




Agentic AI for Cloud-Native Big Data Orchestration in National Food Wholesale Supply Chains

Vinod Battapothu

Independent Researcher

vinod.battapothu.researcher@gmail.com

Abstract

Undertaking empirical research in the food service wholesale domain of national-scale (US) supply chains, this work focuses on agentic AI-enabled orchestration for cloud-native supply-chain optimization. Problems derive from heavyweight supply chains serving capital-- and land-intensive industries, such as agriculture and forestry, with demand patterns poorly driven by underlying ecosystems yet bolstered by growth dynamics in surrounding food service markets. Discussion is framed around orchestration of big data pipelines for supply-chain analytics and decision support. An extended tenet of adaptive governance empowers business-level stakeholders to integrate data sources and keep their national networks in sync. Contributions articulate core supply-chain requirements and actors' analytical needs, introduce a methodological framework for designing data-ingestion and orchestration mechanisms, and recommend a reference system architecture that generalizes to other national-scale economies. Spanning compliance, quality, safety, and cybersecurity, these considerations together speak to agentic AI's accountability, transparency, and risk.

The empirical core comprises synthesis case studies feeding into three use cases for cloud-enabled supply-chain dynamics: forecasting end-customer demand, integrating supplier inventories and resources, and optimizing transport. Even with only a few data sources, agentic AI-enhanced decision support promises genuine advantage: helping generate matching supplier networks and synchronizing transport operations not just for cost but also to mitigate product waste, quality loss, and spoilage. A supporting technical framework addresses the cloud-native supply-chain-analytics components of systems for national-scale economies—conducting core supply-chain planning in the food-service wholesale domain, such as forecasting, integrating supplier inventories, and synchronizing transport across capital-- and land-intensive supply chains that depend on nature's undulating growth cycle.


Keywords : Agentic AI, big data orchestration, cloud-native architectures, supply chain optimization, national food service wholesale, data governance, interoperability, security, resilience.

1. Introduction

Today's complex supply chains present novel challenges across multiple sectors. These challenges originate from a myriad of sources, including climate change, cybersecurity turnover, pandemics, geopolitical tensions, and connectedness. Extreme weather patterns and natural disasters have a growing impact on the availability of resources as the risks of a supplier failing to deliver products on time increases. The COVID-19 pandemic illustrated the dangers of relying on suppliers concentrated in one

geographical region and localizing customer bases. Rapid turnover of skilled personnel increases the risk of failure during these peaks. Transport and logistics furthermore face issues including inflationary pressures, shortages of suitable drivers, and chaotic energy prices, risks that cannot be ignored. Companies must try to absorb these risks and protect their profitability while remaining competitive with respect to the products they sell.

Big Data has the potential to address these challenges by synchronizing data across disparate stakeholders to improve



visibility, confidence, and collaboration. Within the food service sector suppliers and distributors are usually located in defined geographical areas and provide a local service to regional-scale restaurants, hotels, and institutions.

wholesale. Agentic orchestration improves the effectiveness and efficiency of supply chains by enabling multiple use cases in demand forecasting, supplier collaboration, and logistics synchronization.





Fig 1: Intelligent Supply Chain Process Automation with Agentic AI

1.1. Background and Significance Cloud-native supply chains for national food service wholesale require support for orchestrating data across heterogeneous actors, services, and systems. Current orchestration methods lack interoperability with diverse sources and destinations, relying instead on the adaptation of data sources and consumers for each new use. An agentic AI-enabled and declarative big data-orchestration framework overcomes existing limitations by leveraging a cloud-native infrastructure complemented by a modern data fabric. Orchestration requires the ingestion, transformation, and storage of big data, as well as decision-making that leads to intervention in the real world. Together with the data landscape, governance arrangements, and compliance, safety, and quality standards at the intersection of business operations and domain knowledge, these represent the set of conditions that support orchestration in national food service

An agentic AI-enabled orchestration framework for cloud-native national-scale supply chains supports any combination of cloud-native data sources and sinks within the implementation landscape. Declarative workflows and operation-level configuration enable orchestration of data-aware computations and systems. Cloud-native computing with the modern data fabric provides a dynamic, scalable, and performant infrastructure for ingesting, transforming, and storing big data in support of demand predictions and a wide range of additional use cases across different supply-chain actors. As Camus observed during the French Revolution, one can hardly deal the masses without a little automation. This covering can be improved by sharing the data of multiple players to create a National Data Ecosystem at scale to support demand forecasting in order to build a synchronized supply chain in terms of purchasing, stock, transport and logistics decisions. Big Data Orchestration is an emerging concept to exploit the predictive capabilities of artificial intelligence to monitor the incoming streams of data in order to recommend, regular and execute optimized actions on downstream data streams.

1.2. Research design First-of-a-kind investigation of technology-enabled knowledge orchestration for cloud-native supply chain optimization of national food service wholesale products leverages agentic AI, big data orchestration, and cloud-native architecture to create a formalized methodological framework and reference architecture that address specific operational requirements. Complementary exploration of the national food service wholesale ecosystem includes stakeholder mapping, data landscape assessment, discussion of compliance risks and considerations, and evaluation of three demand-forecasting-and-inventory-optimization use cases. Addressing technical requirements for orchestration enabled by modern data fabrics further supports exploration of security, data-loss



prevention, and disaster-recovery resilience-testing measures.

Knowledge orchestration focuses on enabling process- and domain-centric knowledge creation, sharing, and exchange. Responsive to the information-as-a-service paradigm where service consumers are manufacturers and suppliers of data as well as services, demand contexts govern the required service operations and service products. Orchestration processes are composed of multiple interlinked sub-processes in a workflow. Business process workflow orchestration is addressed in the context of agentic AI, where a set of actuators can independently trigger a workflow, whether directly, via connectors from another workflow or via mapped triggers from data products published by domain-subject or service-domain experts.


