Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-114-5_4

DeepScience
Open Access Books

# Chapter 4: Cloud Platforms for Intelligent Operational Systems

## 4.1. Introduction

The importance of advancing intelligence into operational systems—such as intelligent transportation, smart energy, and intelligent supply chain systems—cannot be overstated. Such systems manage demanding physical processes for real-time operations, yet their intelligence has typically been limited to business systems that operate in time frames of seconds to days. Intelligent Operational Systems aim to combine event-driven, real-time data processing with external AI and machine learning for high-frequency, real-time decisions and high-sensitivity supervised learning. To date, the concept has largely been explored using on-premises, private-cloud infrastructure. The analysis here focuses on Cloud Platforms for Intelligent Operational Systems, with an emphasis on definitions, architectural patterns, supporting cloud platform paradigms, enabling data provisioning and processing, cloud support for AI and ML, and key components of reliability, scalability, and availability.

The term cloud provides strong intuitive guidance about what is possible from a cloud platform. That is, the service abstraction of Infrastructure as a Service (IaaS) enables the hosting of compute and storage resources in the cloud, which can then be capitalized, instrumented, and scaled to improve developer productivity and reliability. When resources are hosted on IaaS, the Cloud provider is responsible for deploying, managing, and operating the software packages. Platform as a Service (PaaS) provides an even more compelling operational product abstraction with benefits from both Execution as a Service (EaaS) and Software as a Service (SaaS)—lower operational costs and reduced software management overhead.

### 4.1.1. Overview of Intelligent Operational Systems

Intelligent Operational Systems (IOS) are applications that combine advanced automation with the ability to learn and adapt. To innovate through Operations as a Service, the delivery of Infrastructure as a Service and Platform as a Service must be extended beyond the traditional user-centred formulation. The continued development of serverless environments enables event-driven workflows to span the entire application stack – changes in business rules or operating conditions can thus trigger actions automatically in data ingestion pipelines, large-scale analytics or fleet management systems. The cloud also plays a crucial role in enabling advanced automation through Artificial Intelligence (AI) and Machine Learning (ML). Increased availability of machine learning models is a pre-requisite for the automation of real-time decisions across multiple operational business functions – demand forecasting, resource allocation, inventory maintenance and incident response. The application of AutoML, Federated Learning and transfer learning further eases the deployment and maintenance of ML models for real-time decision-making, supporting delegated autonomy in edge-cloud architectures for latency-sensitive workloads.

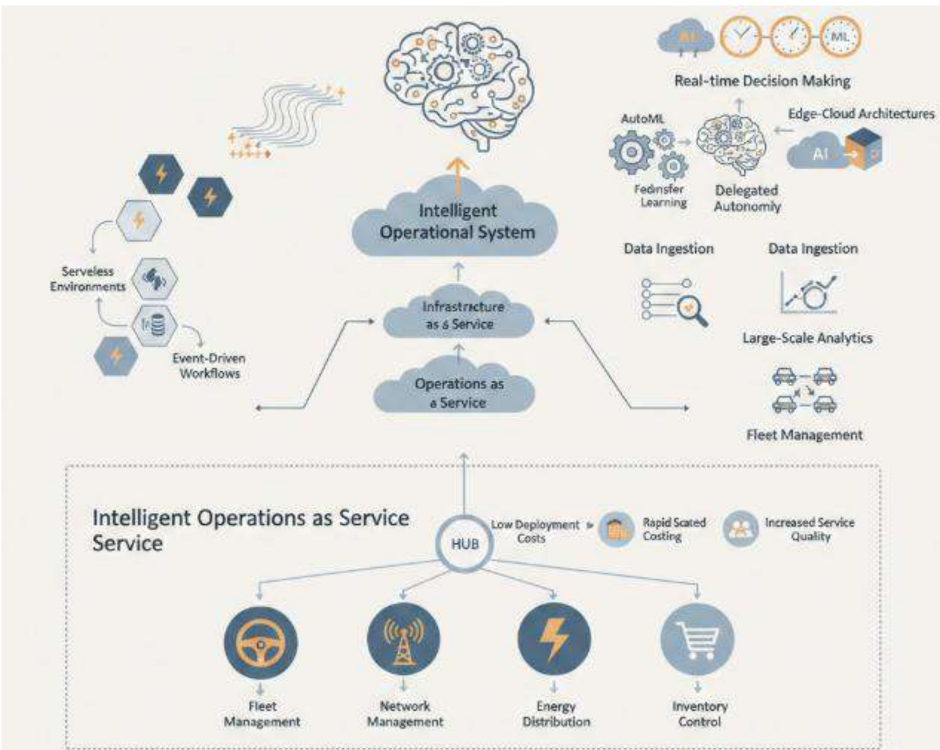


**Fig 4.1:** Augmented Operational Intelligence: A Service-Centric Framework for Adaptive Automation and Edge-Cloud Autonomy

