



Cloud-Native AIOps for Predictive Analytics Across Hospitality and AgriTech

Hiroshi Tanaka

Asst Professor

Abstract

Cloud-Native AIOps enables real-time predictive analytics across hospitality and AgriTech by integrating cloud-native architectures, event-driven processing, and AI/ML governance to deliver scalable, observable, and resilient insights. Predictive machine-learning models for demand forecasts, customer experience, climate resilience, and logistics optimization benefit from improved execution quality and business impact.

Cloud-native AIOps represents a paradigm shift in operating digital services, software, and applications and enables real-time AI/ML-driven predictive analytics across hospitality and AgriTech ecosystems. A cloud-native architecture is combined with an event-driven philosophy, AI/ML deployment governance, and observability concepts for AIOps approaches to event processing. Data ingestion and processing across the technology and business stack are designed for real-time applications, ensuring data is readily available for predictive analytics. These elements unite to enable greater agility, scalability, resilience, and observability in predictive machine-learning models.

Keywords : Cloud-Native AIOps enables real-time predictive analytics across hospitality and AgriTech by integrating cloud-native architectures, event-driven processing, and AI/ML governance to deliver scalable, observable, and resilient insights.

1. Introduction

The monumental rise of artificial intelligence (AI) gives an illusion of solution for all domain problems, but it is not that straightforward. Proper integration of AI life cycle, from data requirements through services to proper model deployment and observability, is required. A walk-through AIOps enables real-time predictive analytical needs for application in both hospitality and AgriTech domains that adopts serverless, event-oriented, cloud-native architecture. Domain knowledge reveals complexity in data collection, so prediction is reliable if done for a shorter time span. The data streams come from sensors in the AgriTech space, so data flow into cloud services is not an issue. The added complexity is that creating automation with the model trained and tested on a single region does not yield reliable results for other regions; predictive models are also to be trained at regular intervals using historical data for different regions. Observability is another area needing attention, be it

for prediction, detection, and resolution of anomaly or any other trend for making business decisions.

Cloud-Native AIOps enable real-time predictive analytics across hospitality and AgriTech by integrating cloud-native architectures, event-driven processing, and AI/ML governance to deliver scalable, observable, and resilient insights. Although the hospitality domain may appear straightforward, operational analytics using more granular data can yield better manageability of distribution throughout the operational hours; real-time prediction of expected guests enhances the orchestration of resources to improve customer service. In the case of AgriTech domains, the focus is more on Crop Yield prediction in the light of climate change—its inevitability or possibility can be factored into the supply chain and logistics optimization to minimize wastage of perishable products. Both domains involve a degree of personalization, whether of service delivery to role-based customers or selection-based delivery of advice, products, and services to farmers.



1.1. Background and Significance

Cloud-Native AIOps enables real-time predictive analytics across hospitality and AgriTech by integrating cloud-native architectures, event-driven processing, and AI/ML governance to deliver scalable, observable, and resilient insights.

Business ecosystems require applications that seamlessly adapt to fluctuations in demand, infrastructure, and risk. Stable Partitions facilitate elastic growth and shrinkage for instant computing and storage services. The growth of publicly deployed Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) enables the emergence of Cloud-Native Architectures. These cross-domain, multi-application architectures reside completely in IaaS Datacentres and enable the creation of platform-based multi-tenant applications. Observability becomes an intrinsic service Property of Cloud-Native Designs. By integrating Event-Driven-Decision-Making, Cloud-Native AIOps enables Real-Time Predictive Analytics across the Hospitality and the AgriTech domains.

The increase in Fan-Out of Digital Applications is changing the way Artificial Intelligence and Machine Learning Model Predictions and Decisions are being made. Instead of AI/ML Models doing Predictive Triggers based on Historical Data and creating steady Stream of Alerts, Events, and Notifications for connected Applications, Explicit Decisions by AI-ML Models for Actions and Remediations relevant to Business Operations must be delivered in a sequential stream of Event Messages across domain-specific Businesses. In Bandwidth-cost-sensitive regions, the need is to deliver Batch Invocations with lower Fan-Out Cost.



Fig 1: AIOps Tools

2. Background and Problem Statement

Historically, traditional on-premises analytics frameworks have been inadequate for new data sources of Type-2 and -3 analytics, including social media, eCommerce transactions, and digital market research, which are generally unstructured and subsequently require a cloud-based environment. In the hospitality industry, for example, Type-4, such as server logs and user activity traces, are often not presently analyzed and collected. At set intervals they may be processed offline. Consequently, the standard approach is neither timely nor effective. Internet of Things and smart technologies are now more prevalent, and structured Type-1 data influx directionally from sensors and measuring devices offers new opportunities to provide predictive analytics. Western Australia is leading the way in AgriTech during the transition from traditional agribusiness practice to real-time supply chain management enabled by Internet-of-Things and smart sensor technologies, while the application of Social Media and Advertising-based analytics in marketing display in hospitality and AgriTech domains aids real-time segmentation and personalization of offerings. The onset of Type-3 Big Data analytics closed the gaps present with traditional Type-1, -2 and -4 but is not able to address the variable-centric demand and supply forecasting in hospitality operations.