2. Theoretical Foundation

Orchestration enables adaptive integration and coordination of data flows. Cloud-native architectures improve responsiveness and resilience by embracing the scalable and elastic characteristics of cloud infrastructures. Agentic AI uses actionable intelligence for self-aware decision-making and independent execution of complex actions. In combination, these areas of research support decisive automation and optimization in supply chains.

Big Data Orchestration. Big data orchestration focuses on the adaptive integration and coordination of distributed components that provide and consume data. A declarative workflow-abstraction model enables business units to compose, manage, and monitor analytic pipelines and expose analytic results and consumable services for downstream consumption. Data and analytic-supply processes adapt to changing requirements and conditions through an event-driven architecture supported by decision trees and rules.

Cloud-Native Architectures. Cloud-native architectures harness the native properties of cloud-computing


infrastructures to improve system responsiveness, resilience, and the user experience. Elasticity, scalability, decoupled operation, and geographic distribution support rapid application-development cycles; improved fault tolerance; latent, on-demand computing capacity; and decisively enhanced user experiences. Such properties, often cited as descriptions of “native” cloud systems, are characteristics of a new generation of multiapplication and multitenant systems that have been introduced in the past five years.



2.1. Big Data Orchestration

Big data management for cloud-native applications goes well beyond facilitating fast data ingestion, storage, and retrieval through modern data lakes. Approaches for big data orchestration provide specialized mechanisms and abstractions to define complex multi-step data processing workflows that can be triggered by external events, scheduling policies, or internal conditions. These workflow orchestrations seamlessly combine sensor data ingestion, data processing, model execution, and execution result handling at any point in such processing pipelines, supporting reliable processing at the national level.

Because infrastructure-as-code facilitates ready-to-use cloud deployments, a standard model may be used to support pilot-scale implementation and testing at regional or local levels.



Reuse of pilot projects at national scale leads to agent-based AI implementation for balance sheet optimization. The broad scope of agentic decision-making in this sense goes well beyond supply chain optimization per se throughout the food service wholesale supply chain, spanning control of multiple warehouses, wholesale distribution centers in “expect the unexpected” and “mitigating risks” use cases.

2.2. Cloud-Native Architectures

Cloud-native architectures enable digital transformation through Business Process as a Service (BPaaS), Software as a Service (SaaS), and related paradigms. Architectures for public cloud infrastructures use commercial-off-the-shelf (COTS) products as platforms, while hybrid and edge cloud infrastructures present more complex design choices. Cloud-native application architectures, independent from the infrastructure, have gained traction within the research community. However, these advancements remain low-level and disconnected from orchestration for Big Data Deployment as a Service (BDaaS).

The heterogeneity of the underlying data in Cloud-Native Architectures poses a further challenge for applications or services operated on top of multiple providers. Virtual institutions and cross-organizational collaboration are hindered by difficulties in sharing and integrating the underlying information within the user-centric Cloud ecosystems. Interoperability is therefore a key requirement for business applications for Cloud-Native architectures. However, Cloud-Native outsourcing remains vulnerable to the proprietary nature of the Cloud Providers hosting the application elements; i.e., Pain points include: coupling to a particular provider hinders the federation of a user-centric Cloud and reduces the fault-tolerance of the BaaS chain. Services and applications in Cloud-Native outsourcing on multiple Cloud Providers for HaaS, IaaS, PaaS, and SaaS are designed and developed independently from the supporting resources.

Equation 1: Forecast accuracy metrics (ties to evaluation discussion)

Let forecast error:

$$e_t = D_t - \hat{D}_t$$

RMSE

1. Square errors: e_t^2
2. Average: $\frac{1}{n} \sum_{t=1}^n e_t^2$
3. Square-root:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (D_t - \hat{D}_t)^2}$$


MAPE

1. Absolute percentage error: $|\frac{e_t}{D_t}|$
2. Average and convert to %:

$$\text{MAPE} = \frac{100}{n} \sum_{t=1}^n \left| \frac{D_t - \hat{D}_t}{D_t} \right|$$

2.3. Agentic AI and Autonomy in Supply Chains

Although research in support of AI-based decision-making has demonstrated artificial intelligence's ability to replace human-centred decision-making, it tends to occur in isolation. Supply-chain management, however, needs to control equal or higher sets of parameters in its decision process—especially unconstrained decisions—through agentic AI-based monitoring, negotiation, proposal, offer, and ordering, which are critical parts of its upper management. The increasingly complex networks of actors and influencing factors in supply chains require that many additional decisions be taken outside the human horizon, especially unconstrained decisions in AI-supported environments where true models of nature are available. Investments in training and labelling data acquisition to



support more human-centred decisions remain vital, but can benefit from AI support to fulfil a near-ideal training and development investment.

Agentic AI is a broad concept that can be defined as AI that has full or specific authority to make decisions or take actions on the business user's behalf. Agentic AI acts autonomously within certain guardrails; higher level processes of strategy and control are supported through monitoring and information provision. It can be small-scale AI fully delegated except for the most critical KPIs, such as crew fatigue or safety, or it can be more complex constructs that can carry out self-optimisation and self-governance while being watched. Agentic AI becomes essential in supply chains whenever external parameters, be it in the markets or in the physical environment, become unpredictable and complex, requiring a true, largely hidden partner decision-maker able to reduce the noise and optimise operations except for unforeseen occasions and thus allow scarce human decision-making bandwidth to focus on those critical moments.

3. Context and Requirements in National Food Service Wholesale

A single national food service wholesaler serves U.S. food service wholesalers. Wholesalers below national-scale deploy cloud-native architectures. Stakeholders include Department of Agriculture, Food and Drug Administration, Centers for Disease Control, and state agencies. A federal data security framework and a national supply chain risk management strategy guide modernizing supporting IT resources. Data involve supply, demand, and price dynamics; vendors; food product properties; incidents and recalls; and drivers (e.g., GDP, weather). Data quality and mapping, suitable standards, and venue normalization promote interoperability. Autonomy, resilience, and cybersecurity enhance vendor decision-making. Compliance-sensitivity, data volatility, privacy, and

protection strengthen firewalls. Data exposure and access rely on risk tolerance and business needs.