An Intelligent Operational System consists of a logical business function that operates a finite, quantifiable resource to achieve a specific operational objective. Examples include fleet management (transportation), network management (telecommunications), energy distribution (utilities) and inventory control (retail). Application of the Intelligent Operations as a Service concept means that the development, deployment and operation of these logic functions are treated as a service offering. Relatively low deployment costs, rapid scaling and autoscaling, and an operating cost model directly aligned to demand are all essential if this logic function, and the process it supports, is to be successfully outsourced. Furthermore, integration of these functions into multi-site, multi-customer, and multi-agency data streams creates the potential for additional insight, new services and increased service quality.

## 4.2. Foundations of Intelligent Operational Systems

The notion of Intelligent Operational Systems refers to systems whose core mission is to enable and deliver intelligent operations, in the sense of efficient, reliable delivery of key operational processes for the entities they serve. A local government department responsible for the management of Water Supply, Wastewater, Refuse Collection, and Street Lighting, and their trial of an Intelligent Street Lighting Operation, is an example. Their deliverables are operations, and the Intelligent Street Lighting Operation enables this deliverable by enabling intelligent operation of the Street Lighting operation. This includes supporting the decision processes by providing information and decisions in real-time, alerting humans only when decisions fall outside a range. Data and Machine Learning models required to enable the desired real-time information, decisions, and alerts are generated from their operations using Advanced Analytical techniques.

Seven architectural patterns have emerged with sufficient frequency in systems supporting Intelligent Operations that they can be considered a common definition for intelligent operational systems. These seven patterns are: Data Lakes, Data Stream Processing, Automated Decision Making, Event-based systems, Infrastructure as Code, A/B Testing, and Predictive and Prescriptive Analytics. Architectures for Intelligent Operations can be built using one or more of these patterns, following the guidance provided by the patterns in the same way as architectures for Intelligent Surveillance and Intelligent Monitoring can be built following the guidance provided by their respective architectural patterns. The patterns of Intelligent Operations relate to Operational Systems more than any of the previous patterns described for Intelligent Surveillance and Intelligent Monitoring.

### 4.2.1. Definition and scope

Intelligent Operational Systems (IOS) execute planned operational processes with data-driven optimization and adaptation guided by Artificial Intelligence (AI) and Machine Learning (ML) models. Self-aware of the current operational state and history, they exploit the full range of historical and real-time data, directly ingesting and analyzing data streams in real time, rather than writing historical data and relying on periodically executed batches. the control signals generated by the automated Real-Time Analytical Models (RTAMs) adapt these processes' parameters to induce optimal behavior. The suitability of cloud platforms for supporting IOS is substantiated by the platforms' own architecture, which provides an operational foundation layer using Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) paradigms.

The adoption of the Serverless paradigm together with event-driven architectures further enhances these foundations by enabling automatic scaling in the temporal dimension according to the actual workload. At the data-in-the-cloud level, the IOS pattern is expected to take advantage of Data Ingestion, Streaming, and Data Storage services for the real-time lifetime phase. Easier, faster, and more efficient still is the direct use of Real-Time Analytics for both Decision Automation and Control Generation. Modern MLOps services thus enable end-to-end support for the entire lifecycle of ML models—and the corresponding models for decision automation—used by an Intelligent Operational System. A Cloud Edge Resource Layer guarantees latency-sensitive workload execution as close as possible to data sources and sinks, which are often distributed over large geographical areas.

### 4.2.2. Architectural patterns

A number of architectural patterns for operations can be delineated. In the Infrastructure as Service (IaaS) model, virtual machines execute traditional workloads, benefiting from elastic resources, database-as-a-service solutions, and assured data recovery provided by the Cloud provider. However, ramp-up, shutdown, and waiting periods introduce inertia, making IaaS less suitable for workloads of short duration. The primary benefit of the Platform as a Service (PaaS) model lies in deployment automation for web-based or RESTful services. Automated Integration neither provides a framework nor library for the user; it introduces a framework instead, actively implementing the external integration via triggers and web hooks. Serverless functions bridge the gap between Infrastructure and Platform as a Service, eliminating virtual machine management, and Event-driven architectures foster truly decoupled implementations. Services are structured as independent equivalents of Discrete Event Systems and communicate via an event bus.

