applications, including Spark, Kubeflow, and TensorFlow.

The team was facing many challenges, including:

- Time and complexity of delivering platform to support complex analytics pipelines and applications.
- Complexity and delays in provisioning new applications and workloads.
- Managing the performance of the deployed applications and scaling to meet the dynamic demands of data analytics requirements.
- Day 2 operations complexity, including storage lifecycle management, data protection and availability.
- Ensuring business continuity and high availability of the deployed solution, as the availability of the workloads is critical to Orange Africa's business.

Delivering Solutions to Meet Customer Requirements

Orange Africa's requirements drove the need for an Application-as-a-Service platform that delivered an agile, scalable analytics solution, meeting the needs of data scientists as well as IT and operational teams.

Data scientists needed a solution that offered:

- Self-service, rapid workload provisioning.
- Simplified application management.
- Performance optimized to achieve the desired SLAs.
- Highly available platform that also meets data protection and security requirements.

Cloud and Information Technology team required a solution that:

- Provided a future-proof and standardized platform.
- Reduced maintenance needs and offered better support for maximized resource utilization.
- Delivered a cost-effective solution compared to using public cloud.

Key Benefits from Deployment of the Solution in Orange Africa

Orchestration of different apps

- Application-agnostic, supports a wide range of applications.
- Application-aware orchestration.
- Optimized network, storage, and workload placement.

Automation of application lifecycle management

- Simplified and automated deployment to reduce service delivery timeframes and increase resiliency.
- API-first, zero-touch automation and orchestration.

Performance at scale

- Scales to thousands of applications, hundreds of users, and hundreds of nodes per cluster.

- Advanced QoS, self-scaling, and auto-healing capabilities built into the platform.

Built-in data & storage management

- SDS (Software-Defined Storage) for simplified management.
- Snapshot, clone, and restore for streamlined data operations.
- Integrated backup and disaster recovery.

Self-service & multitenancy

- One-click provisioning of the entire application pipeline.
- Multitenancy to avoid noisy neighbor issues and maximize resource utilization.
- Enterprise-grade security with RBAC (Role-Based Access Control).

Rakuten Cloud-Native Platform

Rakuten Cloud-Native Platform is an industry-leading Kubernetes® platform that is optimized for running storage and network intensive applications. The platform can automate delivery of complex applications in minutes and deliver cloud-native agility and scalability for complex data-centric applications — perfect for data scientists setting up their own private cloud.

Rakuten Cloud-Native Platform provides an application workflow manager layer on top of Kubernetes®, which can onboard new applications via its one-click interface or API. The software can deploy complete application pipelines seamlessly. Once the software is onboarded, the Application Workflow Manager provides a suite of day 2 management tools including an advanced scheduler, application observability, and monitoring.

