



IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY



VOLUME : 11 ISSUE : 05 Print / Issue Publication Date: September 2026



ISSN : 2455-2143



Indexed In



WWW.IJEAST.COM

editor@ijeast.com



BUILDING RESILIENT FOOD SUPPLY CHAINS: STRATEGIES FOR MANAGING DISRUPTIONS IN THE CONSUMER PRODUCTS INDUSTRY

Asha Kiran Ganesna
Chief Architect
MN-USA

Abstract— Food and consumer packaged goods (CPG) supply chains are highly sensitive to public consumption and structurally fragile: critical, low-margin, easily perishable, and vulnerable to demand fluctuations, supplier failures, transport delays, and climate change impacts. With the pandemic, geopolitical conflict, and the growing rise of extreme weather events becoming the building blocks of future resilience, siloed, transaction-only enterprise resource planning (ERP) systems will no longer suffice. This paper suggests the following four-layer reference architecture where SAP is used as the transactional system of record, Google Cloud Platform (GCP) as the single and governed data foundation, Vertex AI as the predictive and prescriptive intelligence layer, and Palantir Foundry as the ontology-driven decision and control-tower interface. The architecture is applied to four concrete resilience strategies: demand sensing, supplier risk management, logistics visibility and inventory, and production adaptations, and is based on several established supply chain resilience theories, such as Christopher and Peck's resilience principles, Pettit, Fiksel, and Croxton's vulnerability-capability framework, Simchi-Levi et al.'s time-to-recover and time-to-survive metrics, and Ivanov's ripple-effect and digital-twin concepts. To support researchers and practitioners in the consumer products industry, an illustrative scenario, a framework for the measurable set of key performance indicators, an implementation roadmap, governance considerations, and limitations are presented.

Keywords— Supply Chain Resilience; Food Supply Chain; Digital Twin; SAP; Google Cloud; Vertex AI; Palantir Foundry

I. INTRODUCTION

There is special significance in the food supply chain within global value chains. Interruptions to food availability become quickly apparent as a public concern, a political issue, and, in severe cases, as a crisis of food security, as compared to the disruptions that affect durable goods. Food is a basic human need and, as such, emotionally and politically sensitive; it also has a higher profile than other consumer sectors. It also brings

the issue of food security to the forefront of attention much quicker than other sectors if there is a perceived threat to food supply. The COVID-19 pandemic created a sudden scarcity of critical inputs, production delays and unavailability, and sudden price increases in large parts of the economy, as revealed by a survey of senior leaders in dozens of North American consumer goods companies, subjecting food and household-goods supply chains to a sudden and highly visible stress test [13].

Over the years, a series of events has tested the industry, including demand surges during the pandemic and factory shutdowns, strikes in meat processing and fresh-produce facilities, transportation and port congestion, and the constant occurrence of extreme weather and geopolitical events that impact commodity availability. In each of these episodes, a single disruption can trigger events at other nodes in a multi-layer network, changing its structure and performance far beyond the initial impact on the failed node. The impact of each of these disruptions has been to showcase what supply chain scholars call the ripple effect; a disruption that initially affects one node of a multi-tier network will impact other nodes and cause other disruptions that will change the structure and performance of the network at a much larger scale than the first disruption [10], [11]. In the past four decades, numerous years of research in academia have centered on the topic of supply chain resiliency, much of which has been theoretical and conceptual, providing taxonomies of vulnerabilities and capabilities but lacking a specific, actionable blueprint of technology. Industry has been moving at a pace: Hyperscale cloud platforms now have reached a stage of maturity where resilience can be truly predictive, prescriptive, and, for a good proportion, automated—and this is now a realistic prospect for large consumer-product companies.