As the hospitality sector strategizes for recovery from COVID-19, emphasis is being placed on understanding real-



time customer sentiment to map demand fluctuations. Conversational searches are tempered with the Discovery Model, whereby pre-determined venues are considered before topics of general interest. Seasonal occasions (for example, school holidays) heavily influence demand while Social Media analytics are now also part of real-time customer sentiment analytics. Advances in agribusiness from the adoption of smart technologies are producing rapidly increasing data streams aimed at improving margin efficiency of black-and-white milk production and processing. Environment impacts are a major factor in reducing the supply chain from paddock-of-origin to point-of-sale, and consideration for the long-term evaluation of technology investment is producing Climate Change Adaptation Plans (CCAP) and Climate Resilient Decision Assist Tools (DAST) for field and supply chain.

Equation 1: From AR → ARMA → ARIMA → SARIMA

Let y_t be demand at time t (e.g., bookings/day).

AR(p):

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \phi_p y_{t-p} + \varepsilon_t$$

where ε_t is white noise.

Using the backshift operator B where $By_t = y_{t-1}$:

$$(1 - \phi_1 B - \dots - \phi_p B^p) y_t = c + \varepsilon_t$$

Define $\phi(B) = 1 - \sum_{i=1}^p \phi_i B^i$:

$$\phi(B) y_t = c + \varepsilon_t$$

Step 2: Moving Average MA(q)

$$y_t = c + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \dots + \theta_q \varepsilon_{t-q}$$

With $\theta(B) = 1 + \sum_{j=1}^q \theta_j B^j$:

$$y_t = c + \theta(B) \varepsilon_t$$

Step 3: Combine → ARMA(p,q)

$$\phi(B) y_t = c + \theta(B) \varepsilon_t$$

Step 4: Add differencing → ARIMA(p,d,q)

If y_t is non-stationary, difference it:

$$\nabla y_t = (1 - B) y_t$$

For order d :

$$\nabla^d y_t = (1 - B)^d y_t$$

ARIMA:

$$\phi(B) \nabla^d y_t = c + \theta(B) \varepsilon_t$$

Step 5: Add seasonality → SARIMA(p,d,q)(P,D,Q)_s

Seasonal differencing:

$$\nabla_s y_t = (1 - B^s) y_t, \quad \nabla_s^D = (1 - B^s)^D$$

Seasonal AR and MA polynomials:

$$\Phi(B^s) = 1 - \phi_1 B^s - \dots - \phi_P B^{Ps} \quad \Theta(B^s) = 1 + \theta_1 B^s + \dots + \theta_Q B^{Qs}$$

Full SARIMA:

$$\Phi(B^s) \phi(B) \nabla^d \nabla_s^D y_t = c + \Theta(B^s) \theta(B) \varepsilon_t$$

2.1. Research design

Research design aligns with the integrated design science framework for business solutions. The main goal is developing a cloud-native AIOps framework with industrial-scale and operationally exhaustive predictive analytics spanning hospitality and agriculture technology ecosystems. A real-time solution is developed for ResDiary, an online reservation service firm within the hospitality subdomain. Solution guidance is validated through experiential and exploratory design an unstructured agricultural supply chain operating with Agro-IoT for crop yield prediction, demand-supply planning, and geo-temporal crop-agnostic climate-factor analysis.



Research collaboration with ResDiary, a cloud-based reservation management service for the hospitality sector, is substantively assisted by Dustin Fenton, chief technology officer, and Shubham Jharbade, a software engineer. Solutions for ResDiary initial demand forecasting, personnel efficiency tracking, and service quality assessment in the reservation ecosystem exemplify real-time online operational data analytics, while exploratory alignment with the Indian short-duration unstructured agricultural supply chain provides domain breadth across an entire ecosystem. Dynamic requirement supply prediction combines multiple crop yield models into a meta-model that predicts climatic sensitivity regardless of chosen planting. Data integration from remote satellite cloud services enables forward recall for at-risk plantations while aiding planting decisions only a few weeks in advance.


