



Generative AI for Order Management in Edge and Hybrid Cloud Systems

Luca Bianchi

Asst Professor

line 1: 1st Given Name Surname
line 2: *dept. name of organization*
(of Affiliation)
line 3: *name of organization*
(of Affiliation)
line 4: City, Country
line 5: email address or ORCID

Abstract—Next-generation order management with generative AI is a timely concept aligned with ongoing developments in edge computing and hybrid cloud data centers. Most applications contemplate online inference of generative neural network models without strict requirements for low latency. Workloads that demand fast, real-time responses and support distributed execution of edge services, however, are encountering challenges with respective latency-sensitive applications. Reliable, low-delay order management for such settings can greatly benefit from generative AI. The active exploration of Generative AI methods for operational decision automation, particularly in areas such as order scheduling and work order forecasting, may enhance responsiveness and shorten time-to-decision.

Implementation or use-case analyses often remain sketchy or even absent. Performance considerations are commonly limited to standard generative approaches, namely the ability of generative methods to provide accurate outputs. Generic Latency metrics, however, are equally important. Although generative methods are frequently described as costly relative to their discriminative counterparts, this argument merits careful scrutiny in decision-automation scenarios, especially when local deployment at the edge is considered.

Keywords—Generative AI in Order Management, Edge AI Systems, Hybrid Cloud Architectures, Real-Time Order Processing, Low-Latency AI Systems, Distributed Edge Computing, Order Scheduling Optimization, Work Order Forecasting, AI Decision Automation, Latency-Aware AI Models, Edge Deployment Strategies, Operational AI Systems, Real-Time Inference, AI Performance Metrics, Cost-Efficiency in AI, Discriminative vs Generative Models, Time-to-Decision

Optimization, Scalable Edge AI, Intelligent Order Systems, AI-Driven Operations.

I. INTRODUCTION

Generative AI is enabling novel applications in multiple industries, including automotive, consumer electronics, and life sciences. Research in the next few years will look beyond conventional workloads to demonstrate how this class of models could revolutionize business operations by transforming order management into a fully automated, fraud-resistant process. Generative methods, rules-based dialogue systems, and decision support services could simulate patterns and knowledge, enabling task automation across the order lifecycle. Industries can prepare for adoption of these concepts by considering the scope of generative AI, its applications, associated risk and ethical considerations, performance metrics, and a roadmap for implementation.

Advances in AI and immersive technologies will support novel applications like real-time personal assistants with natural language and tone de-biasing, chatbots that anticipate user queries before being explicitly prompted, and order-management systems capable of handling intricate, cross-organizational transactions—or even simulating such transactions entirely. In the latter application domain, generative AI could analyze high-velocity supply chains with real-time product refresh rates. The core underlying idea is to leverage experimental patterns—accelerated by generative methods—for non-player characters in video games. Furthermore, real-time edge computing at the network periphery could enable AI models to make low-latency

inferences relevant to mission-critical processes like order management.

A. Overview of Generative AI in Order Management

Generative AI influences multiple aspects of order management including order classification, forecasting, orchestration, fulfilment, fulfilment assignment, fulfilment scheduling, and anomaly detection. Generative AI models provide contextualized responses, predicting missing parts of data while maintaining semantic coherence. Decision automation populates digital control towers to reduce cognitive workload and facilitate efficient decision-making. Generative AI for order management spans order-to-cash, procure-to-pay, and record-to-report workflows in industrial supply chains. Industrial supply chains adopt order management in a multi-entity, multi-source-to-multi-sink scenario.

Generative AI improves applications sensitive to latency and jitter, particularly in retail and logistics, where tasks executed by a large number of low-cost devices entail stringent service-level agreements. In these domains, latency-sensitive workflows require high-performing solutions that leverage proximity to the action. Traffic management, warehouse robotics, and last-mile delivery fall within this category.

Table 1 — Core variables for mathematical modeling

Symbol	Meaning	Unit
λ	incoming order/request rate	orders/s
D	average source-to-compute distance	km
L	total decision latency	ms
L_n	network latency	ms
L_i	inference latency	ms
L_o	orchestration/control latency	ms
L_q	queueing / acceptance delay	ms
C	total execution cost	monetary unit/order
U	service utility / business value score	normalized
A_f	forecast accuracy	0 to 1
S	scheduling quality / fulfillment score	0 to 1
R	robustness score	0 to 1
F	fairness score	0 to 1
P_{SLA}	SLA violation probability	0 to 1
x_e, x_h, x_c	deployment decision variables for edge, hybrid, cloud	0 or 1

II. BACKGROUND AND MOTIVATION

Recent advances in generative AI are opening new avenues for order management in data centers. By automating demand forecasting, real-time workload management, and resource-scheduling decisions, these models promise to

enhance performance. Tight latency constraints motivate their deployment in edge and hybrid-cloud environments. Generative AI augments edge-to-cloud orchestration engines that govern order-driven provisioning of hardware and software components for large-scale services. Within defined SLAs, reactive forecasting models—trained on historical demand patterns and additional contextual data—automatically anticipate short-term requirements. Forecasts trigger resourcing decisions executed by primary edges or private clouds. As an overlay solution, generative AI alleviates latency-sensitive workloads from main orchestration paths, improving the responsiveness of order management processes.




Fig 1: Edge-To-Cloud Order Orchestration With Generative Ai

Generative AI methods are actively being applied for demand forecasting in supply chains, automated scheduling and management of cloud data centers, and decision-making regarding real-time edge workloads. The applicability of these techniques must be further explored in other time-constrained areas within the order lifecycle of a hybrid-cloud model, as well as for other types of data predictions. Workload management in data centers introduces tight latency requirements. Consequently, the nature of the deployment and the technology choice used to meet latency targets must be carefully considered.

A. Context and Rationale for Integrating Generative AI in Order Management

The accelerating expansion of edge computing and hybrid-cloud data centers fosters the adoption of Generative-AI-enabling models. However, latency-sensitive workflows such as order management could greatly benefit from such advancements, yet remain excluded from related investigations. Generative AI enhances order management by streamlining the end-to-end process of recording, processing, and fulfilling an order within an industrial supply chain. By performing operational tasks such as demand forecasting, order scheduling, resource allocation, and logistics routing, Generative AI reduces both the human effort required and the overall response time for service requests. The integration of Generative AI into order management improves the breadth of its coverage and responsiveness, allowing for round-the-

clock automated operation, and leveraging the data generated to enhance the quality of service. At the same time, the overall architecture embraces explainability and auditability to address the current black-box concerns surrounding Generative AI.

From responsible AI models to regulations around GDPR, a comprehensive set of guidelines governs the design, development, and deployment of AI models across industries. However, responsible AI and transparency take on greater significance in edge computing and hybrid-cloud data centers due to the wider ecosystem of participants involved and the shared nature of data across the cloud. A well-defined data governance model aids in addressing privacy, data ownership, guidelines for data protection, and detailing how compliance is maintained by organizations dealing with customer data.

EQUATION 1 — TOTAL DECISION LATENCY

Step 1: identify all latency components

A decision request experiences:

- network transit delay L_n ,
- model inference delay L_i ,
- orchestration/control delay L_o ,
- queueing/acceptance delay L_q .

Step 2: add them

$$L = L_n + L_i + L_o + L_q$$

Step 3: interpretation

If the system is edge-local, L_n is small.
If the system is cloud-centric, L_n grows.
If orchestration is complex, L_o grows.