This paper focuses on the ensuing gap between the theory of resilience and technology implementation. It proposes a four-layer reference architecture based on four platform categories that are already available in production use at large consumer products and retail companies, used individually or in combination: SAP as the transactional backbone that forms the foundation for master data and core planning processes; the Google Cloud Platform (GCP) as an integrated data



foundation that is now a no-cost option for connecting to SAP, through zero-copy data-sharing capabilities announced by SAP and Google Cloud in 2025; Vertex AI as the machine-learning layer where demand forecasting, anomaly detection, and prescriptive optimization reside; and Palantir Foundry as the ontology-based operational layer that transforms unified data and model outputs into a usable, no-code form of a digital-twin control tower for supply-chain decision-makers [14]. This is a technical architecture paper and so is not about a new algorithm—it's a paper that draws on the literature to position what is known as these four categories of technology in relation to each other to form a coherent deployable system of resilience, with deployments illustrated throughout in the context of food and beverage and other consumer-product applications.

The rest of the paper is structured as follows. In Section 2, the theoretical underpinning of supply chain resilience is discussed, and the proposed architectural structure is positioned in the existing literature. Section 3 establishes the scope and context of food & CPG supply chain disruptions. The four-layer reference architecture is described in detail in Section 4. In Section 5, based on a concrete analysis of methods, the architecture is mapped onto concrete resilience measures. This section explains the following scenario of disruption. Section 7 proposes an evaluation framework of measurable outcomes. A governance model and implementation roadmap are provided in Section 8. Section 9 discusses limitations and future research directions, and Section 10 concludes.

II. THEORETICAL BACKGROUND AND RELATED WORK

A. Supply Chain Resilience: From Vulnerability to Capability

Rooted in Christopher and Peck's 2004 publication, *The Resilient Supply Chain*, the concepts used to build supply chain resilience (SCR) can be summarized in four principles: reengineering the supply chain, collaboration across partners, the need for agility in responding, and creating a risk-management culture [2]. Sheffi went further, saying that resilient companies intentionally built in redundancy, flexibility, and a proper organizational culture, as they have found that organizational, and not only technical, differences between a quick recovery and no recovery after a disruption were quite common [3]. Ponomarev and Holcomb then gave a new definition of resilience as an “adaptive capacity,” which they argued lets firms anticipate, adjust to, and keep operating despite disruptive events [4]. Besides, Pettit, Fiksel, and Croxton combined these elements into a single whole that they called “balanced resilience”—the right balance between the vulnerabilities that put a firm at risk and the capabilities that equip it to predict and resist those vulnerabilities—and ultimately translated this idea into the Supply Chain Resilience Assessment and Management (SCRAM) tool [5].

A common thread in this literature is that resilience capacities—visibility, velocity, collaboration, and redundancy—are as much an information problem as a physical or inventory problem. There is no way to reengineer, to collaborate, or to act quickly and decisively when a disruption comes that isn't seen or forecasted or evaluated quickly by a firm. This observation inspires a technology/architecture view of resilience, which is the perspective taken in the current paper: The vocabulary for the resilience capability of Christopher and Peck, Sheffi, and Pettit et al. is not considered simply an organizational checklist but rather a set of design requirements for implementing an explicit data and decision architecture [2], [6].

B. Measuring Resilience: Time-to-Recover and Time-to-Survive

While the frameworks listed above are mostly qualitative, Simchi-Levi, Schmidt, and Wei came up with a quantitative pair of metrics that have become influential both in academia and in practice: time-to-recover (TTR) is the time it takes a specific node in the supply chain, such as a supplier facility, distribution center, or transportation hub, to resume full operation after a disruption, and time-to-survive (TTS) is the longest time a firm can operate so that a particular node in the supply chain, such as a supplier plant or distribution center, can meet customer demand following a disruption [7]. The framework has been developed and validated in conjunction with Ford Motor Company and gives a simple but powerful decision rule: a node with $TTS > TTR$ exposes the firm to limited risk because other buffers in the system can provide service during the recovery period; a node with $TTR > TTS$ is structurally vulnerable and needs to be addressed through some action, such as multiple supply sources, safety stocks, or capacity flexibility. Subsequent research by Simchi-Levi and Simchi-Levi further developed the idea to suggest that critical supply chains should also be subjected to frequent “stress tests” along the lines of financial-sector stress testing [8]. Wieland also distinguishes between the “bounce-back” and an “ecological/evolutionary” “bounce-forward” outlook on resilience, where the latter focuses on adapting to changing conditions and becoming more resilient over time, as opposed to only being restored faster. Wieland's proposed architecture is good for both views of resilience, since it focuses on adapting system conditions, model-driven, in response to changes, as well as re-optimization, which is a somewhat more advanced form of faster recovery [9].