In a more general sense, the process of generating intelligent operational systems can be categorized into three activities: (1) Data is gathered via Data ingestion services and prepared via Data preprocessing, Data alignment, and Data enrichment services; (2) decisions are then automated, following the flow initiated by Data streaming services in the Real-time analytics state; and (3) such services are generated as part of the Model auto-generation and deployment process, and made available to Data Integration services. These three activities can themselves happen in the Cloud, and the required services are supplied as needed, via a Cloud-based Automated Modeling System.

### 4.3. Cloud Platform Paradigms for Operations

Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) are the prevalent solution-independence delivery models sought when deploying cloud-enabled Intelligent Operational Systems. However, cloud-based Intelligent Operational Systems can be designed and implemented even when the solution is non-persistent and follows an Event-driven Architecture (EDA). Event-driven processing offers an alternative Cloud Platform paradigm: Serverless Computing, also referred to as Function as a Service (FaaS) because users purchase and pay for services offered by event-triggered serverless functions (i.e., actions) within the provider's Cloud Platform. Yet, maintainable and cost-effective implementations of other Architecture Patterns may be better suited to certain workloads. Some of these, primarily those with efficient demand load stabilization for elastic auto-scaling, may execute their time-consuming operations using IaaS, PaaS, or a combination of these models, while the other workloads may also follow Serverless Computing.

Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) are the prevalent solution-independence delivery models sought when deploying cloud-enabled Intelligent Operational Systems. However, cloud-enabled Intelligent Operational Systems can even be designed and implemented under a Serverless Computing paradigm, which follows an Event-driven Architecture (EDA) because users purchase and pay for services offered in an event-triggered manner by short-lived, event-triggered Vendor Functions within the Cloud Platform (the event itself acting as triggering service request). Yet, maintainable and cost-effective implementations of other Architecture Patterns may be better suited to certain workloads, determining a mix of implementation models for the cloud-enabled system. Some workloads, primarily those with efficient demand load stabilization for elastic auto-scaling, may execute their time-consuming operations using IaaS, PaaS, or a combination of these models, while the other workloads may follow Serverless Computing.

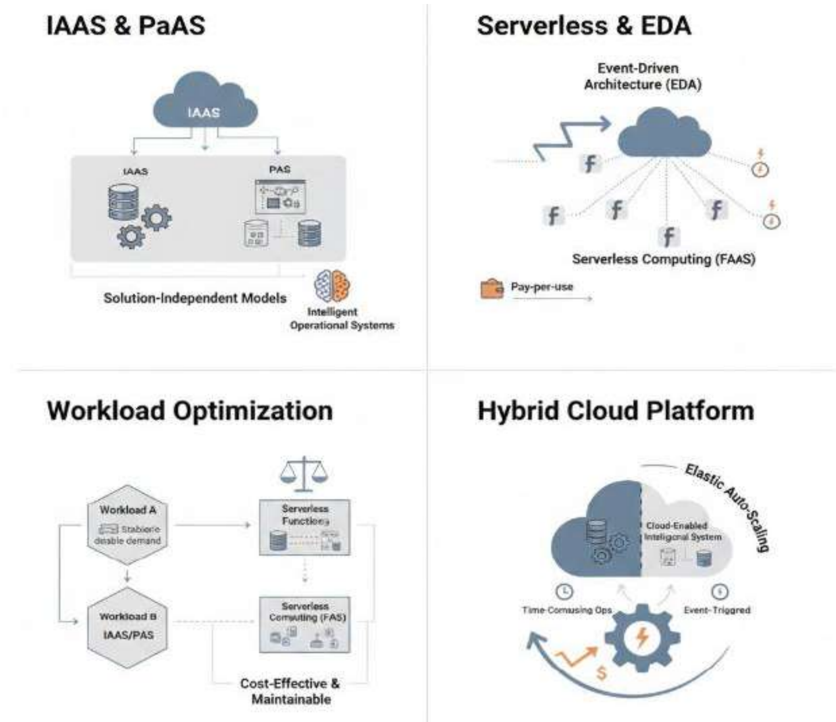


**Fig 4.2:** Hybridizing Cloud Delivery Models: Balancing IaaS/PaaS Persistence with Event-Driven Serverless Paradigms for Intelligent Operational Systems

### 4.3.1. Infrastructure as a Service and Platform as a Service for operations

The three operational paradigms seeding Intelligent Operational Systems are Infrastructure as a Service and Platform as a Service, respectively, for systems and functions. Enables for data ingestion, automated decision-making, and execution of operational functions are commonly implemented using the Serverless and Event-Driven Architecture paradigms.