Final form

$$L = L_n + L_i + L_o + L_q$$

III. METHODOLOGY

Demonstration of the proposed concepts is positioned as an applied research contribution, with the following elements providing a foundation to the real-world evaluation. Generative AI algorithms are used for the creation of key inputs and control attributes. Further, the management of task scheduling is orchestrated by a generative model. The generative AI methods are supervised through an autonomous decision-making capability that determines the delivery mode (compute at edge or cloud) as well as the orchestration deployment that can be either static edge cloud pagination, hybrid pagination, or drop off. These orchestrated decisions are integrated with an AI orchestration controller that supervises the complete decision-making workflow.


Data from the corresponding workflow indicates the latency constraints on the decision-making process and parameters associated with the decision-making. With respect to the elements within the decision-making, these include the reproachfulness of the data model that can

provide true data, the explainability of the page delivery which can help in understanding the model decision, the final acceptance delay which defines the real acceptance time of the information and ultimately these delays if orchestrated can have a reduction effect on the overall acceptance of data at the delivery edges.

A. Generative Approaches and Decision Automation Mechanics

Machine learning (ML) has become a mainstream technology for software-intensive systems. Nevertheless, the operational excellence of these systems has become pivotal to both retention and efficiency, and therefore stands in sharp contrast to the existing unfocused automation solutions based on ML approaches. A next-generation architecture and methodology responding to such needs for responsive Order Management in Edge Computing and Hybrid Cloud Data Centers have thus been presented and subsequently demonstrated using several Generative-AI and Orchestration components, as indicated. These solutions can provide human-like responsiveness, high-throughput automation, low latency, low service-cost, and operate continuously in a voice- or text-based environment.

ML methods can be exploited to automate decision-making in production, formulation, and therefore execution of all or part of the business processes of IT and other enterprise resources represented in the Info-Communication Technology System. The decisions made can by themselves be based on one or multiple ML modules and can be enabled within a Decision Control Loop that configuration can map onto a Decision Engine in an Orchestration solution for enterprise resource management and business operations. Next-generation products are capable of implementing ML methods as Generative-AI applications.



IV. OBJECTIVE OF THE STUDY

This study investigates the integration of generative AI in order management, elucidating technology principles, delineating workflows, discussing business outcomes, and addressing operational excellence and fairness. Three questions frame the exploration: How can order life cycle management be reimagined with generative AI? What new capabilities over the end-to-end process can this technology provide to the business? How does the adoption improve operational excellence and resiliency? Users operate in edge computing environments, enabling real-time automation of customer actions, deliveries, and stock replenishment while managing order fulfillment in hybrid cloud mode.

Order life cycle planning enables organizations to orchestrate business changes, operations, and service

delivery initiatives to sustain technology and services over time. These organizational workflows are complex, and workflow exceptions must occur on time, with attention to possible fines. However, order life cycle management usually focuses only on demand forecasting, neglecting all the other aspects that guarantee fulfillment on time and in full. Automating every phase of order life cycle management can create a strong competitive advantage, making it less dependent on human intervention and more responsive. Generative AI can play a key role in achieving such automation, bringing together deep learning, reinforcement learning, and neural symbolic AI in the same solution.

A. Study Aims and Research Questions

Business operations increasingly emphasize customer experience while keeping costs in check. Generating \$1.29 trillion in US revenue in 2023 alone, online retail is burgeoning, with the rapid expansion of digital marketplaces. Meeting these demands while optimizing costs has become exceedingly complex. Generative AI's understanding and creation capabilities present an opportunity to relieve the burden by automating latent decision-making processes used for handling online orders. Automating these workflows while sensing changes in demand and supply allows reallocation of minimally engaged human resources from maintaining routine orders to handling nonstandard requirements. However, generative AI hinges on quality data design, management, orchestration, and governance. In an era of data deluge, effective data operations reduce latency and preempt a data swamp.

Services are already building pain and gain approaches into their order management processes by integrating AI solutions into the end-to-end order lifecycle, making it possible to process more orders with fewer people. Beyond developing self-healing operational workflows, what's needed is a holistic edge-to-cloud order management system with the probative capabilities of generative AI. Central to such systems are strong latency, fairness, performance, and balancing measures that ensure timely intervention by resources, safeguard customer service levels, lean supply processes, and minimize diverted goods. In the age of everything as a service, a seamless trust-enabled ecosystem with inter-business order management order-handling capability is imperative.




Fig 2: Generative AI in Digital Marketplaces: Automating Latent Order Decisions and Streamlining Logistics

V. RESEARCH SUMMARY

Core findings suggest that, while exploratory, generative approaches are viable candidates for certain Order Management components. Responsiveness and explainability, critical for business impact, are inherently addressed in generative techniques. Empirical performance indicators such as accuracy and latency may expose limitations in other capabilities, such as model reliability and risk-resilience; however, these are primarily controlled by back-end orchestration, not generative components. Consequently, users should prioritize responsiveness and explainability when evaluating performance relative to factory-trained solutions. Sinking latency targets also creates space for Latency-sensitive Generative AI models to match or supersede factory-trained quality. These first principles guide and frame implementation scenarios; Industrial Supply Chain Applications and Retail Logistics Use Cases employing the overall architecture are explored subsequently. Well-defined modeling approaches in support of a specific decision class further enhance de-risking: when a Latency-sensitive Generative AI model is engineered to discover the optimal Resource Planning Schedule, all non-responsive edges become confidence-based decisions.

Key contributions weave Generative AI techniques into Order Management across Edge Computing and Hybrid Cloud Data Centres, demonstrating suitability for low-tolerance-latency Workflow Segments. As with all business-critical processes, generating definitive Decision Automation capabilities remains a key requirement; here, focus is on Logical Templates, Deployment Dependencies, and DataNation to establish and populate the Edge GUI enabling Generative AI interaction, ensuring transparency and Auditability across key transactions.

EQUATION 2 — NETWORK LATENCY AS A FUNCTION OF DISTANCE

Step 1: assume latency grows with distance

Let D be compute distance in km.

Step 2: include base switching/routing overhead

Even at zero distance, some base delay exists:

$$L_n = L_{n0} + kD$$

where:

- L_{n0} = fixed network overhead,
- k = latency increase per km.

Step 3: deployment-specific versions

For edge:

$$L_n^{(e)} = L_{n0} + kD_e$$

For hybrid:

$$L_n^{(h)} = L_{n0} + kD_h$$

For cloud:

$$L_n^{(c)} = L_{n0} + kD_c$$

with

$$D_e < D_h < D_c$$

Final form

$$L_n = L_{n0} + kD$$

A. Key Principles of Generative AI in Order Management

Generative AI enhances order management in edge and hybrid environments by enabling responsive, explainable, reliable, and cost-efficient decision-making, as well as assisting human experts. Responsiveness to business opportunities and threats is essential for success. Enabling AI-assisted decision automation for new order fulfilment and exception handling can significantly increase order throughput compared to conventional approaches. Generative AI helps control the quality of automated decisions by offering a clear model of the decision-making process, providing human operators and customers with explanations. Low-cost execution of order fulfilment activities is a prerequisite for business success. Generative AI enables efficient outsourcing of complex order fulfilment activities to the cloud and helps decide where to process orders with minimal latency. Latency constraints for some processes, such as fulfilment planning and exception handling, demand specialised techniques. Generative AI allows control of these latency-sensitive functions and enables rapid response during critical periods. Generative AI is not designed to replace human operators but to assist their agency.

Systematic guidance for adopting generative AI in order management is lacking. The overview addresses this gap by defining three essential principles and applying them to the context of edge computing and hybrid cloud data centres. Generative AI models are trained to provide fast responses associated with low latency for the most critical order-management decision-making processes. This is achieved by a distributed generative AI architecture capable of executing ready-to-run solutions close to the action site while outsourcing to the cloud those decisions with limited impact on order-path execution time.