3. Conceptual Framework

The Cloud-Native AIOps Framework encompasses the architecture, processes, and components necessary for real-time predictive analytics delivery across hospitality and AgriTech. Cloud-native capabilities ensure that the architecture responds automatically to load, helps application code remain simple and resilient, and conforms to an observable design aligned with open standards. AIOps processes build on scalable streams of real-time events, combining event-driven data ingestion with AI and machine learning to manage predictive models over their full lifecycle. All operational analytics and machine learning governance processes can be applied and automated in a production context.

3.1. Cloud-Native AIOps Architecture

The Cloud-Native AIOps architecture delivers real-time predictive analytics in response to application events. Event-driven cloud-native applications generate and respond to events, enabling new observability patterns that combine AI with scalable microservices-based processing. Cloud-native

AIOps exemplify the combination of event-driven processing and AI that is transforming modern application architectures. The foundational Cloud-Native AIOps framework integrates these three highly complementary areas of development to deliver a coherent cloud-native capability for real-time delivery of predictive analytics across multiple functional domains of the business. These capabilities are realised in the context of hospitality by combining a real-time operational analytics use case with two high-value production demand forecasting use cases, customer experience and personalised recommendation. Repeating this exercise in an AgriTech context highlights different data sources, but the predicted outcomes remain broadly similar: the delivery of AI-informed insight and guidance across the infrastructure ecosystem.



3.1. Cloud-Native AIOps Architecture

A cloud-native AIOps architecture supports event-driven streaming, enabling predictive analytics in hospitality and AgriTech ecosystems. A cloud-native AIOps architecture employs a coherent and unified API for multiple services that are continuously available across data centers, supporting serverless deployments. Cloud-native AIOps-enabled predictive analytic capabilities across hospitality and AgriTech leverage the distributed and elastic capabilities of public cloud service providers. Event-driven streaming requires near-real-time processing of systems, data, and user



events. The cloud-native AIOps architecture supports an event-driven stream ingestion pipeline that tracks and prepares the state of the system, data, user, and other resources to support prediction use cases. Business and user experience enablement outside the core services also demand the rapid delivery of features in a scalable and resilient manner, which can be achieved through cloud-native AIOps. Continuous analyzation and prediction of resource demand and expectations and observability of service health increase the quality of service.

A cloud-native AIOps architecture provides a unified control plane that integrates visibility, optimization, and resilience of delivery across services and products. Managing the expectations of users is key, and collecting these insights can assist in timely product feature placements and enhancements. A hybrid cloud-native deployment model enables the observability and resilience of third-party service dependencies. Although external services can have impact on satisfaction levels, combining observation with prediction can provide a clear view of damage control and compensation required. A comprehensive view of user behavior, activity, and generated content associated with product features can provide direction for improving the underlying systems.



Fig 2: Cloud Intelligence/AIOps

3.2. Real-Time Data Ingestion and Processing

Event-driven processing is central to AIOps because repeated model training is rarely practical due to the

operational cost driving the need for adaptive systems. Such models are therefore generally trained only once or for a limited number of times, such as across an annual cycle; in the interim, only scoring is performed. Event sourcing supports logical event handling for both domain operations and real-time predictive analytics. Improvements, changes in operational behaviour, or data anomalies that may trigger retraining or an alternative operational model are captured as logical events.

Conventional stream processing platforms, including Apache Kafka, Apache Pulsar, Apache Flink, and Apache Spark Streaming, enable parallel processing of event data at scale, but these approaches tend to focus on data movement rather than control of behaviour across distributed components. For example, events with a particular disposition may need to invoke Business Process Model and Notation workflows, whereas failure events may initiate a remediation workflow. Capturing a broader spectrum of events that can have far-reaching consequences for the business, and allowing all supporting processes, both proactive and reactive, to be modelled explicitly with business process technology, can lead to order-of-magnitude improvements in performance and effectiveness. Such an event handling paradigm is most naturally realised through event sourcing, in which state changes are recorded as immutable events in a source of truth.

3.3. Predictive Analytics and Model Management

Data preparation and normalization, along with appropriate data selection, remain key to accurate predictions. Predictive models are typically built and operated within managed (governed) environments, such as those provided by cloud platforms. The needs of predictive AIOps vary by domain and depend on the type of model being supported, for example, whether Internal Model Risk Capital Adjusted VaR (VaR-IMRCA) or anti-money laundering (AML) models are in view, as discussed by. A combination of ADDX and Auto-PMML is proposed and used for VaR-IMRCA. Customers increasingly demand that financial service providers take a wider set of environmental, social, and



governance (ESG) factors into account, as articulated by. Applying Auto-PMML to a wider set of models makes it easier to adapt to new developments.

Demand-forecasting applications in healthcare, retail, and manufacturing must address key challenges: pandemic-related volatility and demand complexity; limited historical data; the unique characteristics of new product introductions; and the rich variety of forecasting targets attributable to complex multi-level product structures. Approaches based on deep learning and hierarchical time series modelling are among those considered.