Common infrastructure in cloud computing for IT workloads extends to operations. Infrastructure as a Service (IaaS) integrates cloud resources into operational functions without demanding consideration of the overall architecture. Allocation of virtual machines, volumes, networks, or add-ons are activated as needed by the custom operational demands of the agency, delivering unlimited resources under a pay-per-use model. Nevertheless, these technologies may be unsuitable for missions where mission cost is unavoidable. An IaaS infrastructure also adds overhead to the operation, requiring specification of machine size, location and price. Additionally, IaaS workloads have assigned lifecycles, demanding the implementation of scripts for shutdown and startup in periods of low use.

Platform as a Service (PaaS) also extends into operations, simplifying operational function development with a combination of hosting, databases, application development environments and common libraries. PaaS functionality accelerates development, reducing risks and ensuring high quality of operations. These two paradigms meet the demand of traditional data-intensive operations, such as custom databases for telemetry elaboration or business intelligence reports.

### 4.3.2. Serverless and event-driven architectures

Cloud Platforms are evolving to support more diverse and sophisticated use cases. Among the new types of Cloud service offered by major Cloud providers, Serverless computing and Event-driven computing servers have emerged as popular paradigms to support IT workloads that are either scalable, elastic, rapidly changing or constrained by other factors. Cloud users deploying such workloads no longer need to provision system resources, monitor their use, or adapt them to changes in demand. Serverless computing and Event-driven computing paradigms are becoming important for Intelligent Operational Systems that take advantage of recent advances in Artificial Intelligence/ Machine Learning and Cloud security and that are generically Fault Tolerant.

The Serverless execution model enables Cloud users to run applications without having to explicitly provision or manage resources. A Serverless application is composed of small pieces of code, called functions, that are executed in response to events. The Cloud service provider is responsible for starting and stopping execution environments based on demand, as well as scaling out to support large increases in load. Serverless computing supports very high degrees of autoscaling and elasticity. Applications are normally billed based on execution time, which follows the actual usage of the application rather than a provisioned capacity that has to be paid for even when unused. Despite the name, an application that uses serverless computing still runs on physical or virtual servers that are managed by the Cloud service provider. Serverless computing is sometimes considered a specific subtype of Event-driven Computing.

Cloud Providers also offer an event-driven execution model in which the Cloud service provider runs server resources and executes user-defined code components in response to events that are generated by the user or by other resources running in the Cloud infrastructure. Recent activities in lightweight application architectures, such as Microservices and Multicloud computing, have also contributed to the rise of Event-driven computing systems.

### 4.4. Data in the Cloud for Operational Intelligence

Numerous data sources feed cloud platforms and their intelligent operational systems. Data ingestion and management require a diverse set of technologies due to the high variability in both data formats and the scale of the ingestion. The cloud enables data streaming, either in batch or micro-batch mode. The capacity to automate decisions based on real-time data analysis enhances operational intelligence. Despite the fast nature of these decisions, the impact may span seconds, minutes, or even hours.

Applications that rely on critical, time-sensitive situations typically deploy decision automation and require very low-latency processing performed on the edge, near the data source. Other applications may leverage a fully cloud-oriented architecture for processing, and the nature of the business logic enables the possibility of higher latencies. However, architectural decisions must consider the business domain. In operational systems that value self-healing and reduction in mean time to resolution, decision automation plays a vital role. A connection with AI is established, encompassing both the ingestion and processing phases of the data.

#### 4.4.1. Data ingestion, streaming, and storage

Cloud platforms provide a plethora of specialized solutions for big data. The most common pattern is to stream data from devices or applications into the platform. In cloud parlance, such data aggregation is usually called "ingestion." Ingestion tends to occur in two distinct streaming service levels: Batch and Near Real-time (NRT) Batch Streaming. With batch ingestion, data are gathered and sent into the platform periodically by a process in the source system.

An example would be plotting the presence of customers in a mall at every midnight instant in order to find patterns. With NRT batch streaming, sources are monitored, and a pipeline processes all data created since the last process execution on a target datastore. It differs from CUDP in that CUDP sends smaller packets more frequently. In terms of cloud resources, it usually triggers a function in a serverless model that executes the pipeline. After ingestion, the deep storage is usually a second-factor data lake.

Cloud platforms provide many ready-to-use deep-storage solutions. These solutions form the second main characteristic for cloud platform operational-data storage architectures. Data lakes relying on deep storage fulfil the main requirements of storing cheap and large amounts of data for long periods. Virtual-connection-based access layers also allow proposing a layer-based partitioning of the data, at no additional storage cost. Consequently, an architect could propose a five-layered virtual-connection-based cloud-data-visualization model. All these layers are increasingly consolidated for event-driven use cases in a unique data-lake layer.

#### 4.4.2. Real-time analytics and decision automation

Advancements in scalable data-storage technologies and improved management strategies for Cloud resources enable Intelligent Operational Systems to perform complex data analysis at scale and near real-time. Browning real-time data to artificially narrow the problem space improves performance while retaining fidelity for certain workloads.