VI. CORE CONCEPTS IN GENERATIVE AI FOR ORDER MANAGEMENT

Generative models and decision automation define the essential mechanisms of next-generation order management. Handling the order life cycle initiates management decisions that span multiple processes with interdependencies and lead to interactions among diverse stakeholders. Generative models incorporate data from the overall life cycle to predict future states and events at any point along the trajectory and support subsequent activities. Consequently, these models respond quickly to triggers and execute decisions at any level in the automation framework. A responsive and explainable generative-model-based policy operates embedded in orchestration engines to deliver decisions at true edge scale. A responsive integrative approach embraces generative models to direct orchestration engines to schedule and

perform the next phase of the life cycle and complete the transition to the next state.

Managing latency-sensitive workflows in edge computing transforms distributed-edge resources into locally available private clouds, shifting these workflows to the edge for execution. Decision latency requirements govern the distribution of generative models: Very tight constraints warrant local edge execution, responsive models reside close to the sources feeding their training data, and other models, although requiring greater latency, can operate in the cloud.

A. Generative models and decision automation

Generative models produce new instances of complex objects, typically based on labelled training sets. Unlike other machine learning models, the outputs may not always be representative of the training set; for example, a generative speech model can produce outputs that sound like the person whose voice was modelled, without using that person's voice data for training. Generative models in order management can be used to produce forecasting and orchestration events. In the industry supply chain domain, a generative forecasting model may generate part demand forecasts for each part and location, while a generative scheduling model may produce schedules for all main assembly facilities on factory holiday. Neural models for forecasting or scheduling are sufficiently flexible and capable that useful results can be produced, including in cases where traditional approaches do not yield satisfactory results. At the same time, these are not the AI equivalent of a self-driving car – while a self-driving car may be as simple as a neural network with successful training data (if the parts work as intended), generative forecasts or schedules should be subject to a decision-loop control mechanism to ensure responsiveness and reliability.

The decision-loops can take various forms. Generative AI approaches can integrate with an existing order management orchestration engine so that a forecast or schedule is only used when later processes agree that it is useful. Alternatively, the generative AI modules can be integrated with a local decision sub-loop that executes discrete deployable functions, such as producing a set of fulfilment orders for delivery to a retailer or fulfilling a retail replenishment request. In this form of integration, use of the generative AI model is governed not only by product definition and software options, but also by fulfilment loads, lineage and provenance of the deployed decision-support models, and any practical supply chain constraints – including risk exposure appetite frameworks being developed for the supply chain.

Table 2 — Architecture components

Component	Role in article	Quantitative interpretation
Edge computing framework	local data generation, monitoring, real-time inference	minimizes L_n
Data orchestration and governance engine	metadata, lineage, compliance	reduces risk, improves traceability

Component	Role in article	Quantitative interpretation
Generative forecasting model	predicts demand / events	increases A_f
Decision automation engine	order validation, resource scheduling, fulfillment orchestration	improves S , reduces L_o
Human-in-the-loop oversight	approval and exception handling	controls fairness/risk
Hybrid cloud deployment	balances responsiveness and scalability	trades off L and C

B. Edge computing implications for latency-sensitive workflows

Latency-sensitive workflows, particularly those involving data processing at the edge, demand responsive and lightweight decision automation. Management by generative models, rather than predictive models, presents multiple advantages. Predictions at the edge require immediate responsiveness, which can be challenging given the long instantiation times associated with prediction serving. Moreover, many edge predictions are short-lived and thus prohibitive to host at the edge. Even for predictions that are responsive and long-lived, keeping them in-memory fashion in location-aware queues can introduce cache alignment problems.

Most edge management predictions are prone to skew, especially if funding or consumption patterns shift. Skew resilience models improve at absorbing shifts, but predicted distributions tend to be too much in the seat of the pants for unforced latencies—e.g., for decentralized internally-forced load-balancing-like scheduling of localized requests (like those for media or games) for which management hopes to maintain rapid response servicing times even though such internally-forced latencies are longer than the source or sink of a forced media). Even for centralized unsupervised long-term unsupervised skew short bursts, it is difficult to impossible to predict safe servicing latencies in real-time.

Generative approaches to distribution induction resolve such major limitations of predictive-based approaches. Using distribution models. For each such long tail of suddenly misguided and shifted requests, the modelling would naturally require a quick-response brain and be relatively fast too, but WAN transit to and from a centralized model seem very uneconomical.




Fig 3: Generative Model Management for Edge Decision Automation

VII. SYSTEM ARCHITECTURE AND REFERENCE FRAMEWORK

A blueprint for order management systems that span edge and cloud data centers is described. It encompasses all relevant components, with accompanying information flows and interaction interfaces. Data orchestration is addressed through a metadata model that governs data provenance, history-tracking, access control, and compliance with legal and regulatory requirements.

The overall architecture for end-to-end order management spans edge-to-cloud environments and involves all participating actors. Generative models run at the edge or cloud as suitable, and domain-specific data governs control decisions throughout the order cycle. Data orchestration, management, and governance are covered in complementary sections. These depend on data lineage and the regulatory context, while the adopted metadata model provides the necessary coverage.

A. Overall architecture for edge-to-cloud order management

To put next-generation order management into practice, a dedicated solution architecture spanning the order lifecycle is essential. By using a data orchestration engine, the capabilities of generative models and a decision automation engine—for demand forecasting, resource scheduling, and application orchestration—are fully integrated into a comprehensive solution. The overall architecture consists of four main building blocks. An edge computing framework handles data generation, monitoring, and real-time inference at the edge. A central data orchestration and governance engine manages metadata and data lineage for production and reporting. A controllable generative model uses orders, inventory, and other contextual data to automatically predict future product demand, and a decision automation engine creates the order management context and continuously adapts it by orchestrating different supporting applications.

Data flow between the architecture components follows the four phases of Infrastructure orchestration, Decision preparation, Order execution, and Decision monitoring. The corresponding sets of interfaces control cross-component

interactions. User-defined data governance policies ensure compliance for data being read, written, or moved between components.

EQUATION 3 — END-TO-END ORDER-MANAGEMENT OBJECTIVE FUNCTION

- low latency,
- high forecast quality,
- high scheduling quality,
- low cost,
- robustness/fairness.

So construct a weighted objective.

Step 1: define positive and negative terms

We want:

- maximize A_f forecast accuracy,
- maximize S scheduling quality,
- maximize R robustness,
- maximize F fairness,
- minimize L ,
- minimize C .

Step 2: convert to one scalar score

$$J = w_1 A_f + w_2 S + w_3 R + w_4 F - w_5 L - w_6 C$$

Step 3: explain weights

The w_i are business priorities:

- larger w_5 means latency matters more,
- larger w_6 means cost matters more,
- larger w_3, w_4 means governance matters more.

Final form

$$J = w_1 A_f + w_2 S + w_3 R + w_4 F - w_5 L - w_6 C$$

and the optimization is

$$\max J$$

B. Data orchestration, governance, and lineage

Data orchestration, governance, and lineage together form an integrated model for managing metadata and controlling data throughout its lifecycle. Orchestration relies on data provenance and business-domain context to ensure the right data are available at the right time and place, while compliance policies—addressing regulatory, audit, privacy, and security concerns—regulate access. Strategies governing


data ownership and sharing facilitate regulated collaboration. A combination of automated and manual processes ensures that appropriate checks are performed without introducing excessive overhead. The resulting model supports data access throughout the organization, generating business value while minimizing risk.