4. Domain Contexts: Hospitality

Recent years have brought significant changes to the hospitality sector, accompanied by notable technological investment. Real-time operational data stream processing, forecasting algorithms, and demand-prediction models are increasingly utilized to manage day-to-day operations. Their integration into enterprise resource planners enables better coordination of tasks such as allocating staff or resources, ordering supplies, adjusting pricing in real time, and organizing marketing and promotional activities. Yet, while predictive analytics resulting from event-driven data processing in combination with machine learning (ML) model hosting and monitoring have matured substantially in the past few years and have reached cloud-native AIOps, the hospitality sector has not yet adopted the state of the art.

Customer experience and personalization have become determining factors for success. The adoption of decision-support systems based on real-time data analysis has enabled faster, more consistent, and more effective responses to customers. Real-time traffic updates, queue times, pricing changes, targeted marketing messages, and improved personalization of recommendations have all been shown to enhance customer satisfaction and loyalty. Even a modest increase in personalization rates can have a significant positive impact on revenues. Yet, despite these benefits, very

few players in the sector are using predictive analytics to enhance customer experience.

Event-driven data processing, predictive modeling, and ML model lifecycle management are essential tools for meeting these new demands. Leveraging these tools makes it possible to inform customers more accurately and quickly through multiple channels, improving the overall customer experience beyond perceived and actual service delivery. Adaptable, optimally timed, and rapidly responsive messages delivered to customers traveling, working, or shopping in proximity to a business can enhance conversion rates effectively.

Equation 2: SVR regression: step-by-step optimization

We build features x_t from calendar, lags, events, etc. Predict $\hat{y}_t = f(x_t)$.

Step 1: linear function in feature space

$$f(x) = w^T x + b$$

Step 2: ϵ -insensitive loss

We want errors within ϵ to be “free”:

$$|y_i - f(x_i)| \leq \epsilon$$

Step 3: allow violations with slack

$$y_i - (w^T x_i + b) \leq \epsilon + \xi_i \quad (w^T x_i + b) - y_i \leq \epsilon + \xi_i^*$$

$$\xi_i, \xi_i^* \geq 0$$

Step 4: minimize model complexity + penalty

$$\min_{w, b, \xi, \xi^*} \frac{1}{2} \|w\|^2 + C \sum_i (\xi_i + \xi_i^*)$$

subject to the constraints above.



Step 5: kernelize (nonlinear SVR)

Replace inner products with kernel $K(x_i, x_j)$ (e.g., RBF):

$$f(x) = \sum_{i=1}^n (\alpha_i - \alpha_i^*) K(x_i, x) + b$$

4.1. Real-Time Operational Analytics

Real-time event sourcing, point-of-interest geocoding, and machine learning-based analytical tools for hotels and other attractions make it possible to deploy special interest groups and events such as social group tours. The return on the investment may be computed in terms of expected revenue generation, financial risk reduction, or carbon footprint reduction. The hotel industry routinely collects data on consumption and revenue. A literature survey on Hotel Management Systems reveals that revenue steadily improved when these systems were systematically used, both for collecting data for prediction and adjusting daily rates according to the price demand curve. The addition of predictive analytics tools for personnel planning enables resources to be allocated according to the predicted demand. Event-driven Real-Time Operational Analytics merge the traditional real-time data metrics (like occupancy count, fuel consumption cost incurred in the previous hour) that get computed every hour and backfilled to create a complete batch in a 24-hour window, with other data input such as geographic source and social group predictions to implement Demand-Driven Resource Control in hospitality.

The effects of Machine Learning (ML) models built on data streams provide daily prediction of the future (that is, the traffic forecast for the next day). For operators of hospitality businesses — who routinely devote staff time and resources to planning day-to-day business staffing, shift timings, and equipment availability — this capability to forecast the future contributes to Business Closure Risk Reduction. With AI/ML/Software-as-a-Service (SaaS) offerings, these models may be developed and performed as a service, reducing the up-front investment in any domain-powered model management service. Service delivered by these Event-

Driven Algorithms may be billed to the customers either through a Combination Fixed-Variable-Coom-Ware Pricing Strategy, which prices the service and customizes the model according to size and operational profile, or Combination of Billable Model Exchange Charges as a Flat Percentage to the Airline Industry Customers through Ticket Price Notification Mailers.

4.2. Demand Forecasting and Resource Allocation

Hospitality demand generally follows predictable patterns determined by cultural and holiday cycles, weather, and the calendar. Accurately forecasting demand enables hotels to allocate resources appropriately, optimize pricing, and develop promotional strategies targeting low-demand periods. The ability to combine and analyze data from multiple properties opens up further opportunities for demand prediction as rooms or services can be consolidated at cloud valleys across locations, and transport chains can be optimized for multi-property inventories.