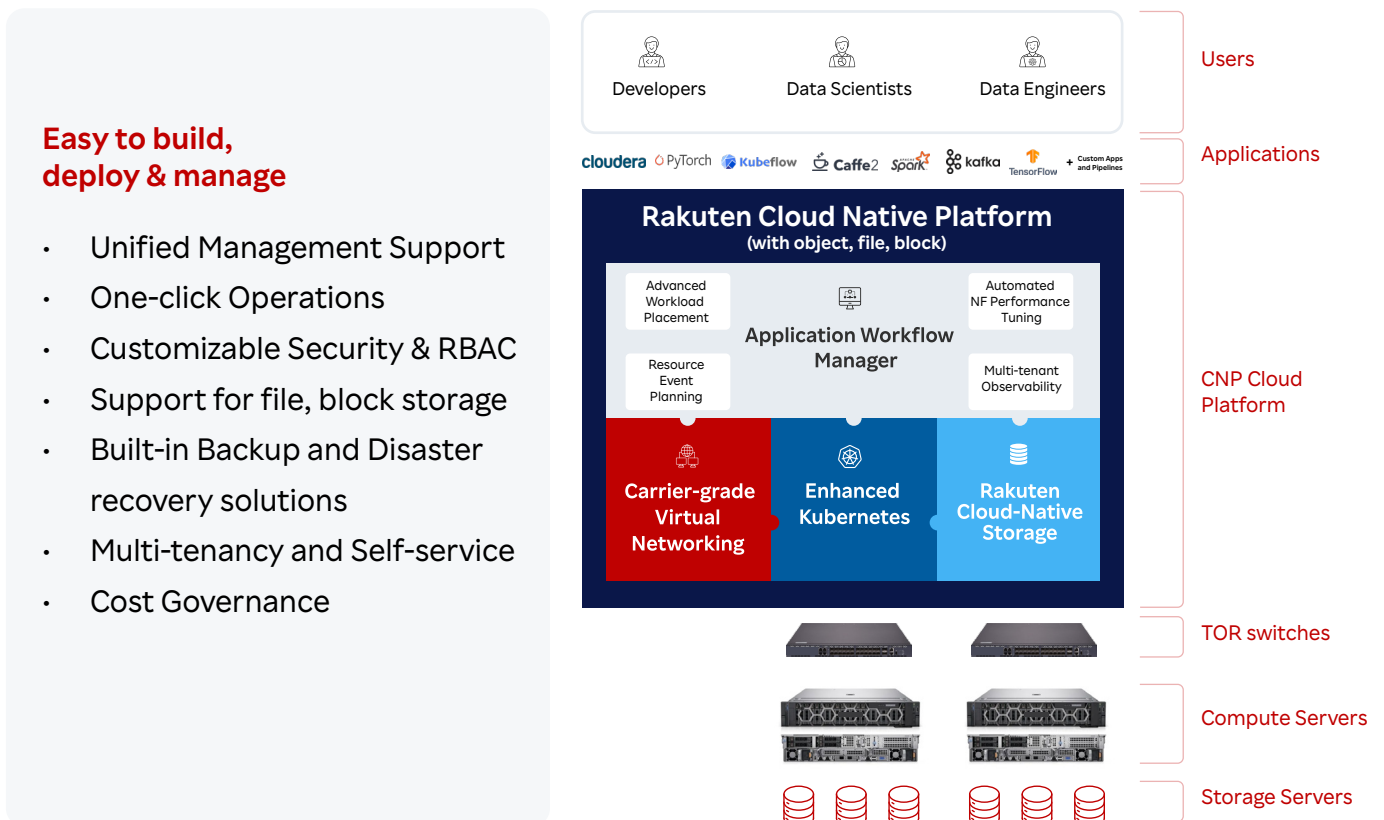


Figure 2 – Rakuten Cloud-Native Platform architecture

Additionally, it has the capability to allocate storage optimized for specific data types.

Rakuten Cloud offers a range of application-aware storage and data management features including snapshots, clones, replication, backup, data rebalancing, tiering, thin provisioning, encryption, and compression.

One-click deployment of AI pipeline



Solution Overview

Rakuten Cloud simplifies the creation of private cloud platforms for AI and analytics applications. The solution leverages Rakuten Cloud-Native Platform, standard hyperconverged servers, GPU-enabled servers, and integrated software-defined storage, providing block, file, and object storage from available servers and HCI platforms.

Easily deploy and manage entire pipelines





Figure 3 – Deployment Management of entire pipelines

The solution reduces the effort required by data engineers and scientists for deploying and managing analytics or AI stacks. Data scientists can easily deploy common AI and analytics applications on demand via the application catalog interface,

without the need for integrations or infrastructure. It also offers an application store-like experience, where a complex analytics pipeline can be deployed on-demand, and many of post-deployment day 2 operations tasks are automated.

One-Click Deployment of Data-Centric App Pipelines and Databases



Application is up and running in minutes, not weeks

- ✓ K8S components auto created (StatefulSets, PVC, Services, ...)
- ✓ Storage & Networking provisioned
- ✓ Data-locality, anti/affinity policies enforced
- ✓ Run workflow hooks to customize application

Common provisioning workflow for any Big Data, NoSQL, Database, AI/ML app

Figure 4 – One-click deployment of data-centric app pipelines and databases

Rakuten Cloud also provides app-aware storage, which is crucial for managing the variety of data types needed to support AI Cloud for Orange Africa. Rakuten Cloud-Native Storage is an enterprise-grade storage stack that aggregates all the available storage and distributes it to all the applications.

Key business Outcomes and Advantages

The solution delivered proven benefits to the IT and data scientists teams including:

- Accelerated time to value for analytics applications, reducing delivery from weeks or months to minutes and achieving the fastest time to value.