Requirements for agentic AI-enabled big data orchestration for cloud-native supply-chain optimization in national food service-wholesale are synthesized. The sector is democratically governed, data cover demand-supply dynamics and quality dynamics, and safety regulations ensure transparency and accountability. Evaluation requires an aligned taxonomy of agentic AI capability. The need for data volume and quality, governance design, and security-resilience correlations are also highlighted.





Fig 2: National Food Service Wholesale of Agentic AI-Enabled Big Data Orchestration

3.1. Stakeholders and Governance

A federated governance model is proposed that incorporates the needs and participation of all stakeholders in the national food service wholesale supply chain. In addition to food wholesalers and suppliers, stakeholders comprising technology providers, system integrators, consulting firms, and public authorities concerned with supply chain security must be recognized. Interoperable vendor-neutral solutions reduce the risk of vendor lock-in, allowing stakeholders to choose alternative cloud services from different providers. The oversight of public authorities is essential for data protection, compliance, and standardization. Finally, a



reliable and trustworthy collaboration between all stakeholders is a prerequisite for a successful deployment of agentic AI system designs.

For organizations to benefit from effective monitoring of their suppliers and other supply chain partners, data from all partners must be made available to the organizations' demand forecasting, risk assessment, and logistics management functions. Appropriate data-sharing agreements among partners help to ensure that data is made available for these collaborative purposes at a low cost while addressing privacy concerns. Moreover, alignment with relevant food quality monitoring standards facilitates the implementation of risk mitigation measures. Standardized data-sharing processes also enhance partners' preparedness for decisions on supply shortage mitigation and related methods of supplier collaboration.

3.2. Data Landscape and Interoperability

Successor national food service supply chain architectures depend on data from multiple sources. The volume, variety, and velocity of data creation demand ingestion into Big Data platforms. Interoperability among data sources and service providers enables governance beyond any single organization. Data fabrics enable efficient, scalable ingestion of data, whether sourced from IOT devices, data providers, or service providers. A cloud-native architecture supports a multi-cloud solution. Multi-agent systems support intelligent, agentic decision-making and actionable intelligence, generating deliverables at the required frequency. Architectural designs adopt implicit cloud-native principles, but the combined complexities of data fabrics, orchestration mechanisms, multi-agent systems, domain use cases, data governance, and security in a national-scale supply chain require explicit representations.


Data ingestion employs an advanced data-fabric abstraction that supersedes modern data pipelines and integrates additional features and capabilities. Data fabrics simplify the infrastructural complexities surrounding advanced ingestion of Big Data. They address core challenges associated with

ingesting complex, heavy-traffic data from multiple sources, meet external-contributory data-traffic challenges, provide fast-track capabilities for perishable-data ingestion, and optimize the services of an advanced data-fabric infrastructure without compulsory reliance on boards and dot-com glass-screen displays. Indeed, so extensive is the embrace of data-overload complexities in the outright scope of data fabrics that data processors are now potentially dispensed with.

3.3. Compliance, Safety, and Quality Standards

Supply chain participants intervene at different stages and create different processes to generate consumables. State and federal laws have established guidelines for these various activities; food service wholesale is subject not only to general income, commerce, and traffic regulations but also to those dealing specifically with food safety and security. Federal laws on the safety and sanitary aspects of foods are enforced mainly by the Food and Drug Administration (FDA), supported by state departments of agriculture and similar agencies in law. FSIS operates under the Department of Agriculture. FDA emphasizes standard specifications for final-consumed foods, while FSIS, so far, operates standards mainly to prevent animal disease. Laws also regulate the establishment and operation of food establishments, sanitation standards for food transportation, pesticide control, and food product inspection.

Besides FSIS, other agencies with food safety responsibilities include the Environmental Protection Agency (EPA), Centers for Disease Control and Prevention (CDC), US Customs and Border Control, National Marine Fisheries Services, and FDA's Center for Veterinary Medicine. Their rules govern areas such as food-contact substances in packaging; modern meat, poultry, and egg products inspection and labeling; poisonous and deleterious substances; interstate movement of diseased animals; pesticide residues, and preparation, serving, and storage of food products. In addition, the FDA also issues guidelines regarding safety and sanitary requirements for the operation of food establishments such as restaurants, snack bars,




delicatessens, supermarkets, stores, bakeries, and food-processing plants.

4. Methodological Framework

Cloud-native supply chain applications operating on a shared national data ecosystem must ingest structured and unstructured big data from diverse sources and in varying formats. Modern data ingestion platforms simplify big-data ingestion through a self-service model and enable ingestion through cleanroom environments in response to request-based workflows or dedicated streams for data in motion. Big data orchestration mechanisms that address cloud-native supply chain requirements—especially the simultaneous execution of multiple workflows and multi-agent collaboration—provide abstraction models. The abstracted workflows can be expressed in a general-purpose programming language and invoke business logic residing in external environments, leading to actionable intelligence within a cloud-native supply chain. To enable agentic decision-making, AI-based components consume both the data and the results generated by the orchestration mechanisms and can autonomously trigger associated actions within the cloud-native supply chain.


Modern data ingestion platforms abstract the underlying complexity of data ingestion into a self-service model and allow ingestion through cleanroom environments. Request-based workflows and dedicated streams for data in motion further support big data ingestion. Actionable intelligence within cloud-native supply chains emerges through the orchestration of all relevant workflows across the supply chain. Such workflows must not only be executed concurrently but also require a combination of agentic workflows triggered by intelligent supply-chain components capable of detecting certain business events. Moreover, the abstracted workflows must support multiple forms of invocation so that they can seamlessly incorporate business logic residing in external environments.