Data provenance encompasses details of a data set's origins and constituents, information on any transformations performed during its lifecycle (that is, during data pipeline execution), and records of any subsequent access. Data pipelines capitalize on such provenance metadata by monitoring data quality and trustworthiness, enabling data selection based on varying levels of accuracy, freshness, and other quality dimensions. Thus, trustworthiness and quality can be seen as expressions of provenance coupled to application-specific criteria. In addition to improving data selection quality, a lineage-aware data-pipeline architecture facilitates support for complex metadata queries (such as "Which user accessed sensitive-personal data yesterday?") and helps answer questions pertinent to data sharing and data-related regulations such as GDPR and HIPAA.

VIII. ALGORITHMS AND TECHNOLOGY STACK

Generative AI methods for forecasting and orchestration include algorithms to forecast demand at appropriate granularity, templates and policies to schedule resources, and rules to automate decisions on allocating cloud and edge resources, service workloads, and accept orders, among others. A technology stack and deployment pipelines address real-time inference requirements at the edge.

Three interrelated goals drive the integration of Generative AI with Edge Computing and Hybrid Cloud Data Centers. Generative AI-based forecasting decomposes production and service demand signals at different spatial and temporal scales, strengthens the orchestration engine by scheduling resources dynamically and on shorter time horizons, and instantiates decision policies that automate order acceptance, workload allocation, and cloud-edge resource assignment. Together, these contributions align the solution with responsive design principles, reduce service lead times and overall system cost, and facilitate real-time inference on sources of latency-sensitive workloads at the edge.



A. Generative AI methods for forecasting and orchestration

Generative AI methods for forecasting and orchestration: Forecasting leverages generative models for sequence prediction, harnessing historical series correlations and exogenous data derived from synchronous campaigns,

weather patterns, and local events. These temporally relevant variables enhance long-term predictions, while shorter-term models capture short-lived effects. Two control loops apply: a slow-burning loop directs inventory-level orchestration, while a fast loop governs decision scheduling.

Generative AI methods for forecasting and orchestration: The technology stack employs Prophets for sales forecast generation and an orchestration engine for order validation and fulfillment scheduling, utilizing either a centralized or edge-local orchestration layer. Decision Automation, depicted in Figure 2, periodically proposes orders based on market forecasts and established rules. LLMs analyze order-related data, dynamically suggesting changes for human approval when needed.

B. Real-time inference at the edge

Latency-sensitive tasks in customer-order management must rely on real-time or near-real-time data associated with multiple and specifically orchestrated flows of events in multiple sources. These real-time activations require the state-of-the-art AI/ML models to work on the closest edge of the data generation sources—most of the cases, on the same site or the same metro area (within a limited defined distance, site, or metro location)—to achieve the best-in-class low-latency response.

As constraints on performance are very strict, utilizing heavy, high-performance, and resource-consuming state-of-the-art AI/ML models is not advisable. When such models must be executed for latency-sensitive inference, their deployment on the closest datacenter sites (DC) must ensure optimization and efficiency. Metrics for continuous monitoring of inference performance must be established to guarantee the stakeholders a commodities-level quality on these services and the required continuous business alignment and non-disturbance.

To achieve the best score on inference latency and therefore continuous generation of real-time triggers to Latency-sensitive automation workflows, some solution architecture ideas can be followed. One of these is a dedicated common framework and pipeline with a clear decision on which model to trigger based upon their performance measured by the defined metrics. To explore this approach, it must be based not only on the latency achieved but also aligned with model cost associated with the total time linked to the generation of the model data used in the inference.




Fig 4: Model Inference Cost-Effectiveness In Distributed Edge Ai

IX. OPERATIONAL EXCELLENCE AND PROCESS IMPROVEMENT

Business processes span the entire company, beginning with order creation and ending only when all associated actions prescribed by operations support are completed. Understanding this end-to-end flow allows for the identification of potential areas for operational excellence and continuous process improvement. The progression of orders from creation, usually by some front office or supply chain orchestration function, through fulfillment, to post-factum closure generates a multitude of handoffs between different organizational entities. Each of these handoffs represents a point where the order custodians have to wait for the next step of the process to take place, and ultimately it is the duration of all these waits that determines the total order lead time.

Consequently, the order lifecycle process is a prime candidate for analysis from an operational excellence and process improvement perspective. Within this context, operational forecasting represents a direct avenue of delivering operational excellence, by providing these order custodians with a more accurate view of future demand and requirements. Generative AI can also deliver improvement by offering adaptive decision-timers to reduce the materialized latency overhead at these handoffs. At the same time, it is crucial to design these time-sensitive decision-timers in such a way that they do not compromise the quality of service during periods of lower demand. A sound fail-safe logic is thus needed for any adaptive decision-timer so that its output can be relied upon even in the face of model errors or untested conditions.

A. End-to-end order lifecycle management

An order plays a crucial role in an enterprise environment, as it connects customers requesting a product/service with suppliers that can provide it. Sequencing and coordinating orders among multiple customers and suppliers can significantly improve operational efficiency by minimizing order delays. Time-sensitive orders (e.g., in-flight aircraft repairs) introduce an additional layer of difficulty, as these

orders not only require coordination but must also be executed quickly. Companies need to manage the end-to-end lifecycle of these orders efficiently to achieve operational excellence and process improvements. Generative methods, which have gained popularity for their ability to learn from and synthesize data, can automate multiple key decision-making stages in the order lifecycle. Businesses benefit from incorporating generative methods into order lifecycle management through improvements to order processing times, latency, and resource allocation.

The order-lifecycle comprises the stages of order request, fulfillment, and completion, along with the handoffs and communication between each stage (see Figure 1). Each of these stages is a good candidate for automation with generative methods, especially those requiring batch processing of requests from multiple customers or suppliers. Further aspects, such as managing exceptions and disruptions, may also benefit from the inclusion of generative methods.

EQUATION 4 — FORECASTING LOSS FUNCTION

Let:

- y_t = actual demand at time t ,
- $\hat{y}_t$ = forecasted demand.

Step 1: define error

$$e_t = y_t - \hat{y}_t$$

Step 2: square the error

$$e_t^2 = (y_t - \hat{y}_t)^2$$

Step 3: average across T time steps

$$\text{MSE} = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2$$

Step 4: optimization target

Training seeks

$$\min \text{MSE}$$

Final form

$$\boxed{\text{MSE} = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2}$$

B. exception handling and resilience

Exception handling encompasses all activities undertaken in managing events that disrupt the regular execution of a process and deviate from its expected course. Transitions between different process states typically introduce constraints that ensure a successful, fault-tolerant completion. When exceptions occur, corrective measures are initiated to recover and continue the desired process flow. An exception is said to be handled if it is either prevented or, once sensed, rectified and recovered from; the handle is the code that executes the deal with the exception. Two default handling strategies are fail-as-fast-as-possible and complete-

or-compensate. In the first approach, execution halts immediately when the occurrence of an exception is detected; if recovery occurs, system execution continues, but in a new process instance instead of the old one. In the second approach, execution continues until one of the states that should be reached for process success is reached, at which point process catastrophe will either happen or an implicit reference state will be executed.

Resilience entails preventing or mitigating failures, allowing the core logic to respond correctly when exceptions occur. In real-life processes, resilience is handled by the human actors through procedures that have been defined informally or ignored in practice. In any process, the speed of fault detection and recovery is crucial. The important attributes are: fault tolerance, the capability to continue to operate in the presence of faults (human or environmental) without interruption in the functional, non-functional, or performance capabilities of the service delivered; fast failure, performing failure recovery as quickly as possible; and preservation of control, maintaining the ability to see all that is happening and intervening appropriately.