- Proven savings in CAPEX by enabling significantly better utilization and efficiency from available hardware.
- A cloud-like experience as a service, simplifying deployment, management, and day to day operations, which reduces OPEX.
- Cloud-like experiences at a much lower total cost of ownership (TCO).
- Accelerated time to insight for selected applications critical to the success of Orange Africa's analytics projects.




Figure 5 – Solution outcomes

Rearchitecting AI Cloud Platforms: A Data-First Approach

The ecosystem of applications in the market covering data analysis lifecycle is complex and diverse. Most of the data scientists are familiar with standard query and analytics engines available as open-source or community editions, such as Spark, Apache Airflow, Trino, Hive Metastore, Jupyter, and others. Even if supported by a cloud-native platform that simplifies the deployment and day 2 management of such workloads, scientists and analysts may still have complex

requirements. These include the storage of vast volumes of data, onboarding of disparate data sources, data governance and management of PII-sensitive information, composition and integration of the applications, and the application of AI/ML techniques to data. This drives a paradigm shift for AI cloud architecture from being application-centric to data-centric.

To reduce the time to value of data pipelines, a comprehensive data lifecycle management platform is essential. Rakuten AI for Telecom simplifies the ingestion, processing, transformation and analysis of data throughout its lifecycle, reducing time to value from months to days. This unified interface is built on well-known open-source tools, such as Trino, Spark, Jupyter, Airflow and others, lowering the entry barrier and costs for data scientists already familiar with these tools, allowing them to concentrate on extracting value from day zero.

Rakuten AI for Telecom is fully pre-packaged with Rakuten Cloud-Native Platform and supports

the onboarding of heterogeneous data sources, including object databases, and relational databases (RDBMS) such as MySQL. Rakuten AI for Telecom provides interfaces for integrating existing data sets, and building aggregates pipelines for day-to-day data workloads. This ensures minimal data duplication by using connectors built on top of an enhanced Trino layer. Supported plugins include Cassandra, MySQL, Kafka, ClickHouse, Hive, Elasticsearch, Iceberg, MongoDB, Oracle, PostgreSQL, Prometheus, Snowflake, and OpenSearch, and more. Rakuten AI for Telecom also supports design of pipelines based on “Canvas-like” plug-and-play function, as well as advanced-level configuration via file or class upload.




Figure 6 – Rakuten AI for telecom capabilities

Capabilities of Rakuten AI for Telecom include

Data Onboarding

A centralized migration tool enabling users to onboard new applications to the S3 object store and data bus. It allows users to efficiently store, manage, and access data based on their business requirements.

Data Discovery

A metadata management tool for data lakes offering classification and categorization capabilities, identifying

sensitive and confidential data, storing complete data lineage to reveal data lifecycle, democratizing data across the organization, publishing inventory details, column types, and data descriptions, enabling easy search and identification of data using search tags and columns.

Query Engine

A powerful tool allowing users to effectively query data across multiple data sources, such as S3 compliant object stores, using simple SQL-like commands and APIs. It offers extensive possibilities for visualizing datasets through interactive dashboards.

Data Transformation

A no-code tool that simplifies the process of defining, submitting, and scheduling Spark jobs through a user-friendly interface. It empowers users to perform data transformations without requiring extensive programming skills, making it accessible to a wider range of users.

Data Governance

A comprehensive platform designed to enable, monitor, and manage data security across platforms. It provides data owners with tools to prevent unauthorized use, mask confidential information, and enhance control for data security teams. Key features include centralized security administration via a user interface or REST APIs, fine-grained authorization for specific actions or operations, standardized authorization methods across components, support for various authorization methods, and centralized auditing of user access and security-related administrative actions.

AIOPS

AIOPS provides real-time visibility and valuable insights into the operational status, performance, and health of an application. With AIOPS, businesses gain the ability to efficiently monitor and manage various aspects of application deployment, execution, and maintenance.

Model Creation

The Rakuten AI for Telecom platform provides a comprehensive set of tools and features that simplify the model creation and execution process for machine learning developers, data scientists, and data engineers. Users can effortlessly develop notebooks and integrate them seamlessly into a pipeline, ensuring a smooth flow of data and operations. The platform provides a convenient scheduling mechanism, enabling users to easily define and automate the execution of their notebooks based on specific business needs and timeframes.