4.1. Data Ingestion and Modern Data Fabrics

The increasing volumes and variety of data demand modern data-fabric solutions to provide an integrated, single-source view of the data landscape across all stakeholders. Such a data view is essential for producing dashboards and reports across different stakeholders; feeding other operations, like demand, procurement, and production forecasting, inventory replenishment, supply, and logistics optimization; and supporting intelligent, agentic decision-making. Industry players are increasingly adopting the modern data-fabric paradigm, which ingests, organizes, prepares, and delivers metadata and data spread across various data silos to analytics ecosystems, so that data is timely, secure, relevant, understandable, and trustworthy for both humans and intelligent algorithms.

Modern data fabrics provide data ingestion processes to connect, prepare, monitor, catalog, and secure data from multiple sources at scale. Data pipelines seamlessly ingest data from systems such as enterprise resource planning (ERP) and supply chain management (SCM); point-of-sale systems; retailer inventory-management applications with forecast information; online sales channels; bulk purchase lists; and order-processing systems to create a comprehensive view of the business. Automating the ingestion of data into the business-intelligence environment



also cuts operational costs by allowing direct analysis of data without moving it to a separate analytics platform.

4.2. Orchestration Mechanisms and Workflow Abstractions

Major advances in cloud and edge computing have accelerated digital transformation efforts across all industries, resulting in growing investments in digital technologies and infrastructure. Businesses are embedding processes, workflows, applications, and services in digital infrastructures, resulting in enormous quantities of multi-structured data generated, stored, processed, and exchanged between devices and systems, both human-operated and autonomous. Enterprise architecture, service-oriented architecture, microservices-based architecture, event-driven architecture, and data-centric architecture can benefit from orchestration to enable business processes, optimize workload management, and synchronize events. Orchestration in cloud computing emphasizes automating the management and coordination of complex cloud services, applications, and resources to optimize workloads and enhance performance, security, and resilience. Built on technological advances such as big data, intelligent cloud, and machine learning, orchestration enables the automation of increasingly complex operations involving extended ecosystems using all available data for timely actionable intelligence. With the emergence of cloud-native architecture, the orchestration of all on-demand interactions becomes important across the life cycles of products and associated services.

Big data cloud orchestrators automate the management, control, and orchestration of data, information, and knowledge in big data operations. Such orchestrators govern the orchestration of a data-centric big data ecosystem in a unified manner, considering data as a first-class citizen, across all data-related activities during the life cycles of data, information, and knowledge assets. Big data orchestration simplifies the coordination of processes and flows for effective cloud resource utilization and cost savings, automates big data technologies through integrated

workflow patterns, and transforms the information generated into knowledge and insight, thus closing the big data circle. The orchestration of processes, workflows, and resources in big data fabric operations provides an additional operational perspective from the data perspective to ensure the stability of a big data fabric. Big data process orchestration automates the building of an information processing pipeline based on predefined patterns.

Equation 2: Inventory optimization equations

Let:

- D = annual demand (units/year)
- S = ordering/setup cost per order (\$/order)
- H = holding cost per unit per year (\$/unit/year)
- Q = order quantity (units)

1) Ordering cost per year

- Orders per year = D/Q
- Ordering cost:

$$C_{order}(Q) = S \cdot \frac{D}{Q}$$


2) Holding cost per year

- Inventory cycles linearly from Q down to 0
- Average inventory = $Q/2$
- Holding cost:

$$C_{hold}(Q) = H \cdot \frac{Q}{2}$$

3) Total relevant cost

$$C(Q) = S \frac{D}{Q} + H \frac{Q}{2}$$



4) Minimize by derivative

$$\frac{dC}{dQ} = -S \frac{D}{Q^2} + \frac{H}{2}$$

Set $\frac{dC}{dQ} = 0$:

$$-S \frac{D}{Q^2} + \frac{H}{2} = 0 \Rightarrow \frac{H}{2} = S \frac{D}{Q^2} \Rightarrow Q^2 = \frac{2SD}{H}$$

So:

$$Q^* = \sqrt{\frac{2SD}{H}}$$

4.3. Agentic Decision-Making and Actionable Intelligence

Delivering actionable intelligence requires not only thought processes plumbing uncertainty and risk but also the ability to choose an appropriate action and follow through.

Different layers of use-case typologies engage various types of decision-making frameworks and response profiles, ranging from structured processes with higher predictability to semi-structured, dynamic processes towards the other extreme, up to fully unstructured, highly reactive situations. These elements may combine into multi-period, multi-level decision architectures featuring synoptic holistic perspectives as well as agentic tip-of-the-pyramid implementation.

Fusing these different forms of deliberation requires a combination of the belief-desire-intention (BDI) agent framework with the related convincing/reaction model capturing acceptance and reactivity in categories of agents and the risk-being framework managerial decision-making in volatile environments. In addition to a clear object of analysis or control, agents need a clear sense of community with associated actors, concerns, emotional interest, and trust variables reflecting a shared sense of belonging and

commitment. Trust variables alert agents to signs of erosion and enable timely and appropriate behavioral decisions and, in turn, a clearly defined execution environment, generating near-term execution content for the agent.


5. System Design and Architecture

Large-scale solutions for orchestrating big data in the context of national food service wholesale require a rigorous architectural design for deployment and operation, together with practical implementations that realise key requirements. A reference architecture shows how agents-as-a-service, digital twins, machine learning models, and data ecosystems encompassing modern data fabrics, data meshes, and federated data governance support supply chain demand forecasting at national and state scales. Data management and support for multiple data-sharing models address security, privacy, vendor lock-in, and regulatory compliance. The architecture serves as a blueprint for developing and deploying a cloud-native supply chain decision-support system across the USA.