For a concise appraisal of cloud properties, demand forecasting models can be based on an ensemble of multi-variable autoregressive integrated moving-average models (SARIMA), XGBoosting regression, and support vector regression supported by a selection of polynomial transformations, particle swarm optimization-abetted optimization, Kalman filtering, and error-correction mechanisms. In a cyclic environment of polyvariant choice, Machine Learning can be harnessed for dynamic demand modeling in highly non-linear data. Neural Network-based models have boosted forecasting performance by reducing overfitting bias while taking into account the spatial conversation of understanding in the Product-Consumption space.

4.3. Customer Experience and Personalization

Real-time data from online booking systems, travel activity feeds and social media enable prediction of customer preferences, sentiment on a property and urgency to bid for accommodation or services (e.g. meals, wellness or activities). Configuration to recognize current guests on



property adds more granularity to provide increased targeting of services. Predictive models are maturing and delivery technologies based on user-defined preferences are emerging. Social posts can be combined with local operational data to surface instantaneous destination alerts. Hybrid-cloud data architectures are enabling hotel operators to break their traditional information monopolies.

Dynamic fare management (Itcher) utilizes known demand patterns, distribution channel-specific demand-driving factors and competitor rates to constantly tune offering rates on aggregators. Integration of major channel management systems into the architecture enables auto-updating of rates at affiliates. Real-time driver demand and inventory imbalance recognition are used to trigger actions to stimulate demand, with a strong focus on pay-per-click strategies and optimising return on marketing expenditure. Teams are exploring the use of airport mobile-waiting-data to predict outer-regional driver demand and negotiate for affiliate links with Ordometer and comparable services, reducing reliance on the reducing-usestring risk.

5. Domain Contexts: Agriculture Technology

Accurate forecasts are vital for the effective management of AgriTech supply chains because raw materials are vulnerable to price fluctuations and perishability. Accurate forecasts support supply chains by facilitating Just-in-Time delivery schedules and optimised stock holdings across multiple storage facilities. AgriTech operations benefit from effective use of low-cost sensors and data-integration solutions. Sensor-based data streams from disparate devices and systems on farms are geospatially associated to produce inputs that can be used to provide real-time insights and enhance data-driven decision making, for example, by monitoring the presence of climatic conditions conducive for diseases and challenges in crops such as blights, pests, or other diseases.

Crop yield prediction models trained on multiple years of data can assist farmers in seeding decisions. These models can use climatic and weather indicators as predictive variables, and careful selection of the indicators can make the model resilient to climate variability. Predictions of future yields can then assist participants and stakeholders in the supply chain and in logistics. The extra surplus may also be supplied to efficient resource-management schemes supporting food security or nutrition, which can allocate supply at periods and places of great demand. High temporal-frequency feedback can assist farmers and enable detailed production monitoring, so that extreme unsatisfied demand can lead to distress purchasing and higher prices.




Fig 3: Domain Contexts Agriculture Technology


5.1. Sensor-Driven Data Streams and Agro-IoT

In the AgriTech domain, Cloud-Native AIOPs is driven by data streams generated by sensors embedded in operational objects during ordinarily monitored processes. When combining these internal data streams with event data from the environment, the result is an Agro-IoT, consisting of a cloud infrastructure supported by sensor networks, enabling real-time detection of any abnormal situation on the farm and immediate alerting of farm personnel. Predictions based on recognized patterns can also be made. Current research is focused on using digital twins created with Agro-IoT technology to forecast future crop yields, considering climate data, soil moisture, and temperature information captured in the present.



Agriculture is the sector most affected by the climate crisis and requires major advances to adapt and increase its resilience. In normal years, agronomic decisions such as crop variety selection, planting date, and fertilization, among many others, are based on the experience of producers, agronomists, and technicians. However, with such a large number of information variables, modeling becomes increasingly complex. The greater the uncertainty for farmers, the greater the care and the safer the crop. Actual production in an area should therefore always be lower than the potential, and the supply must be established between this actual production and demand to avoid oversupply of the commodity, which leads to heavy losses for producers.

Currently, AIOps models are used to estimate crop yield for crops commonly grown in Brazil, such as soybean, maize, and sugarcane. In these models, climate data, soil moisture, and temperature information are used as input. In a final stage, the models are consolidated, and the regional climate forecast becomes the model's main input, allowing for forecasting at a regional level. In the Agro-IoT context, the model will be executed in a digital twin created for each monitored area.