C. Ripple Effect and Digital Twins

Ivanov, Sokolov, and Dolgui described a localized disruption propagating through a multi-tier supply network and changing its structure and parameters well beyond the point of origin as the ripple effect [10]. On the same note, Ivanov and Dolgui then claimed that observing ripple effects, preemptive simulation of disruptions, and supporting real-time replanning



in case of disruption are among the main applications of the digital supply chain twins, which are live, data-driven replications of a supply network's actual status [11]. The overarching vision of a digital twin, which is designed to be a sensing and stress-testing device, is a direct inspiration for the ontology-based, simulation-ready layer, described in section 4.

D. Food and CPG-Specific Vulnerabilities and COVID-19 Lessons

Food supply chains have vulnerabilities that are shared to a lesser degree with other groups of consumers. In reflecting on the first year of the COVID-19 pandemic in Canada and the United States, Hobbs recognized that overall, agrifood supply chains did particularly well and that food trade across the border has continued to operate, though with the largest interruptions in specific nodes—within the meat processing industry and in fruit and vegetable production—as well as that the nature and ubiquity of food as a basic need makes any disruption of food supply far more politically relevant than an equivalent disruption of a durable goods industry [12]. An FTC review of the U.S. grocery supply chain that happened earlier in the pandemic reached the same conclusion at a wider level, noting that the shutdown of factories and ports during this period exacerbated shortages of critical food inputs that the FTC witnessed in consumers' shopping carts, which was evident in their food prices.

One of the most interesting insights provided by Burgos and Ivanov's digital-twin simulation study of German food retail during the pandemic is that the impact of the surge in demand and supplier shutdowns was the largest measured by the company, with the chaos caused by bullwhip-like inventory-ordering dynamics also a key reason for the sharp rises in

transportation costs [1]. This finding is important, as it implies that the most important resilience investments in the food retail sector are not just those in transportation tracking but also accurate demand sensing and supplier risk visibility, which will be leveraged directly in this paper (see Sections 4 and 5). The complementary activities conducted for short food chains during the pandemic further underline three elements of food chain resilience, namely, readiness, responsiveness, and recovery.

E. The Gap Between Frameworks and Architecture

Current literature on supply chain resilience focuses mainly on capacity, taxonomy, or simulation-based studies and utilizes digital twins (DE) to model disruption scenarios, but fails to clearly define an enterprise technology stack that is deployable to scale and can implement the underlying theory. In contrast, vendor and industry documentation for specific platforms (SAP, Google Cloud, and Palantir) exhibits point capabilities in isolation but rarely maps them with the academic usage of the vocabulary of resilience theory presented above. This paper attempts to redress that balance in the rest of the paper.

III. PROBLEM FORMULATION: THE DISRUPTION LANDSCAPE IN FOOD AND CPG SUPPLY CHAINS

A. A Taxonomy of Disruption Types

The five broad categories summarized in Table 1 are disruptions that can occur in food and CPG supply chains. This taxonomy combines the vulnerability factors identified by Christopher and Peck and Pettit et al. and the evidence reviewed in Section 2.4 for each food [2], [6].

Table- 1 Categories of Disruptions That Can Occur in Food and CPG Supply Chains

Category	Representative Examples	Primary Resilience Lever Implicated
Demand-side	Panic buying, foodservice-to-retail channel shift, promotion-driven bullwhip amplification	Demand sensing/forecasting
Supply-side	Raw-material shortage, labor outbreak or shortage, single-source supplier failure, commodity price shock	Supplier risk visibility, sourcing flexibility
Logistics	Port congestion, carrier capacity constraints, cold-chain breakdown, border or customs delay	Network visibility, routing flexibility