X. EVALUATION, METRICS, AND VALIDATION

Success metrics for the AI-driven generative order management solution are defined. Focus extends to measuring the system's capability to support efficient decision-making across the complete end-to-end workflow, including exception handling. Order management generates a tailored fulfillment plan based on system input regarding demand, supply, and transportation networks. Real-time end-to-end orchestration then seeks to execute the fulfillment plan with optimal efficiency while ensuring excellent service delivery. Key metrics for the fulfillment plan generation stage factor in autonomy, speed, and accuracy. Robustness and fairness metrics gauge order management resilience, with emphasis on supply chain ethics and bias reduction.

An order management solution capable of real-time decision-making does not relieve humans of the responsibility for decisions exceeding the bounds of acceptable failure but should ideally reduce the frequency of such decisions. Generative order management encompasses all steps from order initiation to fulfillment plan execution, with steps often conducted in parallel. All exception cases – such as unavailability and cost limits for transport or service resource allocation – and the decision-making surrounding them are therefore in scope. Metrics thus need to include performance across the entire exception handling process. Robustness and fairness dimensions also demand attention, especially for diversity and inclusion in data-driven learning.

A. Performance metrics for generative order management

Performance metrics for generative order management should align with key business goals. For example, in a discrete manufacturing scenario, it may be essential to minimize backorder and stockout situations, while also striving to reduce total inventory holding costs. An appropriate objective would therefore be to minimize a weighted sum of these two quantities. Similarly, for a measurable retail supply chain, an objective function for order management may weigh demand forecast accuracy more heavily than revenue potential from reducing inventory stockouts. Also, in fund or investment management, an important aspect of performance is total transaction costs

associated with security trades, while risk is often quantified in terms of large deviations from the expected result. The objective functions could therefore reflect the minimization of expected total transaction costs subject to suitable bounds on risk.

In integrated generative order management, performance assessment becomes more complex, as relevant decision parameters are often determined by a hierarchy of models. It thus becomes necessary to evaluate the robustness of the holistic effect of these parameters. Generative models that support the critical decisions can be executed using the appropriate data so as to provide a real-time recommendation platform to enable next-best-action guidance, including forecasting of near-term scenarios considering end-to-end supply chain considerations. Further, in context-aware generative orchestration, specific supply-demand scenarios and configurations can trigger orchestration between the supply chain layers and spur reconfiguration of deep supply-demand matching of market products and services.

Table 3 — Typical optimization goals implied


Objective	Mathematical target	Why it matters
Minimize latency	minimize L	real-time order handling
Maximize forecast quality	maximize A_f	better demand planning
Maximize scheduling quality	maximize S	timely fulfillment
Minimize cost	minimize C	efficient operation
Minimize SLA violation risk	minimize P_{SLA}	service-level compliance
Maximize fairness and robustness	maximize F, R	responsible AI deployment

B. Robustness, fairness, and bias considerations

Generative AI models play a considerable role in managing business-to-business supply chain processes through automated yet robust and fair decision-making. However, business-to-consumer retail and logistics sectors can benefit from more tailored offerings, such as via chatbots and journey mapping. Many approaches in these segments are inherently imbalanced, as low failure rates imply imbalanced class distribution. Any condition with the potential to cause order failure requires a robust and resilient process capable of handling such exceptions. Ensuring the process is resilient to business-as-usual failure conditions provides the basis for safety and well-managed resources, allowing for support from real-time machine learning.

Fairness, risk, and bias are also important considerations in Generative AI, especially when dealing directly with customers via chatbots that can learn and evolve. Discipline and governance are paramount to ensure responsible AI, which is achieved in business-to-business supply chains by keeping human-in-the-loop decision-making in the forecast and orchestration process. Nevertheless, the learnings from Generative AI in the supply chain can be extended to retail by considering the nuances of fairness in the chatbots implementing business-to-consumer supply chain processes.

The preparation of the chatbots must consider journey mapping to avoid bias, and all interactions via the chatbots must be recorded with lineage tracing for later analysis and learning for future versions.



XI. IMPLEMENTATION SCENARIOS AND CASE STUDIES

Deployment scenarios can help verify capability and performance within a particular domain. Two examples describe decision-making authority and flow within an industrial supply chain and its supporting ecosystem. Advanced storage location forecasting serves a warehouse management system, while a data center forecasting-ordering-scheduling pipeline represents IT and services domains.

Edge computing for retail encompasses locations with high-volume and low-margin transactions. Although sales order processing, demand forecasting, and automated replenishment strategies are widely adopted, Latency-sensitive end-customer-facing applications are often required to address situations not covered by the planning process and direct customer demands that must occur close to the selling location. Rapid fulfillment is needed for supermarket orders placed by private customers. Generative order fulfillment refers to handling only some of the sales orders generated and covers a small fraction of total orders. The retail concessionaire or wholesaler handles much of the planning and allocation of physical units within the supply chain to keep service levels high, supported by direct replenishment programs. The generative process complements longer-term planning by optimizing the current status of physical units and demand patterns.




Fig 5: Edge-Integrated Fulfillment & Domain-Specific Deployment Architectures In Retail And Industrial Ecosystems

A. Industrial supply chain applications

In the future, local buyers or customers could make a purchase request via automated social channels (e.g. messaging apps or social media). The selling organization would directly acknowledge the request and attempt to fulfill it via local resources. This type of sales channel shifts supply chain management from the core to the edge of the organization. The company expects that, with the advances in generative AI and edge computing, such applications could soon become a reality.

Companies that operate industrial supply chains could deploy that model for order management. Local buyers may initiate purchase requests via automated social channels such as messaging applications or social media. The selling organization could acknowledge their requests directly and seek to fulfill them using local resources, thereby shifting supply chain management from the core to the edge of the organization. With recent advances in generative AI and edge computing, such applications could soon become a reality.

B. Retail and logistics use cases

Generative AI can improve the customer experience and the financial performance of companies operating in logistics, retail, hospitality, and food services. Four representative applications for these sectors illustrate the potential benefits and impact.

Retailers need to integrate dynamic pricing with omnichannel inventory availability. Some products can be picked at the store, while other less-selling items must be shipped from another location. When utilizing third-party logistics, the real-time transport offer must also be considered to optimize both price and delivery time. Dynamic pricing is optimized by integrating external information and generating internal information that affects demand.

Suppliers combine demand and supply information to optimize product availability where it is most demanded at the lowest-possible cost. Offering good margin products all

the time everywhere is not realistic. When demand for a product in a certain area is extremely low or when margins fall below a certain level, the product should not be available. Automation is needed to control the decisions that affect the supply chain, sales, and customer satisfaction. Optimization of delivery times can be a differentiating factor. Real-time transport offers allow dynamic routing and real-time reassessment of promised delivery times.

XII. RISK, ETHICS, AND GOVERNANCE

Generative AI opens the door to powerful, previously unattainable technologies, but these advances raise pressing questions about responsible use and governance mechanisms. Businesses, customers, and regulators all have a stake in AI systems that are explainable, auditable, and capable of maintaining societal trust. Examining the entire solution—the data, models, and decision process—facilitates the governance applied to AI-driven decision making.

Responsible deployment of generative AI processes across edge-to-cloud supply chains requires careful planning. Potential privacy infringement and unfair bias must be considered and mitigated. Generative models introduce new risks, and common sense dictates the use of all possible tools. It is hardly surprising, then, that these systems demand adversarial risk mitigation. Auditing is essential in order to validate the training records and that the decisions made are fair. Moreover, a fail-safe approach to such decision making must be established: all decisions based on generative models—or any model for that matter—are latent. A key aspect of supply chains is data availability, and any request for data that is not currently available must be satisfied via other means, such as a query to a generative model.