The design and architecture of a cloud-native Agentic AI-enabled big data orchestration platform meets key requirements identified in the context of US national food service wholesale food products. Cloud service delivery choice is complemented by a modular, extensible reference architecture based on agentic services, a data mesh, a federated data-sharing model, a modern data fabric, and security-by-design principles. By supporting data governance that embraces algorithmic transparency, accountability, and an explicit data-sharing policy, the architecture mitigates organizational and regulatory risk, creates a positive environment for foodborne disease mitigation and product recall, and addresses vendor lock-in.

5.1. Reference Architecture for National Scale

With its broad reach and wide range of stakeholders, national federal-level orchestration of data flows in the food supply chain requires a reference architecture that supports a



specific data mesh design pattern and data-as-a-product paradigm. Each regional operational territory should implement its own regional data mesh, and a federated governance structure must ensure overall interoperability between implementation, meet harmonization requirements, and guarantee the fulfillment of overriding national service-level agreements.

In the data-as-a-product paradigm, supply chain actors are not only consumers of internal data but also their producers. They validate the quality against their own business rules. These data products can be made available to all interested parties in a business-to-business approach. In addition to actors directly participating in the food chain, administrative bodies (e.g., health and inspection) can generate data products for quality control, while planning authorities can provide noise maps and cartographic products to support demand forecasting. Quality technical-scientific information for risk analysis is constantly supplied by knowledge networks and research bodies. Moreover, all these data products, as well as relevant managerial information of public interest, may finally converge into an open-data platform.




Fig 3: Reference Architecture for National Scale

5.2. Data Mesh and Federated Governance

National food service wholesale encompasses a wide spectrum of products and services ranging from cake bases and flavored syrups to paper products and sanitation


supplies, supplied by multiple specialized vendors. Achieving genuine demand-supply matching—timely availability of the right products in the right amounts and location—necessitates synchronized supply-side capacity and demand forecasts. High stakeholder diversity implies more complex data governance than vendor-controlled retail operations. Data meshes foster federated ownership, with product-specific vendors accountable for data quality. Fabric security must balance privacy, competitive advantage, and confidentiality concerns while minimizing vendor lock-in.

Food services are subject to compliance with safety and quality standards governing temperature, humidity, and allergen risks. Supply chain services for perishables require support for freshness data modeling with duration, reliability, and condition variations, and for freshness-optimized routing and allocation. Any shortage exception should trigger support for product substitution and supplier collaboration while implementing resilience-enhancing distribution designs.

5.3. Security, Privacy, and Resilience

Supply chains in the wholesale food service sector must comply with strict security, privacy, safety, and quality mandates. Vendors such as the U.S. Food and Drug Administration, the U.S. Department of Agriculture, and the Food Safety Inspection Service are tasked with protecting public health and welfare, and a single food commodity may be subject to oversight by multiple agencies. Security and safety regulations cover areas including intellectual property, cybersecurity, product certifications, and food quality and safety. These factors pose risks to data security and privacy; thus, the establishment of clear policies identified in the data governance framework is paramount.

The system must mitigate risk associated with confidentiality violations and be resilient against malicious attacks. Security and privacy should be enforced from the lower levels of a data architecture, for instance, through the use of hardware security modules for sensitive cryptographic functions and the adoption of intrinsically secure



confidential compute services. Moreover, privacy-preserving mechanisms, such as differential privacy and federated learning or federated transfers, should be in place. Resilience involves ensuring the continued operation of essential services during disturbances and the accessibility of the entire supply chain during crises. During non-crisis scenarios, careful planning of supply routes in the presence of alternative suppliers can help reduce visible costs.

6. Use Cases in Wholesale Food Service

The national wholesale supply chain for food service products comprises a multitude of stakeholders, major components, and countless interactions across many types of products. Demand signals, inventory levels, and transport movements are seldom synchronized. Each wholesaler, retailer, supplier, and logistics business in the supply chain applies its own forecasting, inventory-planning, and routing algorithms separately, and imperfect or missing data hinder decision quality. Regulators and consumers are increasingly onerous about food quality and safety, with rising expectations of data transparency and accessibility. These facts create visibility gaps, but at the same time opportunities exist to share data among parties capable of mutual benefit. Modern digital technologies can help across all stakeholders, but adoption costs and data silos impede progress.

Agentic AI-enabled big-data orchestration and its supporting cloud-native technology fabric open the door to actualizing control signals that are better than individual contributors working independently. Agentic decision-making, along with agent-mediated approaches to actionable intelligence in orchestrated data-fabric systems, allow for joint optimization of demand forecasting, inventory management, supplier collaboration, logistics synchronization, and more. Within food service wholesale, data mesh principles help ensure that data is shared when others can benefit, while federated data-governance mechanisms reduce the risk of vendor lock-in. Data-sensitive applications such as demand forecasting can use data-sharing patterns that optimize for algorithmic transparency and accountability, while compliance with

strict food-quality and safety regulations can be supported through transport documentation and cooldown-chain monitoring.

Equation 3: Transportation & logistics synchronization equations

Let:

- Nodes i, j in a distribution network
- $x_{ij} \geq 0$ = flow shipped from i to j
- c_{ij} = cost per unit shipped
- s_i = supply at node i
- d_j = demand at node j

Objective (min cost):

$$\min \sum_i \sum_j c_{ij} x_{ij}$$

Constraints

- Supply limits:


$$\sum_j x_{ij} \leq s_i \forall i$$

- Demand satisfaction:

$$\sum_i x_{ij} \geq d_j \forall j$$

- Non-negativity:

$$x_{ij} \geq 0$$



6.1. Demand Forecasting and Inventory Optimization

Reliable and transparent demand forecasting models are essential to managing supply chain dynamics and enabling optimal inventory levels. While each wholesaler has developed its own forecast models that reflect its knowledge of local market conditions, optimal stock levels are affected by synchronizing overall regional demand. AI-enabled methods are increasingly used for demand forecasting across all supply wholesaler levels and regions, as described by Goh & Daryanto, and empirical evaluations are presented in Huang et al., Goh et al., and Saberi et al. Such models use both bottom-up and top-down demand data.