5.2. Crop Yield Prediction and Climate Resilience

Predicting crop yields is vital for many stakeholders, including farmers, traders, and government planners, and

sensor data from multiple sources can assist in building reliable crop yield prediction models. Seasonal weather forecasts, soil moisture levels, and satellite imagery provide information on the occurrence and intensity of key climatic and biophysical drivers, and conditions and disturbances throughout a growing season in near-time. Agri-Tech Internet of Things (IoT) solutions generate real-time data streams from soil, weather, and irrigation sensors, and combine these with pest-risk warnings from satellite-acquired data. The daily data streams are subjected to real-time predictive analytics for forecasting crop yields or risks of infestation of specific crops.

Machine learning models applied on a state-level basis, and based on locally determined relationships between crop yields and the above factors, have demonstrated high predictive power for certain crops. Yield forecasts also assist import, export, and stockpiling decisions. Climate change is resulting in increased production variability for many crops and large-scale population movement in pursuit of livelihood opportunities. Yield predictions that utilize future climate-science forecasts may assist food security and livelihood protection in the turbulent times ahead.

5.3. Supply Chain and Logistics Optimization

In addition to predicting crop yields, ML models trained on sensor-driven data can also support the management of institutional supply chains, enabling the timely delivery of crops when locations are prepared for harvest. With the role of members as suppliers of their respective crops, the AgriTech companies management must plan accordingly to meet customer demand. Supply chain management (SCM) focuses on optimizing the delivery of services and products using a range of business processes, such as operations, revenue, supply, and distribution network planning (Duan et al. 2019). Furthermore, the interconnected nature of SCM and logistics requires separating these two areas, with logistics focusing on the efficient and timely delivery of products according to demand. These areas can be trained using IoT signals by using historical information (planting and harvest schedules, distribution flows, collection dates,



and expiration periods) combined with other signals like climate weather forecasts and urban flows.

Supply chain and logistics operations can benefit from the use of ML, such as SCM increasing customer satisfaction levels by matching operations with demand through the accurate prediction of inventory levels and delivery times, while multiple logistic companies started to invest in such techniques to expand their service levels (Duan et al. 2019). The operations can be defined with the aid of the Agro-IoT systems, evaluating the minimal cost operation during delivery accordingly to the fuel price along the periods.

6. Conclusion

Cloud-Native AIOps enables real-time predictive analytics across hospitality and AgriTech by integrating cloud-native architectures, event-driven processing, and AI/ML governance to deliver scalable, observable, and resilient insights. The continual growth of cloud-native architectures and increased data-driven decision making has brought forth a compelling need for real-time AIOps monitoring and management. Making all of an organization's operational data observably available and ready to produce actionable insights is a large task. Real-time operational analytics, demand forecasting and resource allocation, and customer experience and personalization are key drivers in the hospitality domain. Simultaneously, sensor-driven data streams, crop yield prediction and climate-resilient agriculture, and supply chain and logistics optimization are pivotal in the AgriTech sector.

Emerging trends like the wider adoption of sensor-driven data streams, a greater focus on climate- and environment-driven business practices, increased MLOps adoption strategies, the growing scope for integrating sensor-driven data streams, AIOps deployment monitoring platforms, and growing interest in decentralized cloud storage are pushing organizations toward full-fledged cloud-native AIOps implementation. After all, practically all organizations are

scaling their cloud-native architecture despite the on-going economic uncertainty, and the demand for fully cloud-native AIOps implementation is genuine. Cloud-native AIOps can ensure near real-time analytics for capable organizations while also ensuring near real-time capability monitoring and management.

Equation 3: "XGBoost regression": objective and additive model

XGBoost predicts with an additive ensemble of trees:

$$\hat{y}_i = \sum_{k=1}^K f_k(x_i), \quad f_k \in \mathcal{F}$$

($\mathcal{F}$ is the space of decision trees.)

Objective (generic form):

$$\mathcal{L} = \sum_i \ell(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k)$$

with a regularizer like:

$$\Omega(f) = \gamma T + \frac{\lambda}{2} \|w\|^2$$

(T leaves, w leaf scores.)