A. Responsible AI and transparency

Responsibility incorporates fairness, accountability, and transparency. Fairness concerns arise from the possibility of biased data leading to decisions harmful to individuals and communities. Since the automated processes in edge computing and cloud data centers are generally closed to endpoint users, limited auditability can negatively impact trustworthiness. Therefore, auditors should be able to assess how models reason, make decisions, and adapt. Transparency enables key stakeholders to understand how AI impacts their roles. Some end users also require comprehensive understanding and awareness of how models infer command outputs. Management audit should thus cover the entire command lifecycle and involve data, model, user, and decision audits.

Fairness implications emerge from the datasets employed to train specific generative models and the resulting decision policy. Probabilistic decision policies that weight past patterns equally are traditionally considered fair. For example, if an ordering policy consistently sanctions a group variation, perhaps because such variations were not submitted initially, a concentration assumption allows models to remain fair. In many cases, however, this level of fairness is insufficient. Robustness and reliability against adversarial shifts can also be interpreted as facets of fairness. Although temporal lists provide some protection by forcing models to remember all patterns, other types of generalization, such as spatial proximity in space and action, potentially mitigate this risk.

B. Data governance in edge and hybrid environments

Responsible use of AI requires data governance frameworks addressing ownership, provenance, lineage, privacy, and compliance. In edge and hybrid cloud environments, governance policies should clarify boundaries between public and private sectors, mandate consideration of multiple user perspectives, ensure policy visibility, and specify compliance mechanisms. These elements guide data preparation, collection, orchestration, and dissemination.

Compliance with data privacy regulations like the General Data Protection Regulation (GDPR) necessitates robust governance frameworks. Such provisions enable AI training systems to respect data subjects' privacy preferences through appropriate anonymization, pseudonymization, or removal while maintaining data usability. Enhanced governance facilitates the distribution of AI training across data centers, ensuring privacy protection during content creation and subsequent global knowledge consolidation. In addition to data subjects, data donors, data infrastructure owners, and AI designers must support heightened governance.

Table 4 — Deployment comparison derived

Criterion	Edge	Hybrid	Cloud
Network latency	very low	moderate	high
Real-time responsiveness	excellent	very good	weaker
Scalability	moderate	high	very high
Local autonomy	high	medium-high	low
Governance complexity	medium	high	medium
Best use case	strict latency	balanced workloads	bulk/global processing

XIII. ROADMAP AND FUTURE DIRECTIONS

Emerging technologies and standards: List promising innovations and relevant standards bodies.

Next-generation order management will leverage latency-sensitive generative AI techniques supported by the customised orchestration of edge-to-cloud data and computing resources. Consequently, future directions for the technology area are manifold. Emerging technologies such as—blockchain, fast in-memory databases, secure enclaves, next-generation telemetry protocols (for example, TRON), event-mesh data processing frameworks and data abstraction layers, 5G networks, complex event processing, natural-language processing (NLP) techniques, federated-learning and related privacy-preserving machine-learning solutions, and summarisation-oriented large-language models— have major impacts.

The importance of these technologies is underscored by their treatment as key priorities for research and development by the International Telecommunications Union (ITU-T). They are central to next-generation implementations that address the high-throughput, low-latency and privacy-

preserving requirements of users through the orchestration of distributed data movement, processing and model inference.

Standards bodies such as the Telecommunication Industry Association (TIA), International Telecommunication Union (ITU-T), 3rd Generation Partnership Project (3GPP), specializing in mobile technology, and the Crypto Valley Association (CVA), aligned with the Swiss Financial Market Supervisory Authority (FINMA) and guarding the interests of the Crypto Valley ecosystem, have prioritized these technologies. The increasing interest in these technologies is reflected in the growing body of contributions to the open-source community across Github, TensorFlow, Caffe2, PyTorch and comparable platforms associated with established cloud vendors and normalised under the umbrella of the Observability Working Group of the Cloud Native Computing Foundation (CNCF). Scenarios for edge-to-cloud deployment constitute another area demanding further focused research and effort. Suitable indicators, critical success factors and metrics for technology adoption, operation governance and ongoing performance improvements should also be identified and adapted for these end-to-end use cases. Best practices for ongoing implementation of end-to-end edge-and-cloud readiness and execution should also formulate.




Fig 6: Global Standards & Community Influence

A. Emerging technologies and standards

The interplay of edge computing and hybrid cloud data centers introduces new challenges for real-time processing and decision-making. Many workloads running in these environments require low inference latency to enable the detection or prediction of events and execution of actions in a timely manner. Furthermore, some prediction tasks can be effectively solved with relatively small datasets, thus making it feasible to train dedicated models for specific applications at the edge. Among these are new generation order management functions that support supply chain, e-commerce, and logistics applications. Generative AI techniques can accelerate these workloads by providing large language models capable of anticipating multiple types of events as well as scheduling, sequencing, differentiating, and orchestrating the fulfillment of all types of orders in the system.

Generative AI is creating interest across industry verticals, especially with respect to the development of foundation models capable of answering a variety of questions. However, very few industrial use-case deployments are currently active, and the technology still

lacks presence in edge environments, where real-time data provenance and smart data orchestration are required. Recent industrial supply chain disruptions are pushing firms to develop new capabilities. Generative AI can assist the tailoring of products and services to individual customer preferences, detect and predict the occurrence of costly-to-serve events, and closely monitor adherence to logistics service-level agreements. Such historical, context-oriented information can only be generated and reasons about if enabling functions are error free.

EQUATION 5 — FORECAST ACCURACY SCORE FROM ERROR

Step 1: start from normalized error

Let normalized error be

$$E_n = \frac{1}{T} \sum_{t=1}^T \left| \frac{y_t - \hat{y}_t}{y_t + \epsilon} \right|$$

where $\epsilon > 0$ avoids division by zero.

Step 2: convert error to accuracy

$$A_f = 1 - E_n$$

Final form

$$A_f = 1 - \frac{1}{T} \sum_{t=1}^T \left| \frac{y_t - \hat{y}_t}{y_t + \epsilon} \right|$$

B. Adoption pathways and best practices

A phased adoption roadmap encourages gradual engagement, culminating in the whole-systems approach outlined for end-to-end order lifecycle management. Organizations with established domain expertise should initially concentrate on embedding generative capabilities into those decision processes that are latency-sensitive (e.g., forecasting, scheduling, order fulfilment) rather than introduce additional generative capabilities across the entire decision-making chain.

Each stage should aim to contain the data management and governance complexity by using copies of data from cloud-based public repositories. These data must be used with care and subsequently discarded or retained only for specific compliance reasons—not necessarily for the decision model—where bias is prevented by ensuring an even distribution of samples across provenance classes. Furthermore, such a strategy avoids investing effort in building and maintaining complete descriptions of the data, as it suffices to attribute the provenance of the generated local data to the source data of cloud repositories. The only part of the model that requires specific location data is the need to identify these repositories serving as cloud data sources.

KPI for the generative approach becomes easy: whenever customer demand behaves in a manner consistent with the historical behaviour of the key available predictors contained in a generative-type forecasting model, the whole set of decisions governing the supply chain can be executed in real time. This is particularly true at the low-volume end of the pyramid. It is when the supply chain behaves in an unordered

and less predictable fashion that the burden of the high-speed-order-fulfilment decisions is lifted from the engines.

XIV. CONCLUSION

Generative AI radically transforms order management along the edge-to-cloud continuum. Responsiveness, reliability, and explainability are paramount for latency-sensitive workflows, prompting a rethink of process design and the technology stack. Generative approaches underpin responsive forecasting and scheduling, while traditional methods ensure correctness and compliance. The analysis identifies principles, concepts, and algorithms for forward-thinking applications and establishes a reference framework that integrates technical orchestration, data lineage, and governance.