When incoming data must be used immediately to implement an operation, it needs to be analyzed and a decision made in a short amount of time. Challenges include automatically deciding which analytical preparation is needed for the data and which analytical function should be applied — if the latter is not expressed as a Service — and selecting a cloud resource that will minimize execution time, including for model training. For some operations, data from a limited period may be fed downstream even if it is not of sufficient fidelity. Parallel execution of Decision Assistant Services for multiple input streams with their own Decision Fusion Services can also be useful.

Insights from AI and ML may also be incorporated to eliminate or combine Decision Assistant Services. Graphical models that represent and map the relations between these Services for specific types of operations can help outline these decisions. Such a model should not be viewed as a formally proven design, but rather as inspiration for operations in that field. More complex operations inevitably require such a decision model to be developed specifically for them.

Operational Autonomous Agents are capable of executing operations directly without human supervision once needed conditions have been modeled appropriately. Low levels of attack detection in Operations Support Services may call for the Executive Director of Operations to define, prepare, and deploy Defence-based Autonomous Agents. They may also request Help Services to prepare a Supply Chain Defence System or rectify a failure detected in a single stream, and these keys may inductively develop into Homeopathic Operational Agents.

### 4.5. AI and Machine Learning in Operational Systems

Data-driven decision-making in intelligent operational systems relies on AI and ML models. These systems utilize models to classify data, recognize images, forecast demand, recommend products, detect fraud or churn, and optimize business processes. The ML lifecycle involves data collection, model training and testing, model deployment and monitoring, and model maintenance. Internally developed ML models undergo rigorous inspection before being deployed in production, while externally provided models are tested against an exhaustive set of expected inputs and corresponding outputs. With the deployment phase complete, ongoing monitoring and maintenance become crucial to keep models in top condition. MLOps (ML Operations) provide tools and methodologies to facilitate the ML lifecycle.

The operational use of data is primarily associated with complex decision-making queries, often constituting the final step in business flows. For tasks demanding low decision latency, e.g., fraud detection during online transactions and traffic signal control, the decision-making phase preferably occurs at the edge of the operational systems. All prior stages of these processes remain in the cloud, where large amounts of data are stored and processed. These complex, long-lasting data processes periodically build models that migrate from the cloud to the edge to reduce latency and make real-time decision making possible.

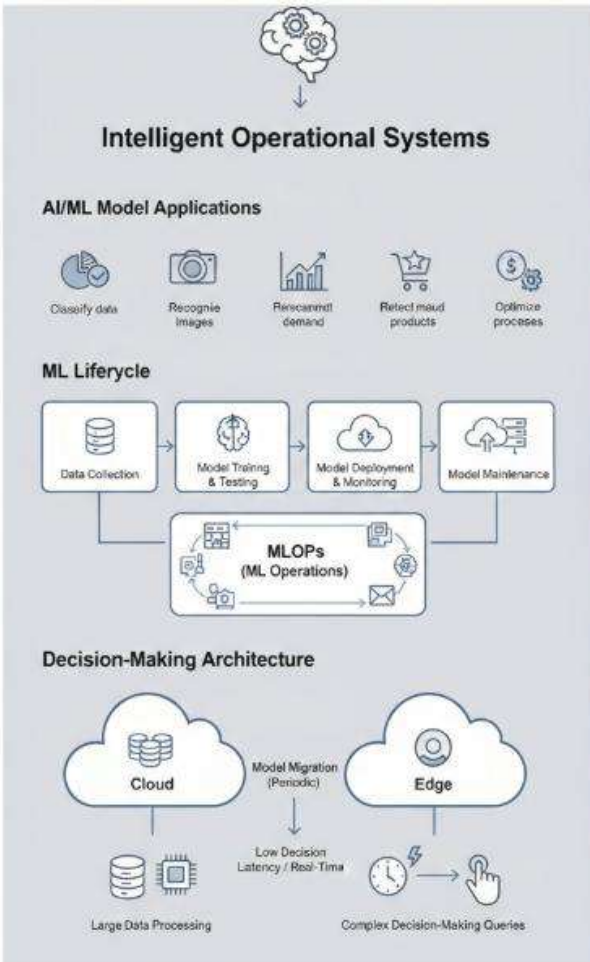


Fig 4.3: From Cloud to Edge: Optimizing the AI Lifecycle for Low-Latency Decision Support in Intelligent Operational Systems

### 4.5.1. Model lifecycle and MLOps

Operations encompass many tasks and processes, but one of the most captivating is that of anomaly detection supported by machine learning models. These models learn from past observations and experiences, mapping regular scenarios to a normal class while mapping anomalies to a set of abnormal classes. Recognizing the rarity of the latter category, machine learning models exploit generalization abilities to learn a normal pattern.