B. Why Food and CPG Chains Are Structurally Exposed

The food and CPG supply chains are structurally more prone to such disruptions when compared with other industries. Inventory perishability contrasts with durable-goods categories and supports Burgos and Ivanov's finding that inventory-ordering dynamics themselves become a source of cost and instability during a shock [1]. Thin margins restrict the investment of capital and exacerbate the classic trade-offs of efficiency vs. resilience that Christopher and Peck consider key to resilience strategy [2]. Multi-level, cross-border supply chains expose planners to the cascading effect of propagation, because a loss of a tier-two or tier-three supplier will not become apparent to planners until it manifests itself in the

finished goods. Last but not least, the number of SKUs is high, promotions are common, and there is a stream of new products added to the shelves, which often makes forecasting a challenge, such as predicting demand for new products that have little or no sales history, a problem known as a "cold start."

C. The Data and Visibility Problem as the Binding Constraint.

For all disruption categories in Table 1, the primary constraint to slow or inadequate response is that data about suppliers, production, logistics, and demand are scattered among enterprise resource planning systems, warehouse and

transportation management systems, supplier portals, Internet-of-Things sensors, etc., and external data like the weather or news feeds. As industry analyses of the challenges AI-driven supply chain transformation poses explain, the first problem is that this data lives in distinct systems that were built without the ability to interoperate; the second is that a lasting resilience solution lies in first bringing this data together before predictive and prescriptive can be added on top. The five design requirements build from this diagnosis and become the design logic for the architecture in Section 4 below: (1) Real-time, unified data visibility across transactional and external sources; (2) Predictive signals (not just reporting that's backward-looking); (3) Prescriptive, scenario-simulation capability to explore response options before committing to them; (4) Fast, governed decision execution that completes the loop back into transactional systems; and (5) Continuous, closed-loop learning—continually optimize the architecture for each disruption cycle.

IV. PROPOSED REFERENCE ARCHITECTURE

A. Design Rationale and Overview

It is seen here that the architecture is intentionally layered, where each layer deals with one or more of the five design requirements outlined in Section 3.3. The system is organized around four layers, as shown in Figure 1: there is a transactional backbone defined by SAP, a unified, governed data foundation provided by Google Cloud Platform, a layer of predictive and prescriptive intelligence called Vertex AI, and a layer of ontology-based orchestration and decision interface called Palantir Foundry that supply chain planners, procurement and category managers, and operations leadership directly consume. It is crucial to stress that this is not an SAP replacement; in fact, it's a solution that sits on top of SAP. Migration away from an existing ERP system comes with a significant price tag and risk and is typically infeasible simply because of resilience, as it is usually warranted on a broader operational basis. Rather, the architecture has two layers built on top of SAP: a sensing, intelligence, and orchestration layer and a system of record for transactions and master data, SAP.




Fig. 1. Four-layer reference architecture for resilient food and CPG supply chains.

B. Layer 1 — SAP: Transactional and Master Data Backbone

The central transactional data used for procurement, production planning, inventory movements, and sales order processing are stored in SAP S/4HANA and legacy ECC environments, with

SAP Integrated Business Planning (IBP) offering native sales-and-operations planning capability. Every additional level of the architecture of the supply network ultimately references the master data, such as materials, bills of material, and plant and supplier data records, upon which the single version of the truth



of the physical and organizational supply network is based. A major recent development has been SAP Business Data Cloud (BDC), a fully managed, software-as-a-service data platform that provides integration with third-party data for more comprehensive and robust decision-making while unifying and governing SAP data [14], [15]. SAP introduced BDC commercially in small numbers in early 2025, with it available in more widespread GA later in that year. However, BDC serves as the starting point (or endpoint) for the cross-cloud, zero-copy connectivity detailed in the next layer.

C. Layer 2 — Google Cloud Platform: Unified, Governed Data Foundation

At SAP Connect 2025, SAP and Google Cloud announced SAP Business Data Cloud Connect for Google BigQuery, a bidirectional, zero-copy data fabric that provides secure, direct access to curated SAP data products from BigQuery without duplicating data through conventional extract-transform-load (ETL) pipelines, with general availability planned for the first half of 2026 [14], [15]. A key benefit of this zero-copy pattern for resilience architecture is that it maintains SAP's business semantics and hierarchies when the data becomes available for analytical and artificial intelligence applications and helps eliminate the latency and staleness risk that is present in periodic batch replication, which is more of an advantage in the context of real-time disruption sensing vs. retrospective reporting.