The increasing focus on ESG goals across all sectors of the economy, including global transportation and national food wholesaling services, drives an acceleration of demand-sensitive cross-wholesaler stock management. Additionally, enhanced synchronization of region-sensitive cross-wholesaler product transport logistics is key to achieving consequent cost, CO2 emission, and total production-distribution time reductions. Most importantly for food wholesalers, unsold product use-by date expiration and consequent disposal losses are also minimized.

6.2. Supplier Collaboration and Risk Mitigation

Formal cooperation between suppliers is a common, though frequent-informal mechanism aimed at optimizing supply chains. Such cooperation can help to mitigate risks associated with potential demand shocks and to exploit economies and development of scale. Akin to demand forecast sharing, supplier collaboration can be more broadly indicated for visibility along the supply chain with respect to supply and production capacities. Using the available data from supply and transport chain partners, agent-enabled systems can identify potential production and capacity shortfalls and, using assessments of stakeholder risk-appetite, evaluate alternative supplier sources in the direct area or across national borders. Holistic analyses can validate the impact of strategic-sourcing considerations regarding supplier location. The presence of a supplier

network enables complex mitigation strategies for production sites in remote geographical regions.

Transportation and logistics are significant cost factors for any production company. For the food service sector, sustainability mandates underscore the need to minimize the transport and logistic footprints. Efficient transport planning is clearly supported by transport demand-supply matching, which can maintain the high fill-rates of delivery vehicles across all supply chain constituents. Assuring transportation capacities can enhance trade relations and offer more visibility while helping to limit freight price volatility. Using a dedicated store-locator database enriching address information, the network approach enables intelligent transport-logistics optimization including transport sourcing of drivers.

6.3. Transportation and Logistics Synchronization

The government and transportation agencies supplying national-scale disaster recovery also need support from the wholesale segment of the food service supply chain when standing up CDCs. During recovery from natural disasters, shifting commodity flows to temporary visual demand patterns through distribution centers in areas with low availability of food service product categories helps ease food insecurity. These centers typically transport food to local agencies that distribute supplies to residents in impacted communities. As needs continue to fluctuate, their timing and volumes must be closely collaborated on with multiplex relationships among suppliers and cross-segment freight service providers. In addition, some wholesalers can leverage efficient service via bonded carrier agreements with the port authority and other agencies.

The actual need for demand forecasting accuracy depends on the costs of storing commodity inventories and the marginal costs of procuring and transporting supplies in non-homogeneous demand environments. Minimizing such costs requires serious consideration of risk-adjusted measures, especially for non-homogeneous demand that can swing to near zero or negligible service levels—compared to the

annual total demand. Therefore, it is also important to establish optimal quantities of forward contracts for procurement and transportation, supplemented by FIFO dispatch together with the regular network to minimize total costs.

7. Evaluation and Validation

A multi-level evaluation and validation framework enables assessing components and their interconnections alone, as well as the impact of orchestration services on agentic decision-making. Several use-case drafts serve as assessment benchmarks, including a case study synthesizing demand for food service in Switzerland based on data from the Food and Agriculture Organization classification.

A coherent, structured flow of information within and across supply chains creates an instrument enabling actionable intelligence at all levels, from suppliers to service companies, consumers, and public authorities. Using seamless data orchestration across a diverse prime data landscape, supervisors can promptly detect trends in the development of food products, consumption, health and environmental issues. This capacity strongly contributes to policy and planning processes in the food sector—typically hard to model and predict—offering an essential supporting function for the actual stakeholders. At the same time, these elements enhance day-to-day decision-making and steering of the food service sector across its diverse and often conflicting components. Quantitative demand forecasting in food service reduces forecast error, allowing optimization of storage and supplies ITUC delivery and transport planning, furthermore supplier cooperation on risk mitigation with shared storage capacities.

7.1. Metrics and Evaluation Protocols

Quantifying the benefits of orchestrating and improving cloud-native supply-chain processes is challenging. Resilience, security, compliance, and data quality require trade-offs and are often expensive to implement; their

benefits can lie outside the enterprise and only materialize in crises. The complications are compounded in a federally orchestrated data mesh where validation relies on the logical coherence across multiple agents. A volume of events might come together to define a use-case domain. In this context, benefits can be exposed via indirect synthesis—integrating minority reports from different perspectives to expose common ground across otherwise distinct spaces of events supporting different use cases.

To illustrate, the culmination of demand forecasting and inventory optimization in periodic refresh cycles can be simulated. The growing divergence of inventory versus demand in prioritized regions reveals any emerging readiness for action. Enriched risk-temperature maps enrich demand, freshness, and monetary cost dimensions. By contrast, orchestrating multiple agents exposes synergies for suppliers in maintaining freshness, response capacity, or quality, while drivers and logistic raters rely on inverse, antagonistic risk maps. With integrated reviews of actionability, cross-supplier tangential routes, and alternative logistic Arrangements emerging as real-world suggestions, achieving friendly road-transportation conditions become actionable beneath the authentic authority required.





Fig 4: Metrics and Evaluation Protocols



7.2. Case Study Synthesis

The requirements support the design of a reference architecture for national-scale agentic AI-enabled big data orchestration of cloud-native supply-chain optimization in national food service wholesale products. Scalability, trust, security, privacy, and resilience guide the choice of a federated data mesh design and associated federated data-governance framework. It supports the AI-enabled analysis of demand patterns, inventory, food safety, compliance and quality, supplier capabilities, and logistics, while minimizing risks and formulating collaborative action plans.