A key challenge lies in operationalizing models for detecting abnormal situations in production, enabling them to monitor real-world scenarios trained on historical data and escalate alerts when anomalies are identified. Considering that models are not conceived once and for all but rather for a given operational scenario and period, automated pipelines can manage the model lifecycle in an MLOps exercise. During a specific training period, a decision can be made whether to maintain or delete the operational model(s) in the cloud by following a decision-making-first approach: operates it when it is needed, not before. These autonomic pipelines exploit Cloud Infrastructure as a Service and Cloud Function as a Service both for the decision and operationalization stages.

Automating the complete model lifecycle provides several advantages, notably for anomaly detection and related tasks. The approach encapsulates the training and operationalization capability along with data streaming/data-at-rest processes.

### 4.5.2. Edge-cloud integration for latency-sensitive workloads

As more applications become “intelligent” by integrating AI and ML to analyze, predict, and automate decisions, edge-cloud computing continues to evolve. Although latency-sensitive, low-bandwidth data-generative applications may use edge devices for decisions and actuation, many services are open-ended and resource-intensive and therefore remain in the cloud. In such cases, Fast And Send Mostly (FASyM) image and video services are architecturally similar to established edge-cloud solutions, where serveless edge inference services preprocess input, reduce bandwidth requirements, and allow cloud ingestion, storage, processing, and decision automation in real-time.

These services also provide a new entry point to implement a highly photorealistic image synthesis pipeline. Other emerging applications, however, will require an even larger set of edge services, either because they use large amounts of data that cannot be processed directly on the edge appliance or because their definition is complex enough to demand a very large training dataset. In these cases, the emerging solutions must face the challenge of distributing ML inference over the edge-cloud path to reduce latency between data ingestion, processing, and action. In such a solution, processing is still iteratively addressed from edge devices and directly actuatable in the physical environment. The only difference stems from the necessity of introducing edge services also for the cloud data processing tasks: decisions that must comply with tight deadlines, such as object detection in intruder alerts or facial recognition in access controls, are still executed on edge premises, even in a distributed manner.

## 4.6. Reliability, Scalability, and Availability

Reliability, scalability, and availability are core requirements of Intelligent Operational Systems, especially when deployed in the Cloud. Technologies to address these requirements are abundantly available, yet few cloud-native intelligent operations projects use these technologies, often with serious consequences.

Fault tolerance of cloud-hosted systems can be achieved by replicating software components, using techniques already used for two decades in non-cloud deployments. High availability, rapid recovery from natural disasters, and support for business continuity plans rely on system redundancy plus off-site, up-to-date backups. Effectively implementing these techniques in public Clouds is easier than in non-Cloud deployments, since the Cloud provider handles physical-site redundancy, geo-clustering services, and rapidly changing Virtual Machine images for recovery.

Many vertical Cloud providers also recommend or enforce autoscaling by monitoring system utilization metrics and adaptively allocating Virtual Machine instances. Autoscaling has a less-well-known daughter technology for development and deployment: capacity planning. Capacity control can deliver timely response in intelligent operations while staying inside budgets, automating all the data engineering involved and delivering continuous deployment of Dashboards and Reports. When applied during development in Data Engineering, Capacity Planning ensures timeliness within budgets, by providing all necessary Data Engineering in advance, since it operates on historical data. The complete lack of provision for capacity control is a main reason so many business still use BI solutions that are primarily super-spreadsheets.

### 4.6.1. Fault tolerance and disaster recovery

Cloud Platform Paradigms for Operations emphasize the importance of fault tolerance and disaster recovery to ensure the ongoing operation of Intelligent Operational Systems. Resilience is required at many levels, including application architecture, Cloud Provider resiliency, fault-tolerant design principles, data redundancy, network planning, and policy management. The analysis draws both on the business side of the operations ecosystem and on the Cloud Providers' core offerings.

Intelligent Operational Systems automate critical aspects of business operations with real-time analytics powered by data from the operational environment. Failure of these systems can have dramatic financial and reputational costs, and can affect services for potentially millions of clients. Failure and recovery modes therefore need to be carefully understood on both the operational and service provision sides. Any single point of failure on the service provision side must be identified, and resiliency mechanisms built into the design. Cloud Provider mechanisms for redundancy should be fully considered, and their applicability, limitations, and costs understood. Cloud Provider disaster recovery solutions should be carefully configured, and Recovery Point and Recovery Time objectives aligned with operational requirements.