The BigQuery analytical core integrates unified SAP data with external food- and CPG-relevant signals such as point-of-sale data from retail partners, weather and climate data relevant to agriculture and cold chains, Google Maps geospatial data and traffic data streaming through pipelines to BigQuery for cold-chain and logistics-lane risk assessment, Google Trends and search-behavior signals as early indicators of demand shifts, and Internet-of-Things telemetry from cold-chain sensors or agriculture and production sensors ingested through Pub/Sub and Dataflow streaming pipelines. Dataplex and the Knowledge Catalog offer governance, lineage tracking, and access control for this blended dataset—a material prerequisite following food-safety requirements and multi-jurisdiction data-residency rules—and have been deliberately designed to abide by regional data-residency laws across markets, including Singapore, Mumbai, and Sydney, making these cases the proof of the pattern's viability for multinationals with a strong presence in multiple regions. This layer satisfies design requirement (1)—unified, real-time visibility—and provides the governed substrate that requirement (2), predictive modeling, depends upon.

D. Layer 3 — Vertex AI: Predictive and Prescriptive Intelligence

Google has been continuing to develop its forecasting backend, such as the TimeSeries Dense Encoder (TiDE) architecture recently added to enhance the accuracy and efficiency of forecasting massive, sparse retail datasets—a use case that

Google says has helped today's leading retail brands with forecasting in hundreds of stores and vast product offerings. Notably, the ability to create forecasts for an item that has limited or no sales history is an important consideration for CPG, known as the cold-start problem. Vertex AI Forecast is being embedded in third-party retail and CPG-specific solutions focused on reducing forecast risk, increasing financial performance, and improving supply chain resilience at optimal cost and inventory [16]. A recent example is a major home-improvement retailer's deployment of Vertex AI-based demand forecasting across more than 200 stores, showing that it works at a meaningful scale.

The same single-source-of-truth GCP data is also used for anomaly detection and supplier or transport-lane risk scoring: Anomalies that result from delivery-lane fluctuations, quality-defect rates, financial-health metrics, and geopolitical or news signals can be detected earlier as a risk indicator affecting a supplier or transport lane, rather than waiting for the shortages to be reported downstream.

A second prescriptive-optimization capability, which integrates not only the Vertex AI custom models but also operations-research solvers, can be used for inventory allocation, production sequencing, and network reconfiguration in the event of disrupted capacities, and combines optimization logic originally formalized by Simchi-Levi and colleagues at Ford Motor Company for supplier-node risk assessment [22]. Lastly, a generative model based on Google's Gemini models can convert complex multi-signal forecasts and risk scores into narrative alerts and reasoning for planners, removing the barrier to interpretability and making human decision-making more efficient in the middle of rapidly evolving disruptions like a pandemic outbreak.

E. Layer 4 - Palantir Foundry: Ontology, Simulation, and Decision Interface

Palantir Foundry imports SAP- and GCP-based data into a live, object-oriented digital-twin model of plants, stock-keeping units, suppliers, distribution centers, and customers, and eliminates the need to query ERP tables or use no-code applications. In a consumer products-enabled, documented deployment among the Fortune 100, Foundry's ERP Suite was integrated with over seven SAP data sources to build a digital twin of the value chain from supplier to customer, with an out-of-the-box bill-of-materials workflow process enabling purchasing teams to start optimizing raw-material purchase and spot-buy substitution opportunities within one week of deployment, showing how rapidly an ontology-based layer can transform previously siloed ERP data into actionable decisions [17], [18].

Foundry's Object Explorer and Contour and Workshop tools let users explore, analyze, and build applications directly on top of the ontology without needing to code, while Foundry's AI platform (AIP) extends scenario-simulating functionality to the same digital twin—in one scenario, simulating the consequences of one of the key ports closing and instantly

identifying affected shipments, alternative routing, and the implications for cost and delivery time, for example, or a sudden surge of demand, the impact it will have on inventory, production, and supplier capacity [20]. These capabilities are often combined into a supply chain control tower: A dynamic, comprehensive operating picture that is not just a dashboard but actively identifies deviations and brings up suggested corrective steps. Independent implementation reports show that implementing visibility and suggested actions, not just reporting, resulted in a 50 percent drop in out-of-stocks for a global retailer.