Network Monitoring Case Study for a Tier-1 Operator

Anomaly detection and KPI forecasting use case is a prime example of how the Rakuten AI for Telecom addresses critical business challenges in network monitoring and performance management.

Delivering Solutions to Meet Customer Requirements

To ensure optimal service quality and network reliability, there was a need for a solution that could monitor network performance KPIs, detect anomalies, and forecast trends. The goal was to proactively identify and address network issues before they impact business operations or end-user experience.

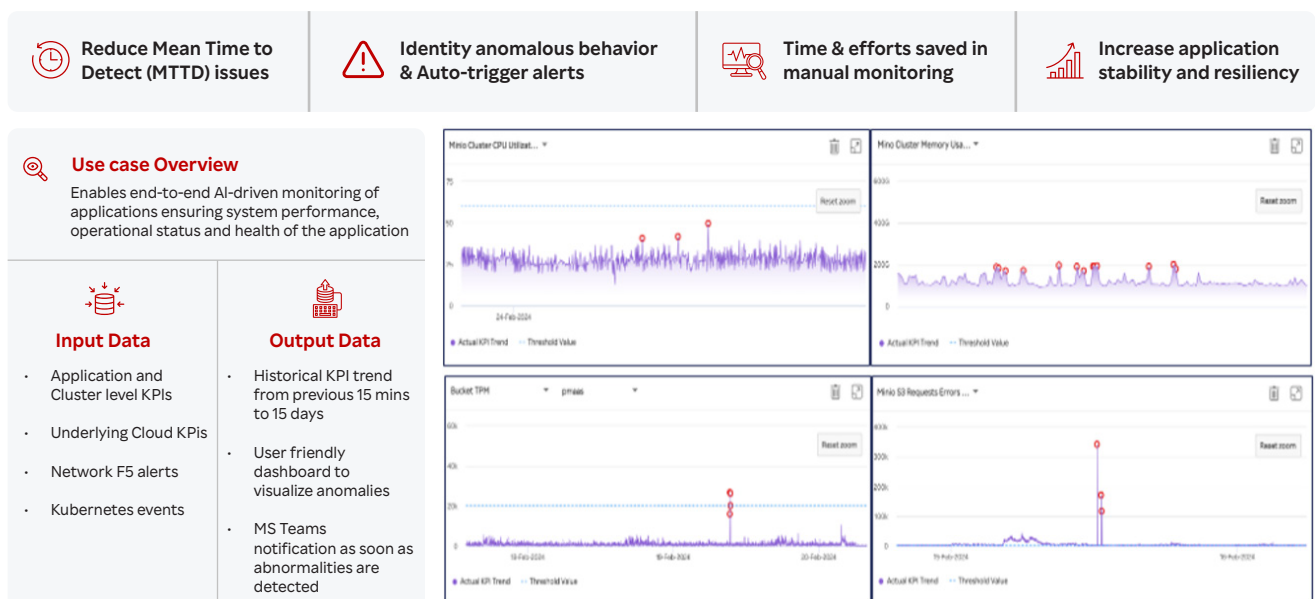


Figure 7 – KPI monitoring

Solution Overview

For Network Monitoring Teams:

- Automated anomaly detection without the need for manual thresholds.
- Visualization of anomalies and forecasted KPI trends.
- Severity-based anomaly scoring to prioritize actions.
- Alerts to enable proactive issue resolution.

For Security Operations Center (SOC):

- KPI forecasting to prepare for potential anomalies.
- Seamless integration with existing Performance Monitor applications for visualization.
- Confidence band calculation to indicate uncertainty in anomaly detection.

Anomaly Detection & KPI Forecasting

Anomaly detection and KPI forecasting solution leverages advanced machine learning to monitor network performance KPIs, detect anomalies, and forecast trends. It eliminates reliance on manual thresholds, delivering automated, real-time insights into network health. The solution integrates seamlessly with Performance Monitor applications, allowing teams to visualize anomalies and forecasted values directly within the application. Detected anomalies are highlighted in red, and an anomaly score is assigned based on severity. Alerts are triggered to ensure prompt action, while confidence bands indicate the level of uncertainty in the detected anomalies.