Key drivers of agentic decision-making and orchestration of actionable intelligence include orchestration automata that simplify the specification of complex user-defined workflows to govern data and metadata movement and transformation along the supply chain. Allegiance to agent interest groups formalizes trust-reputation-based decision-making among agents in the absence of an independent trusted third party. Market mechanisms ensure capacity allocation and payment accountability among agents. Real-time situational assessment enables AI-enabled scoping of supply-chain risk zones based on actionable intelligence such as supplier alerts and natural disaster warnings.

8. Conclusion

Agentic AI-enabled orchestration facilitates responsive, data-driven decision-making and knowledge generation across the entire food-service wholesale industry by ingesting, processing, and analyzing heterogeneous data from a wide variety of sources. Collective or joint activities, such as demand forecasting or supplier collaboration, benefit from coordinated efforts of connected partners. By defining common approaches, sharing expertise, or harmonizing operational practices, organizations can enhance the quality of their decisions and mitigate risks.

Three specific use cases demonstrate how agentic AI-enabled orchestration facilitates active optimization in a national food-service wholesale context. Mechanisms for


demand forecasting and inventory optimization, supplier risk mitigation and collaboration, and synchronous transportation and logistics represent generic categories for development. These joint tasks involve supply-side and demand-side stakeholders integrating their data and expertise with a focus on either supply or demand. Together, they illustrate how responsive decision networks, powered by connected supply-side and demand-side agents, can support food safety and quality beyond the data capabilities of individual agencies.

8.1. Algorithmic Transparency and Accountability

Agentic AI-enabled orchestration of cloud-native supply-chain optimization in national food service wholesale food service systems support industry requirements of trust, compliance, and transparency

Cloud-native architectures federate proprietary and third-party cloud and on-premises services into a seamless digital ecosystem. Agentic AI-enabled orchestration of cloud-native supply-chain optimization in national food service wholesale food service systems support industry requirements of trust, compliance, and transparency. These include regulations on data ethics, privacy, and protection, Quality Assurance Program (QAP) Testing Orders from Health Canada, and the Qualité Supérieure de la viande (QSV) program of the ministère de l'Agriculture, des Pêcheries et de l'Alimentation du Québec (MAPAQ). The latter two institutions govern: within and outside Québec, respectively—the quality of ground beef served in Canada's schools. These standards enhance demand-side confidence, incentivizing meal purchases and optimizing ingredient use.

Algorithmic transparency and accountability in Big Data systems require an adequate model lifecycle. The data ingestion layer employs a modern data fabric that provides self-service data-engineering capabilities, reducing complexity and cost. An ingestion framework promotes consistent data preparation for analytics, utilizing a futuristic data abstraction that integrates and harmonizes the processes supporting Big Data pipelines. A PaaS offering incorporates



service-discovery repositories utilized at runtime by the orchestrators. Orchestration mechanisms and workflow abstractions formalize the implementation of synchronization and coupling within supply-chain datasets, supporting the effective generation of actionable intelligence.


8.2. Data Governance and Vendor Lock-In

The agents governing the national food service wholesale supply chain must ensure that their decisions remain local but optimized on a national scale, answering the core challenge of preserving local decision ownership while benefiting from a nation-wide optimization process. These two opposing requirements can create a risky situation when an agent position is assigned to private, for-profit organizations due to the inherent need to maximize their profit. If the design also includes a methodology for vendor lock-in, maintaining a data-driven service-oriented approach during the decision process, non-local data can easily become concentrated with the same provider, from both the offer and demand sides.

The risk relates to possible vendor lock-in situations in the demand side consolidation by the agents' commercial organizations. Forecasting models need to access data from suppliers to guarantee availability; therefore, demand-side models should not be concentrated with any single vendor but rather orchestrated as an agnostic collaboration guaranteeing data from every supplier to all demand agents. Private business operation naturally optimizes services offered integrating the best models; however, it is critical to guarantee demand-side models' data access and avoid concentration in order to mitigate contract portfolio risk and optimize redundancy.

9. References

- [1] Abdel-Basset, M., Mohamed, R., Elhoseny, M., Chiclana, F., & Zaied, A. N. H. (2021). Intelligent decision support for supply chain management under uncertainty: A systematic review. *Expert Systems with Applications*, 170, 114517.
- [2] Akkaya, K., Guvenc, I., Aygun, R., Pala, N., & Kadri, A. (2020). IoT-based smart supply chain management: A review and future directions. *IEEE Internet of Things Journal*, 7(5), 3956–3970.
- [3] Al-Hawari, T., Aqlan, F., & Al-Zu'bi, Z. (2021). Big data analytics-driven supply chain resilience: A systematic literature review. *International Journal of Production Research*, 59(14), 4211–4232.
- [4] Alqahtani, A., Gupta, A., & Nakashima, K. (2022). Cloud-native analytics architectures for scalable supply chain decision-making. *Future Generation Computer Systems*, 128, 161–176.
- [5] Anand, A., Sharma, S., & Bhatia, M. (2022). Intelligent workflow orchestration in cloud-native environments: A survey. *Journal of Cloud Computing*, 11(1), 92.
- [6] Angappa Gunasekaran, A., Yusuf, Y. Y., Adeleye, E. O., & Papadopoulos, T. (2020). Agile manufacturing practices: The role of big data and AI in supply chain competitiveness. *International Journal of Production Research*, 58(8), 2373–2390.



[7] Arora, S., Sanyal, S., & Saha, S. (2022). Explainable AI in demand forecasting: A review and framework for supply chain applications. *Computers & Industrial Engineering*, 170, 108312.

[8] Baryannis, G., Dani, S., & Antoniou, G. (2019). Predictive analytics and artificial intelligence in supply chain management: Review and implications. *Computers & Industrial Engineering*, 137, 106024.