Concrete deployments in the food and beverage sector demonstrate its relevance for the food sector in particular. Palantir's work with General Mills extends the use of Foundry and AIP to logistics and freight optimization, waste reduction, supply-chain risk management, sales-and-operations planning, and sourcing and procurement functions powered by AI [21]. By using Foundry AIP, Aramark achieves about a ninety-nine percent degree of confidence in the matching process on a small sample of point-of-sale product records that previously required significant manual effort for the Aramark reconciliation team to match. Independently reported, case data for a food manufacturer revealed that it has been able to lower inventory costs by 15 percent and that it has improved product freshness by adapting its production schedules based on AI forecasts of demand, while for a bakery, it cut customer complaints by 40 percent after introducing AI-driven visual

quality control [19]. This layer satisfies design requirements (3) and (4)—prescriptive scenario simulation and fast, governed decision execution—and provides the human interface that keeps a person in the decision loop for consequential actions, consistent with Christopher and Peck's emphasis on risk-management culture: the technology accelerates and informs organizational judgment rather than replacing it [2].

F. The Closed-Loop Resilience Cycle

Figure 2 shows the operational architecture as a continuous 6-stage loop, with disruption and demand signals being received through the GCP ingestion, then unified with SAP's transactional backbone via the zero-copy fabric, followed by predictions of demand, risk, and anomalies via the Vertex AI layer, simulation of response options in the Foundry's digital-twin ontology, and finally recommendations to take one or more actions—whether an adjusted purchase order, reallocated stock, rerouted shipment, or a reset of the safety-stock value—via the control-tower interface before executing and writing accepted decisions back to SAP—the authoritative system of record. The closed-loop nature of both parts of the Simchi-Levi framework—reducing the sense-to-act cycle time reduces the time a firm has for effective time-to-recover in the event of a disruption; the predictive and prescriptive elements provide additional options for substitution and buffer, increasing a firm's time-to-survive before a disruption is detected within its transactional data.




Fig. 2. Sense-unify-predict-simulate-decide-act resilience loop



G. Integration Patterns and Governance

There are two kinds of integration patterns in the architecture. In addition, while most non-streaming pipelines like ETL are adequate for most high-volume, non-SAP sources like IoT telemetry and point-of-sale feeds, the SAP BDC Connect zero-copy pattern is preferred in environments where duplicated, stale data is not permitted and where reducing integration risk and cost is important. Identity and access should be federated in a consistent manner throughout SAP, GCP, and Foundry, and role-based access should be consistent with food-safety and financial-controls requirements. Getting to Layer 3's

consequential decisions that are automatically write-back eligible—model governance (explainability, versioning, and continuous monitoring for drift) is not something you can simply overlook, as decisions with this magnitude of impact on food available to eat and supplier relationships are sensitive to auditability. Lastly, the architecture had to be compliant with data-residency laws across the multinational global presence often seen in large CPG companies, and the underlying SAP–Google Cloud zero-copy fabric had to be designed with this in mind.

Table-2 Layer Responsibilities in the Reference Architecture

Layer	Primary Technology	Design Requirement(s) Served
1. Transactional backbone	SAP S/4HANA, ECC, IBP, Business Data Cloud (BDC)	System of record; master data integrity
2. Unified data foundation	BigQuery, BDC Connect (zero-copy), Pub/Sub, Dataflow, Dataplex	(1) Real-time unified visibility
3. Predictive & prescriptive intelligence	Vertex AI Forecast (AutoML/TiDE), anomaly & risk scoring, OR-based optimization, Gemini	(2) Predictive signals; (5) continuous learning

V. MAPPING ARCHITECTURE

A. Demand-Side Resilience: AI-Driven Demand Sensing

The use of promotional signals, pricing information, weather, and trend signals directly tackles what Burgos and Ivanov found to be a major source of the bullwhip effect in food retail during the pandemic—operational and cost drivers [1]. The ability to anticipate fluctuations in demand, even for new products that have limited history, means planners are less likely to reorder reactively, as the same study identified reordering as a bigger number one source of disruption impact than transportation congestion. It's reliably used outside of a lab, like production deployments at a retailer of more than two hundred stores, using the Vertex AI built-in Vertex AI Forecast to forecast stock levels, as well as dedicated CPG-focused planning-platform integrations with Vertex AI Forecast.