Emerging applications in industrial supply chains and logistics underscore responsiveness, with credible forecasts enabling proactive mitigation of disruptions. Retail-oriented scenarios demonstrate cost benefits and reductions in unsold stock, showcasing the potential for next-gen order management. Progressing beyond demonstrators, tangible implementations are now an attainable objective.

REFERENCES

1. Alsurdeh, R., Calheiros, R. N., Matawie, K. M., & Javadi, B. (2021). Hybrid workflow scheduling on edge cloud computing systems. *IEEE Access*, 9, 134783–134799.
2. Ayoade, G., Karande, V., Khan, L., & Hamlen, K. W. (2021). Decentralized IoT data management using edge computing: A survey. *IEEE Internet of Things Journal*, 8(20), 15145–15159.
3. Mattaparthi, R. (2024). Transformer-Based Fault Diagnosis for Large-Scale Standby Power Generators: Partial Discharge Pattern Recognition at Hyperscale Data Center Installations. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8781–8799.
4. Gill, S. S., Xu, M., Ottaviani, C., Patros, P., Bahsoon, R., Kaur, R., Garraghan, P., & Buyya, R. (2022). AI for next-generation computing: Emerging trends and future directions. *Internet of Things*, 19, 100514.
5. Buyya, R., Srirama, S. N., Casale, G., Calheiros, R. N., Simmhan, Y., Varghese, B., & Gelenbe, E. (2022). A manifesto for future generation cloud computing. *ACM Computing Surveys*, 55(5), 1–38.
6. Oladosu, S. A., Ige, A. B., Ike, C. C., Adepoju, P. A., Amoo, O. O., & Afolabi, A. I. (2022). Revolutionizing data center security: Conceptualizing a unified security framework for hybrid and multi-cloud data centers. *Open Access Research Journal of Science and Technology*, 5(2), 086–076.
7. Varghese, B., & Buyya, R. (2022). Next generation cloud computing: New trends and research directions. *Future Generation Computer Systems*, 79, 849–861.
8. Li, Y., Zhao, T., & Zhang, H. (2022). Intelligent task scheduling for edge-cloud computing using deep reinforcement learning. *Future Generation Computer Systems*, 126, 128–140.

9. Davuluri, P. N. (2019). Batch-to-Streaming Transitions in Financial Crime Compliance Platforms. *International Journal Of Engineering And Computer Science*, 8(12).
10. Mourad, A., Tout, H., Wahab, O. A., Otok, H., & Guizani, M. (2022). Edge computing for the Internet of Things: A survey. *IEEE Internet of Things Journal*, 9(2), 949–972.
11. Khan, W. Z., Rehman, M. H., Zangoti, H. M., Afzal, M. K., Armi, N., & Salah, K. (2022). Industrial internet of things: Recent advances, enabling technologies, and open challenges. *Computers & Electrical Engineering*, 81, 106522.
12. Ivanov, D., & Dolgui, A. (2022). The industry 5.0 framework: Viability-based integration of resilience, sustainability, and digitalization. *International Journal of Production Research*, 60(5), 1688–1704.
13. Reddy, V. A. R. (2023). Predictive Healthcare Administration Using Advanced Payer Analytics and Population Health Data Engineering. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(2), 7967-7978.
14. Javaid, M., Haleem, A., Singh, R. P., Khan, I. H., Suman, R., & Rab, S. (2023). Unlocking the opportunities through ChatGPT and generative AI technologies. *Benchmark Transactions on Benchmarks, Standards and Evaluations*, 3(2), 100061.
15. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI. *International Journal of Information Management*, 71, 102642.
16. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126.
17. Kolla, S. K. (2022). Engineering Healthcare Data Infrastructures for Predictive Clinical Analytics and Evidence-Based Decision Making. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 4(5), 5370-5380.
18. Jackson, I., Ivanov, D., Dolgui, A., & Namdar, J. (2024). Generative artificial intelligence in supply chain and operations management: A capability-based framework for analysis and implementation. *International Journal of Production Research*, 62(18), 6120–6145.
19. Tian, Y., Zhang, Z., Yang, Y., Chen, Z., Yang, Z., Jin, R., Quek, T. Q. S., & Wong, K.-K. (2024). An edge-cloud collaboration framework for generative AI service provision with synergetic big cloud model and small edge models. *IEEE Network*, 38(5), 37–46.
20. Gill, S. S., Golec, M., Hu, J., Xu, M., Du, J., Wu, H., Walia, G. K., Murugesan, S. S., Ali, B., Kumar, M., Ye, K., Verma, P., Kumar, S., Cuadrado, F., & Uhlig, S. (2024). Edge AI: A taxonomy, systematic review and future directions. *ACM Computing Surveys*.
21. Mangalampalli, B. M. (2024). Transparent Intelligence Explainability Frameworks for AI-Driven Clinical Decision Support in Healthcare Business Intelligence. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 7(3), 10566-10579.
22. Raisch, S., & Fomina, K. (2024). Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review*, 49(2), 241–265.
23. Khaledian, N., Völz, M., Azizi, S., & Hussain, F. (2024). AI-based and heuristic workflow scheduling in cloud and fog computing: A systematic review. *Cluster Computing*, 27, 13099–13133.
24. Zhang, C., Zhang, J., Sangaiah, A. K., Li, D., & Li, W. (2024). Evaluating edge artificial intelligence-driven supply chain management platforms using collaborative large-scale fuzzy information fusion. *Applied Soft Computing*, 156, 111686.
25. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. *Metallurgical and Materials Engineering*.
26. Jiang, Y., & Sun, X. (2024). Efficient workflow scheduling in edge cloud-enabled space-air-ground-integrated information systems. *International Journal on Semantic Web and Information Systems*, 20(1), 1–29.
27. Patel, D., Raut, G., Cheetirala, S. N., Nadkarni, G. N., Freeman, R., Glicksberg, B. S., Klang, E., & Timsina, P. (2024). Cloud platforms for developing generative AI solutions: A scoping review of tools and services. *arXiv*.
28. Wamba, S. F., Guthrie, C., Queiroz, M. M., & Minner, S. (2024). ChatGPT and generative artificial intelligence: An exploratory study of key benefits and challenges in operations and supply chain management. *International Journal of Production Research*, 62(16), 5676–5696.
29. Sharma, P., Gunasekaran, A., & Subramanian, G. (2024). Enhancing supply chain: Exploring and exploiting AI capabilities. *Production Planning & Control*.
30. Li, L., Liu, Y., Jin, Y., Cheng, T. C. E., & Zhang, Q. (2024). Generative AI-enabled supply chain management: The critical role of coordination and dynamism. *International Journal of Production Economics*, 277, 109388.
31. Pamisetty, V., & Amistapuram, K. Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning.
32. Jackson, I., Ivanov, D., Dolgui, A., & Namdar, J. (2024). Generative artificial intelligence in supply chain and operations management: A capability-based framework for analysis and implementation.