Step-by-step boosting update

At iteration t , current prediction $\hat{y}_i^{(t-1)}$. Add new tree f_t :

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + f_t(x_i)$$

Second-order Taylor expansion around $\hat{y}_i^{(t-1)}$:

$$\ell(y_i, \hat{y}_i^{(t)}) \approx \ell(y_i, \hat{y}_i^{(t-1)}) + g_{if_t}(x_i) + \frac{1}{2} h_{if_t}(x_i)^2$$

where

$$g_i = \frac{\partial \ell(y_i, \hat{y})}{\partial \hat{y}} \Big|_{\hat{y}=\hat{y}_i^{(t-1)}}, \quad h_i = \frac{\partial^2 \ell(y_i, \hat{y})}{\partial \hat{y}^2} \Big|_{\hat{y}=\hat{y}_i^{(t-1)}}$$



So the tree f_t is chosen to minimize:

$$\sum_i \left(g_i f_t(x_i) + \frac{1}{2} h_i f_t(x_i)^2 \right) + \Omega(f_t)$$

6.1. Emerging Trends

AI/ML governance, data lake optimization, and adaptable model lifecycles are crucial as demand for real-time predictive insights escalates. Data pipelines incorporating ingestion, transformation, storage, and delivery are no longer sufficient; real-time data explosion—driven by growth in telemetry, sensors, and connected devices—requires Continuous Data Engineering (CDE) to ensure quality, trustworthiness, and proper treatment during analytic modeling and consumption. Governance must now extend beyond ML, AI/ML Ops, and DataOps to include AIOps, where operations meet analysis and monitoring to better inform actions, especially critical during real-time adaptive processes. Alerts, triggers, recommendations, and actions no longer depend solely on past events or on humans; models embedded in the CDE process become predictors, foreseeing demand fluctuations so responses can be pre-emptively implemented.

Cloud-Native AIOps supports the complete lifecycle of operational analytics. A novel CDE framework emerges for Hospitality and AgriTech considering AI/ML Ops, DataOps, AIOps, real-time data processing, and event-driven architectures. Addressing the full lifecycle in an adaptable manner, with Observability and Data Trustworthiness as focusses, ensures scalable and resilient organizations that offer improved experience. In Hospitality, AIOps, Demand Forecasting, and Customer Experience are in the spotlight; AgriTech considerations include the Agro-Internet of Things (Agro-IoT), Climate-Resilient Crop Management, and Supply Chain Management.

7. References

1. Ahmed, S., Singh, M., Doherty, B., Ramlan, E., Harkin, K., & Coyle, D. (2022). AI for Information Technology Operation (AIOps): A review of IT incident risk prediction. In Proceedings of the 9th International Conference on Soft Computing and Machine Intelligence (ISCMi 2022) (pp. 253–257).
2. Bernabé García, G., García Clemente, F. J., et al. (2021). FARMIT: Continuous assessment of crop quality using machine learning and deep learning techniques for IoT-based smart farming. *Cluster Computing*, 24.
3. Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. *Universal Journal of Computer Sciences and Communications*, 1(1), 1-21.
4. Choudhury, S., et al. (2022). Agriculture field automation and digitization using Internet of Things and machine learning. *Journal of Sensors*, 2022, 9042382.
5. Das, R., Chadha, H., & Banerjee, S. (2021). Multi-layered market forecast framework for hotel revenue management by continuously learning market



- dynamics. *Journal of Revenue and Pricing Management*, 20(3).
6. Doborjeh, Z., Hemmington, N., Doborjeh, M., et al. (2022). Artificial intelligence: A systematic review of methods and applications in hospitality and tourism. *International Journal of Contemporary Hospitality Management*, 34(3), 1154–1176.
 7. Dong, W. (2022). AIOps architecture in data center site infrastructure monitoring. *Computational Intelligence and Neuroscience*, 2022, 1988990.
 8. Junaid, M., Shaikh, A., Hassan, M. U., Alghamdi, A., Rajab, K., Al Reshan, M. S., & Alkinani, M. (2021). Smart agriculture cloud using AI based techniques. *Energies*, 14(16), 5129.
 9. Reddy, V. A. R. (2021). Challenges in Standardizing Member Eligibility Data Across Multi-Payer Healthcare Ecosystems. *International Journal of Medical Toxicology and Legal Medicine*, 24(3), 1-19.
 10. Katiyar, S., & Farhana, A. (2021). Smart agriculture: The future of agriculture using AI and IoT. *Journal of Computer Science*, 17(10), 984–999.
 11. Kong, H., Wang, K., Qiu, X., Cheung, C., et al. (2022). 30 years of artificial intelligence (AI) research relating to the hospitality and tourism industry. *International Journal of Contemporary Hospitality Management*, 35(6), 2157–2177.
 12. Li, Y., Zhang, X., He, S., Chen, Z., Kang, Y., Liu, J., Li, L., Dang, Y., Gao, F., Xu, Z., Rajmohan, S., Lin, Q., Zhang, D., & Lyu, M. R. (2022). An intelligent framework for timely, accurate, and comprehensive cloud incident detection. *ACM SIGOPS Operating Systems Review*, 56(1), 1–7.
 13. Nobre Pereira, L., & Cerqueira, V. (2022). Forecasting hotel demand for revenue management using machine learning regression methods. *Current Issues in Tourism*, 25(17), 2733–2750.
 14. Saydam, M. B., Arici, H. E., & Koseoglu, M. A. (2022). How does the tourism and hospitality industry use artificial intelligence? A review of empirical studies and future research agenda. *Journal of Hospitality Marketing & Management*, 31(8), 908–936.
 15. Shaikh, F. K., Karim, S., Zeadally, S., & Nebhen, J. (2022). Recent trends in Internet-of-Things-enabled sensor technologies for smart agriculture. *IEEE Internet of Things Journal*, 9(23), 23583–23598.