### 4.6.2. Autoscaling strategies and capacity planning

The Cloud offers a promising paradigm for implementing intelligent operational systems, enabling scalable environments that respond to demand fluctuations. Autoscaling strategies allow cloud services to scale up and down according to predefined metrics. These metrics can include resource load—be it CPU, memory, disk, or network interface—or custom application-specific metrics defined using the cloud provider's specific Application Programming Interface. Cloud functions are inherently serverless, automatically scaling with the number of concurrently running instances according to user traffic or event volume. Nevertheless, fail-safe and resilience strategies are paramount. Traffic spikes can overwhelm the underlying resources, prompting a breakdown, while temporary failures yield degraded services, such as long latencies, before system recovery.

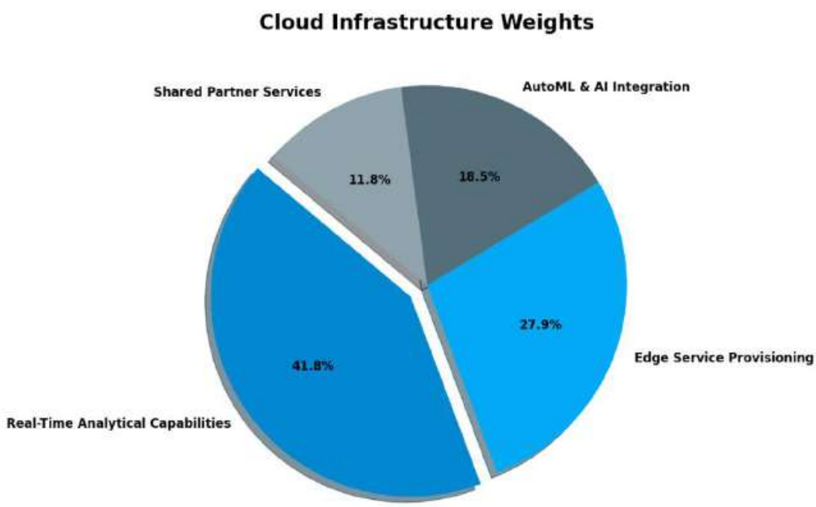


**Fig 4.4:** Cloud Infrastructure Weights

Autoscaling decisions rely on choosing ideal thresholds, periods, and increments for scaling the resource in question. Suboptimal settings can fail to fulfil demand, resulting in service disruption, or trigger unnecessary scaling, incurring operational costs. Thus, resource usage must be modeled and predicted over time, informing future capacity changes. Nevertheless, despite being critical, capacity planning often remains undeveloped in operation systems, as Scaling Operations can require complex preparations, especially for data transfer around Europe due to data protection implications. Detecting misalignments among resource status, configuration, and user demand can yield service disruptions and mistakes. Consequently, scheduling periodic capacity-planning reviews enables proactive action, particularly for major operational changes or for anticipating holiday traffic spikes.

### 4.7. Conclusion

Unlike many recent trends in information systems and services, Cloud Platforms have been proven essential for operational systems. Such platforms are responsible for the delivery of many services that an intelligent operational system needs. Cloud Platforms cover the Infrastructure as a Service, Platform as a Service, Software as a Service, and Marketplace-as-a-Service computing paradigms. Moreover, Cloud Platforms provide additional conveniences for operational intelligence: support for support for event and serverless architectures; ingestion, streaming, and storage modes for data; real-time analytical capabilities; MLOps for Artificial Intelligence; resiliency, scalability, and availability services. The future is in real-time operations, from offering services to propel business processes to provide data and insights to Customers and Partners for them to sell products and services and manage their processes.

A real-time Cloud Platform delivers low-latency responses from data captured or computing Servers that reside close to the Core of the Internet and enables operations to add AI and ML models, either as services or for Auto ML, to the business processes while reducing the cost of Data Engineering and serving by being closer to operations. Edge services of the Cloud Platform offer Infrastructure-as-a-Service for easy provisioning of new services closer to operations while End-Users and Partners benefit from a simplified Business-Process-as-a-Service interface to configure the needed share of the services offered through the Platform. At the same time, Demand Forecasters and other MLOps services facilitate a Customer-driven offer of new services for all Partners of a sales and supply process chain to carry out their processes with tailored Data and Insights services.

### 4.7.1. Final Thoughts and Future Directions

Intelligent operational systems strive to automate or support operational tasks in businesses and provide services to users in near real time. The operations are usually managed by management information systems, enterprise resource planning systems, or other systems residing in the enterprise datacentre. The Intelligent operational system intelligent operational systems use ML models to make predictions and also include servers used for hosting applications. Such operational systems can be deployed and managed in the cloud, thus taking advantage of the elasticity provided by cloud platforms. These systems can browse and search for cloud resources and services, use IaaS and PaaS offerings, adopt a serverless architecture, and handle unpredictable and variable workloads.