[9] Ben-Daya, M., Hassini, E., & Bahrour, Z. (2019). Internet of Things and supply chain management: A literature review. *International Journal of Production Research*, 57(15–16), 4719–4742.

[10] Bhosale, V. A., & Yadav, S. (2023). Big data-driven supply chain analytics: Architectures, tools, and applications. *Journal of Big Data*, 10(1), 44.

[11] Bojarski, P., & Pękała, P. (2020). Data pipelines for industrial analytics: Orchestration patterns and reliability. *Software: Practice and Experience*, 50(12), 2225–2245.

[12] Bousdekis, A., Lepenioti, K., Apostolou, D., & Mentzas, G. (2021). A review of data-driven decision-making in supply chain management: Techniques and applications. *Computers & Industrial Engineering*, 154, 107158.

[13] Brintrup, A., Wang, Y., & Tiwari, A. (2020). Supply chain digitalization: Agent-based and AI approaches for adaptive logistics. *International Journal of Production Research*, 58(7), 2096–2118.

[14] Cao, G., Duan, Y., & Li, G. (2015). Linking business analytics to decision making effectiveness: A path model analysis. *Decision Support Systems*, 74, 70–80.


[15] Carvalho, H., Barroso, A. P., Machado, V. H., Azevedo, S., & Cruz-Machado, V. (2012). Supply chain redesign for resilience using simulation. *Computers & Industrial Engineering*, 62(1), 329–341.

[16] Chai, J., & Ngai, E. W. T. (2020). Deep learning for predicting retail demand: A review and future research agenda. *Decision Support Systems*, 137, 113316.

[17] Chen, X., Zhou, H., & Li, Z. (2022). Cloud-edge collaboration for real-time supply chain optimization: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 24(3), 1569–1602.

[18] Chopra, S., & Meindl, P. (2022). *Supply chain management: Strategy, planning, and operation* (8th ed.). Pearson.

[19] Christodoulou, P., Christodoulou, K., & Andreou, A. S. (2021). Reinforcement learning for



inventory optimization: A systematic literature review. *International Journal of Production Research*, 59(14), 4233–4257.

[20] Corbett, C. J., & Klassen, R. D. (2006). Extending the horizons: Environmental excellence as key to improving operations. *Manufacturing & Service Operations Management*, 8(1), 5–22.

[21] Dalal, R., & Mali, A. (2024). A stacking ensemble model for food demand forecasting to reduce waste. *Results in Engineering*, 21, 101032.

[22] Dam, H. K., Tran, T., & Ghose, A. (2019). Explainable software analytics: A survey. *ACM Computing Surveys*, 52(4), 1–36.

[23] Dey, A., LaGuardia, T., & Srinivasan, S. (2022). Data lineage and governance for cloud-native big data pipelines. *IEEE Access*, 10, 43455–43472.

[24] Ding, Y., Jin, M., Li, S., & Feng, D. (2020). Smart logistics based on IoT and AI: Challenges and opportunities. *IEEE Transactions on Industrial Informatics*, 16(8), 5009–5020.

[25] Dolgui, A., Ivanov, D., & Sokolov, B. (2018). Ripple effect in supply chains: An analysis and recent literature. *International Journal of Production Research*, 56(1–2), 414–430.

[26] Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of big data. *International Journal of Information Management*, 48, 63–71.

[27] Erol, S., Sencer, S., & Sari, R. (2021). A review of ML-based demand forecasting in food and grocery supply chains. *Computers & Industrial Engineering*, 162, 107747.


[28] Fahimnia, B., Jabbarzadeh, A., & Sarkis, J. (2018). Supply chain sustainability: A review and future agenda. *Transportation Research Part E: Logistics and Transportation Review*, 118, 395–407.

[29] Fakhimi, M., Kwon, K., & Yoon, S. (2023). Kubernetes-native data orchestration patterns for large-scale analytics. *Future Internet*, 15(2), 58.

[30] Fountas, S., Espejo-Garcia, B., Kasimati, A., Mylonas, N., & Darra, N. (2020). The future of digital agriculture: Big data and AI in supply chains. *Computers and Electronics in Agriculture*, 175, 105598.

[31] Fu, S., Tang, Y., & Zhang, Y. (2022). A survey of cloud-native workflow scheduling and orchestration. *ACM Computing Surveys*, 55(7), 1–36.

[32] Gharehgozli, A., Zanjirani Farahani, R., & Dullaert, W. (2020). Green logistics and



perishable supply chains: Optimization models and analytics. *Transportation Research Part E*, 136, 101877.

[33] Gong, Y., Janssen, M., & van Hillegersberg, J. (2019). Blockchain in supply chains: A systematic review. *International Journal of Production Research*, 57(18), 5523–5546.

[34] Govindan, K., Mina, H., & Alavi, B. (2020). A decision support framework for resilient supply chains using AI and analytics. *European Journal of Operational Research*, 283(3), 906–927.

[35] Guo, X., Li, C., & Zhang, X. (2022). Federated learning in supply chain analytics: Privacy-preserving demand forecasting and optimization. *IEEE Transactions on Engineering Management*, 69(6), 3212–3226.

[36] Gupta, A., & Kumar, V. (2023). MLOps architecture patterns for enterprise-scale supply chain AI. *IEEE Access*, 11, 112211–112232.

[37] Harrison, T. P., & van Hoek, R. (2020). *Logistics management and strategy* (6th ed.). Pearson.

[38] He, Z., Zheng, X., & Zhou, J. (2021). Deep reinforcement learning in supply chain management: A review. *Computers & Industrial Engineering*, 158, 107406.

[39] Helo, P., & Hao, Y. (2022). Digital twin-driven supply chain planning: A systematic review. *International Journal of Production Research*, 60(9), 2689–2710.

[40] Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of resilience in supply chains: Quantitative models and analytics. *International Journal of Production Research*, 57(13), 4191–4213.