Abnormality Detection

Identify abnormalities in the system early using AI models thus reducing MTTD (Mean Time to Detect) and ultimately MTTR (Mean Time to Resolve)

Features in production: -

- Anomaly or abnormality detection feature is implemented for both MiniO cluster and bucket-level KPIs, with graphs that display the severity & date time with value of anomalies.
- Allow users to add and visualize relevant KPI trends in graphical formats, enabling informed decision-making aligned with business requirements.



Figure 8 – Anomaly detection and KPI forecasting

This solution is set up on Rakuten Cloud-Native Platform with access across all network domains, ensuring comprehensive coverage and proactive issue resolution.

Here are some of the key benefits of the solution:

- Automated Anomaly Detection: Detects anomalies across all network domains without the need for manual thresholds.
- KPI Forecasting: Predicts future KPI trends, enabling proactive planning and preparation for potential issues.
- Anomaly Scoring: Assigns severity scores to detected anomalies, helping prioritize actions based on impact.
- Confidence Band Calculation: Provides a measure of uncertainty in anomaly detection, aiding in more informed decision-making.
- Integration with PM Applications: Seamlessly visualizes anomalies and forecasted values directly within an existing Performance Monitor app.
- Automated Alerts: Triggers alerts for relevant teams to address issues before they escalate.
- Scalable and Adaptable: Designed to handle large-scale network data and scale with growing demands.

Storage and Data Management

The solution processes and analyzes large volumes of network performance data using advanced machine learning algorithms.

It supports efficient storage and data management by:

- Handling both historical and real-time data for anomaly detection and forecasting.
- Ensuring scalability to accommodate growing network demands.
- Providing automated data processing and visualization for simplified operations.
- Leveraging Rakuten Cloud-Native Platform data management and data protection capabilities for resiliency and automation.

Conclusion

In today's data-driven landscape, organizations face a critical challenge not in computing power, but in the reliability and management of data pipelines and analytic workflows that power AI-driven innovation. According to a joint study from data.world and Data Kitchen [ref], 50% of surveyed data engineers spend too much time maintaining pipelines and manual processes.

Rakuten Cloud-Native Platform, coupled with Rakuten AI for Telecom, delivers a fully integrated, cloud-native platform for running AI and big data workloads at scale. Built on an industry-leading Kubernetes® foundation, it abstracts platform complexity and provides an Application-as-a-Service experience for deploying and managing analytics pipelines, machine learning workloads, and network operations use cases. It enables seamless orchestration of CPU-based and GPU-based workloads, and ensures data sovereignty through private or hybrid cloud deployment models.

The platform supports the complete data lifecycle—from onboarding and transformation

Here are some of the key benefits of the solution:

Operational efficiency

- Automated anomaly detection and alerts reduce manual effort.
- Faster issue resolution minimizes downtime, enhancing service reliability.

Improved service quality

- Proactive detection and resolution of network issues enhance end-user experience.
- KPI forecasting prepares teams for potential anomalies, ensuring consistent performance.

Enhanced decision-making

- Confidence bands and visualized data provide actionable insights.
- Anomaly scoring enables prioritization of issues based on severity.

Cost savings

- Reduced downtime and manual intervention lower operational costs.
- Scalable design ensures optimal resource utilization.

to querying, governance, and real-time processing — through a unified interface powered by open-source, solution supports heterogeneous data sources including RDBMS and object stores.

Rakuten AI for Telecom also includes pre-integrated use cases such as anomaly detection, KPI forecasting, and AIOps. The solution enables fast, self-service deployment of complex pipelines via a one-click interface, and includes Day 2 operations tools such as observability, policy enforcement, and lineage tracking.

By delivering Analytics-as-a-service with built-in automation, scalability, and cost control, Rakuten Symphony helps enterprises and telecom operators accelerate time to value, simplify operations, and transition to a data-centric architecture purpose-built for AI-driven decision-making.

Start your AI and automation journey today.

Interested in a demo or have questions?

Reach out at cloud-tsp1@rakuten.com

Rakuten Symphony

symphony.rakuten.com

KUBERNETES® is a registered trademark of the Linux Foundation in the United States and other countries, and is used pursuant to a license from the Linux Foundation.

All other trademarks and copyrights are property of their respective owners and are only mentioned for informative purposes. Trademark indications such as ™ and ® may be omitted in the description of the trademarks of each company.