The relationship between Clouds and intelligent operational systems can be studied in multiple ways. First, Clouds provide IaaS and PaaS to host operations. Serverless technologies, such as Function as a Service and event-driven architectures, can enable the automation of operational decision-making. Data management and configuration storage can support the high frequency of operational data ingestion and data dependency. Furthermore, the natural integration of analytics and AI/ML services in Clouds enhances the reliability of the Cloud services and fosters the automation of operational task execution.

## References


Buyya, R., Dustdar, S., Ranjan, R., et al. (2021). Cloud computing: A distributed internet computing paradigm for the future. Future Generation Computer Systems, 29(3), 599–616.

Vadisetty, R., Polamarasetti, A., Rongali, S. K., kumar Prajapati, S., & Butani, J. B. (2025, May). Blockchain and Generative AI for Cloud Security: Ensuring Integrity and Transparency in Cloud Transactions. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-6). IEEE.

Bai, T., Zheng, Z., Ren, K., & Shi, S. (2024). Cloud-native machine learning systems. IEEE Software, 41(1), 50–58.

Guntupalli, R. (2025, August). 5G and AI-Powered Cloud Security: Safeguarding Ultra-Low Latency Networks. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-4). IEEE.

Newman, S. (2021). Building microservices (2nd ed.). O'Reilly Media.

Agentic AI in Data Pipelines: Self OptimizingSystems for Continuous Data Quality, Performance,and Governance. (2024). American Data Science Journal for Advanced Computations (ADSJAC) ISSN: 3067-4166, 2(1). https://adsjac.com/index.php/adsjac/article/view/23

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems, 33, 6840–6851.

Varri, D. B. S. (2024). Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems. Available at SSRN 5774785.

Lebcir, I., Shah, C. A., Nagubandi, A. R., Dhoke, S. M., sikh, G. S. & Mishra, M. K. (2025). FinTech and Financial Inclusion in Emerging Economies: An Empirical Assessment. Advances in Consumer Research, 2(6), 2005-2011.

Kingma, D. P., & Welling, M. (2020). An introduction to variational autoencoders. Foundations and Trends in Machine Learning, 12(4), 307–392.

Keerthi Amistapuram. (2024). Federated Learning for Cross-Carrier Insurance Fraud Detection: Secure Multi-Institutional Collaboration. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 6727–6738. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/3934

Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. Information Sciences, 479, 448–455.

Rani, P. R. S., Kummari, D. N., Yellanki, S. K., Meda, R., Reddy Koppolu, H. K., & Inala, R. (2025). Blockchain and AI for Securing Electrical Infrastructure. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1–6). IEEE. 2025 2nd International Conference on Computing and Data Science (ICCDS). https://doi.org/10.1109/iccds64403.2025.11209487

Newman, S. (2021). Building microservices (2nd ed.). O'Reilly Media.

Davuluri, P. S. L. N. (2021). Event-Driven Compliance Systems: Modernizing Financial Crime Detection Without Machine Intelligence. Journal of International Crisis and Risk Communication Research , 339–354. https://doi.org/10.63278/jicrcr.vi.3636

Armbrust, M., Xin, R. S., Lian, C., et al. (2020). Delta Lake: High-performance ACID table storage. Proceedings of the VLDB Endowment, 13(12), 3411–3424.

Avinash Reddy Segireddy. (2022). Terraform and Ansible in Building Resilient Cloud-Native Payment Architectures. International Journal of Intelligent Systems and Applications in Engineering, 10(3s), 444–455. Retrieved from https://www.ijisae.org/index.php/IJISAE/article/view/7905

Garapati, R. S. (2023). Optimizing Energy Consumption in Smart Build-ings Through Web-Integrated AI and Cloud-Driven Control Systems.

Blier-Wong, C., Cossette, H., & Marceau, E. (2021). Tree-based machine learning models for insurance ratemaking. North American Actuarial Journal, 25(3), 383–407.

Charpentier, A., Denuit, M., & Trufin, J. (2021). Explainable machine learning in insurance pricing. Scandinavian Actuarial Journal, 2021(7), 565–594.

Aitha, A. R. (2023). CloudBased Microservices Architecture for Seamless Insurance Policy Administration. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 607-632.

Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., & Stoica, I. (2020). Spark: Cluster computing with working sets. Communications of the ACM, 59(11), 56–65.

Kreps, J. (2021). I heart logs. O'Reilly Media.

Deep Learning-Driven Optimization of ISO 20022 Protocol Stacks for Secure Cross-Border Messaging. (2024). MSW Management Journal, 34(2), 1545-1554.

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 17(3), e0265480.