16. Shaikh, T. A., Rasool, T., & Lone, F. R. (2022). Towards leveraging the role of machine learning and artificial intelligence in precision agriculture and smart farming. *Computers and Electronics in Agriculture*, 198, 107119.
17. Suhaimi, A. F., Yaakob, N., Saad, S. A., Sidek, K. A., Elshaikh, M. E., Dafhalla, A. K. Y., Lynn, O. B., & Almashor, M. (2021). IoT based smart agriculture monitoring, automation and intrusion detection system. *Journal of Physics: Conference Series*, 1962, 012016.
18. Thota, R. C. (2022). AIOps in cloud computing: Automation performance monitoring with AI and New Relic. *International Journal of Scientific Research and Management*, 10(3), 809–817.
19. Zhu, H., et al. (2021). Multi-horizon accommodation demand forecasting: A New Zealand case study. *International Journal of Tourism Research*, 23, 442–453.
20. Mattaparthi, R. (2021). Unified Data Lineage and Quality Governance Framework for Multi-Source Sensor Streams in Heavy-Duty Powertrain Manufacturing. *Online Journal of Mechanical Engineering*, 1(1), 1-15.
21. Ayoub Shaikh, T., Rasool, T., & Lone, F. R. (2022). Towards leveraging the role of machine learning and artificial intelligence in precision agriculture and smart farming. *Computers and Electronics in Agriculture*, 198, 107119.
22. Rathod, S., et al. (2022). Cloud computing and networking for SmartFarm AgriTech. *Journal of Nanomaterials*, 2022, 6491747.
23. Goyal, L., Malhotra, S., & Kumar, A. (2022). Study of applications of Internet of Things and machine learning for smart agriculture. In *Proceedings of the 2022 International Conference on Computing, Communication, and Emerging Technologies*.
24. Li, X., et al. (2021). Internet of Things for the future of smart agriculture: A comprehensive survey of emerging technologies. *IEEE/CAA Journal of Automatica Sinica*, 8.
25. García, G. B., et al. (2021). FARMIT: Continuous assessment of crop quality using machine learning and deep learning techniques for IoT-based smart farming. *Cluster Computing*, 24.
26. Hussain, A., et al. (2021). Survey for smart farming technologies: Challenges and issues. *Computers and Electrical Engineering*, 92, 107104.



27. Suhaimi, A. F., Yaakob, N., Saad, S. A., et al. (2021). IoT based smart agriculture monitoring, automation and intrusion detection system. *Journal of Physics: Conference Series*, 1962, 012016.
28. Pereira, L. N., & Cerqueira, V. (2022). Forecasting hotel demand for revenue management using machine learning regression methods. *Current Issues in Tourism*, 25(17), 2733–2750.
29. Das, R., Chadha, H., & Banerjee, S. (2021). Multi-layered market forecast framework for hotel revenue management by continuously learning market dynamics. *Journal of Revenue and Pricing Management*, 20(3).
30. Doborjeh, Z., Hemmington, N., Doborjeh, M., et al. (2022). Artificial intelligence: A systematic review of methods and applications in hospitality and tourism. *International Journal of Contemporary Hospitality Management*, 34(3), 1154–1176.
31. Loganathan, R. (2021). Integrated Risk and Compliance Frameworks for Global Data Center Operations: A Governance-Centric Approach. *Universal Journal of Computer Sciences and Communications*, 1(1), 1-26.
32. Kong, H., Wang, K., Qiu, X., & Cheung, C. (2022). 30 years of artificial intelligence (AI) research relating to the hospitality and tourism industry. *International Journal of Contemporary Hospitality Management*, 35(6), 2157–2177.
33. Dong, W. (2022). AIOps architecture in data center site infrastructure monitoring. *Computational Intelligence and Neuroscience*, 2022, 1988990.
34. Li, Y., Zhang, X., He, S., Chen, Z., Kang, Y., Liu, J., Li, L., Dang, Y., Gao, F., Xu, Z., Rajmohan, S., Lin, Q., Zhang, D., & Lyu, M. R. (2022). An intelligent framework for timely, accurate, and comprehensive cloud incident detection. *ACM SIGOPS Operating Systems Review*, 56(1), 1–7.