- International Journal of Production Research, 62(17), 6120–6145.
33. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126.
 34. Tian, Y., Zhang, Z., Yang, Y., Chen, Z., Yang, Z., Jin, R., Quek, T. Q. S., & Wong, K.-K. (2024). An edge-cloud collaboration framework for generative AI service provision with synergetic big cloud model and small edge models. *IEEE Network*, 38(5), 37–46.
 35. Gill, S. S., Golec, M., Hu, J., Xu, M., Du, J., Wu, H., Walia, G. K., Murugesan, S. S., Ali, B., Kumar, M., Ye, K., Verma, P., Kumar, S., Cuadrado, F., & Uhlig, S. (2024). Edge AI: A taxonomy, systematic review and future directions. *ACM Computing Surveys*.
 36. Wang, Y.-C., Xue, J., Wei, C., & Kuo, C.-C. J. (2023). An overview on generative AI at scale with edge-cloud computing. *arXiv*.
 37. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI. *International Journal of Information Management*, 71, 102642.
 38. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. *Journal of Information Systems Engineering and Management*.
 39. Javaid, M., Haleem, A., Singh, R. P., Khan, I. H., Suman, R., & Rab, S. (2023). Unlocking the opportunities through ChatGPT and generative AI technologies. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 3(2), 100061.
 40. Raisch, S., & Fomina, K. (2024). Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review*, 49(2), 241–265.
 41. Khaledian, N., Völz, M., Azizi, S., & Hussain, F. (2024). AI-based and heuristic workflow scheduling in cloud and fog computing: A systematic review. *Cluster Computing*, 27, 13099–13133.
 42. Kolla, T. (2024). Intelligent Discovery and Governance of Healthcare Data Assets Through AI-Powered Catalog Architectures. *International Journal of Emerging Trends in Engineering and Management Research*, 9(4), 16083.
 43. Zhang, C., Zhang, J., Sangaiah, A. K., Li, D., & Li, W. (2024). Evaluating edge artificial intelligence-driven supply chain management platforms using collaborative large-scale fuzzy information fusion. *Applied Soft Computing*, 156, 111686.
 44. Jiang, Y., & Sun, X. (2024). Efficient workflow scheduling in edge cloud-enabled space-air-ground-integrated information systems. *International Journal on Semantic Web and Information Systems*, 20(1), 1–29.
 45. Patel, D., Raut, G., Cheetirala, S. N., Nadkarni, G. N., Freeman, R., Glicksberg, B. S., Klang, E., & Timsina, P. (2024). Cloud platforms for developing generative AI solutions: A scoping review of tools and services. *arXiv*.
 46. Bandi, V. D. V. K. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. *Metallurgical and Materials Engineering*, 30(4), 1011–1027.
 47. Sharma, A. J., & Rathore, B. (2024). Examine the enablers of generative artificial intelligence adoption in supply chain: A mixed method study. *Journal of Decision Systems*.
 48. Ivanov, D., & Dolgui, A. (2022). The industry 5.0 framework: Viability-based integration of resilience, sustainability, and digitalization. *International Journal of Production Research*, 60(5), 1688–1704.
 49. Dolgui, A., & Ivanov, D. (2022). 5G in digital supply chain and operations management: Fostering flexibility, end-to-end connectivity and real-time visibility through Internet-of-Everything. *International Journal of Production Research*, 60(2), 442–451.
 50. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.
 51. Buyya, R., Srirama, S. N., Casale, G., Calheiros, R. N., Simmhan, Y., Varghese, B., & Gelenbe, E. (2022). A manifesto for future generation cloud computing. *ACM Computing Surveys*, 55(5), 1–38.
 52. Gill, S. S., Xu, M., Ottaviani, C., Patros, P., Bahsoon, R., Kaur, R., Garraghan, P., & Buyya, R. (2022). AI for next-generation computing: Emerging trends and future directions. *Internet of Things*, 19, 100514.
 53. Wang, Y.-C., Xue, J., Wei, C., & Kuo, C.-C. J. (2023). An overview on generative AI at scale with edge-cloud computing. *IEEE Open Journal of the Communications Society*, 4, 2952–2971.
 54. Tian, Y., Zhang, Z., Yang, Y., Chen, Z., Yang, Z., Jin, R., Quek, T. Q. S., & Wong, K.-K. (2024). An edge-cloud collaboration framework for generative AI service provision with synergetic big cloud model and small edge models. *IEEE Network*, 38(5), 37–46.
 55. Gill, S. S., Golec, M., Hu, J., Xu, M., Du, J., Wu, H., Walia, G. K., Murugesan, S. S., Ali, B., Kumar, M., Ye, K., Verma, P., Kumar, S., Cuadrado, F., & Uhlig, S. (2024). Edge AI: A taxonomy, systematic review and future directions. *ACM Computing Surveys*.
 56. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126.
 57. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.

58. Li, L., Liu, Y., Jin, Y., Cheng, T. C. E., & Zhang, Q. (2024). Generative AI-enabled supply chain management: The critical role of coordination and dynamism. *International Journal of Production Economics*, 277, 109388.
59. Wamba, S. F., Guthrie, C., Queiroz, M. M., & Minner, S. (2024). ChatGPT and generative artificial intelligence in operations and supply chain management. *International Journal of Production Research*, 62(16), 5676–5696.
60. Sharma, A. J., & Rathore, B. (2024). Examine the enablers of generative artificial intelligence adoption in supply chain: A mixed method study. *Journal of Decision Systems*.
61. Khaledian, N., Völz, M., Azizi, S., & Hussain, F. (2024). AI-based and heuristic workflow scheduling in cloud and fog computing: A systematic review. *Cluster Computing*, 27, 13099–13133.
62. Haddadha, P. K., Rezvani, M. H., MollaMotalebi, M., & Shankar, A. (2024). Machine learning methods for service placement: A systematic review. *Artificial Intelligence Review*, 57, Article 61.
63. Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8589-8601.
64. Fabri, L., Häckel, B., Oberländer, A. M., Rieg, M., & Stöhr, A. (2023). Disentangling human-AI hybrids: Conceptualizing the interworking of humans and AI-enabled systems. *Business & Information Systems Engineering*, 65, 623–641.
65. Alkouz, B., Al-Rahayfeh, A., & Alouneh, S. (2023). Distributed artificial intelligence: Taxonomy, review, framework, and reference architecture. *Intelligent Systems with Applications*, 18, 200231.
66. Ivanov, D., & Dolgui, A. (2022). The industry 5.0 framework: Viability-based integration of resilience, sustainability, and digitalization. *International Journal of Production Research*, 60(5), 1688–1704.
67. Dolgui, A., & Ivanov, D. (2022). 5G in digital supply chain and operations management: Fostering flexibility, end-to-end connectivity and real-time visibility through Internet-of-Everything. *International Journal of Production Research*, 60(2), 442–451.
68. Buyya, R., Srirama, S. N., Casale, G., Calheiros, R. N., Simmhan, Y., Varghese, B., & Gelenbe, E. (2022). A manifesto for future generation cloud computing. *ACM Computing Surveys*, 55(5), 1–38.
69. Gill, S. S., Xu, M., Ottaviani, C., Patros, P., Bahsoon, R., Kaur, R., Garraghan, P., & Buyya, R. (2022). AI for next-generation computing: Emerging trends and future directions. *Internet of Things*, 19, 100514.
70. Javaid, M., Haleem, A., Singh, R. P., Khan, I. H., Suman, R., & Rab, S. (2023). Unlocking the opportunities through ChatGPT and generative AI technologies. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 3(2), 100061.
71. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI. *International Journal of Information Management*, 71, 102642.
72. Irshad, R. R., Hussain, Z., Hussain, I., Hussain, S., Asghar, E., Alwayle, I. M., Alalayah, K. M., Yousif, A., & Ali, A. (2024). Enhancing cloud-based inventory management: A hybrid blockchain approach with generative adversarial network and elliptic curve Diffie-Hellman techniques. *IEEE Access*, 12, 25917–25932.
73. Nezami, Z., Hafeez, M., Djemame, K., & Zaidi, S. A. R. (2024). Generative AI on the edge: Architecture and performance evaluation. *arXiv*.
74. Chen, D., Youssef, A., Pendse, R., Schleife, A., Clark, B. K., Hamann, H., He, J., Laino, T., Varshney, L., Wang, Y., & colleagues. (2024). Transforming the hybrid cloud for emerging AI workloads. *arXiv*.