B. Supply-Side Resilience: Supplier Risk Scoring and Sourcing Flexibility

The Foundry ontology, populated with risk scores derived from Vertex AI, enables planners to determine which suppliers to look for bottlenecks in, rather than by intuition or by seldom manually reviewing risk scores. The consumer-products bill-of-materials and cost-of-goods-sold optimization example in Section 4.5 illustrates the speed with which this capability can be realized in real, concrete sourcing decisions, such as quickly evaluating the feasibility of spot-buy substitutions for product formulation, existing stock levels, and forecasted demand.

C. Logistics Resilience: Control-Tower Visibility and Rerouting

Control-tower, what-if simulation of a port closure or transportation disruption operationalizes Ivanov's ripple-effect stress-testing concept as a production tool rather than a purely academic exercise, allowing logistics teams to identify affected shipments and alternative routing options before committing resources, and to quantify the cost and delivery-time consequences of each option in advance of a decision.

D. Production and Inventory Resilience: Dynamic Safety Stock and Waste Reduction

AI-adjustable production scheduling analyzes the perishability constraint identified in Section 3.2 as a structural constraint that necessitates a new buffer-stock strategy, introducing an adjustment that balances freshness constraints with demand forecasts, provided by independent food-manufacturing case data. AI-based visual quality-inspection systems have the potential to process much larger product volumes than manual sampling, and save product samples from customer complaints—this is another resilience aspect that is unique to food and beverage. Whilst many of the other vulnerabilities discussed in this section have a close analogue in manufacturing, food-safety risk is an entirely different breed for which there is little or no overlap in existing technology.

E. Organizational Resilience: Culture, Governance, and Human-in-the-Loop Decisioning

The organizational aspect of resilience, as noted by Christopher and Peck (risk-management culture) and Sheffi (supportive corporate culture), is not replaceable by the on-screen or technical aspects listed above [2], [3]. Independent implementation guidance on the matter of AI-powered operational efficiency in food and beverage settings clearly



states that improving business outcomes rather than technical aspects like system uptime is the basis for the success of a deployment, and that it takes time for teams to understand how their individual tasks impact downstream partners in the newly

exposed system. So, the architecture is more correctly seen as a layer for decision support and orchestration to keep a person in the loop to make consequential decisions, to speed up and inform the person's judgment, not supplant it.

Table -3 Mapping of resilience strategies to architecture capabilities

Resilience Strategy	Enabling Architecture Capability	Supporting Evidence
Demand sensing	Vertex AI Forecast (AutoML/TiDE) with exogenous features	Retail deployment across 200+ stores; CPG planning-platform integration
Supplier risk management	Foundry ontology + Vertex AI risk scoring	Fortune 100 CPG BOM/COGS optimization within one week
Logistics rerouting	Foundry AIP scenario simulation / control tower	Port-closure what-if simulation; ~50% out-of-stock reduction case

VI. MAPPING ARCHITECTURE

For purposes of making the architecture tangible, imagine a hypothetical middle-sized CPG producer called NutriFoods International that produces shelf-stable snack product items and has a single supplier plant for a major specialty oil ingredient. Planners can pre-emptively test this supplier's node-level logic (time-to-recover and time-to-survive) based on the Foundry ontology's assumptions: for NutriFoods, its estimated time-to-survive is about nine days, since it has a single plant with an estimated finished-goods and work-in-progress buffer, while its historical and industry data estimates that twenty-one days will pass before the incident is fully addressed (time-to-recover). The time-to-survive is less than the time-to-recover, so this supplier node is flagged as a structural vulnerability during the regular/routine scenario simulation—before a fire, flood, or labor incident—not afterward.

The closed loop in Section 4.6 helps speed up and simplify the response to any disruption—such as a fire at the supplier's plant that stops production—that moves from a multi-day process of exchanging e-mails and spreadsheets to making hours seem more like minutes. Pre-existing knowledge of the ontology relationship between this supplier and all downstream plants, SKUs, and all customers, combined with news and supplier-portal signals ingested through the GCP layer, enables Foundry to identify impacted products and customer commitments instantly before the disruption reaches the point-of-sale or ERP days later. Vertex AI also reforecasts at-risk demand and scores the available alternative suppliers with the same risk-scoring ability as described in Section 4.4. Foundry AIP then simultaneously executes a variety of what-if scenarios, each with estimated cost, lead time, and service-level impact, including partial product reformulation,

qualification of alternate suppliers, and temporary rationalization of lowest-velocity SKUs impacted. In the control tower, a planner scans through these options, chooses a combination of actions, and the purchase orders are updated, the safety-stock parameters are adjusted, inventory is reallocated across distribution centers, and within hours it goes back into SAP.

In the absence of this architecture, any disruption would typically be discovered when it became evident in inventory reporting or downstream sales; cross-functional coordination between procurement, planning, and operations would take place via email and phone calls, pushing discovery of such disruption even further out in time and delaying the coordinated response even more. The most significant contribution to resilience of the architecture in this scenario is the compression of the latency between detecting the underlying supply risk and taking action to ensure the firm's continued effectiveness, which directly impacts the firm's effective exposure in light of its time-to-recover; the systematic identification of suppliers with substitution options that push the firm's effective time-to-survive further must be considered as a contribution to resilience as well.

VII. EVALUATION FRAMEWORK: MEASURING RESILIENCE IMPACT

The reliability of a resilience architecture depends on the results that it can deliver. Table 4 suggests some KPIs across various engineering-based metrics of resilience presented in Section 2.2 and operational metrics more relevant to food and CPG supply chains in particular, along with benchmark values that have been used in various deployments shared throughout this paper.

Table-4 Evaluation framework — resilience KPIs and benchmark evidence

KPI	Definition	Repeated Benchmark Evidence
Time-to-recover (TTR)	Days for a disrupted node to return to full functionality	Framework validated at Ford Motor Company
Time-to-survive (TTS)	Days a firm can sustain service if a given node is disrupted	Node-level supplier stress-testing
Forecast accuracy (WAPE/MAPE)	Weighted or mean absolute percentage forecast error	Improved accuracy across 450+ hypermarket stores (Vertex AI case)

These figures should be interpreted as directional technical and operational feasibility rather than as generalizable causal effect sizes, as not all are necessarily peer-reviewed and controlled, but rather come from vendor case studies and from industry reporting, and may be subject to selection bias by virtue of the fact that they are likely to represent successful and well-resourced implementations. The same caveat as a limitation of the present analysis is repeated in section 9. Nonetheless, taken together, they show that there is a measurable operational improvement for each layer of the architecture associated with their respective deployments, at

least in part from deployments that resemble the reference architecture proposed here.

VIII. EVALUATION FRAMEWORK: MEASURING RESILIENCE IMPACT

Deploying the architecture in a real organization is best sequenced in phases that build a data foundation before intelligence, and intelligence before broad decision-automation so that each phase can be validated against the KPIs in Table 4 before the next is funded. Table 5 outlines an indicative roadmap.

Table – 5 Implementation roadmap by phase

Phase	Indicative Timeframe	Key Activities
1. Data foundation	0–6 months	Establish SAP BDC / BDC Connect or equivalent pipeline to BigQuery; data governance and cataloging; identity federation across SAP and GCP
2. Predictive pilots	6–12 months	Vertex AI Forecast pilot on one or two priority product categories; supplier anomaly/risk-scoring pilot on top-tier suppliers
3. Ontology & control tower	12–18 months	Build Foundry ontology for priority plants, SKUs, and suppliers; deploy control tower to planning teams; begin AIP scenario simulation

Governance mechanisms, along with delivery, are essential to a successful implementation of this roadmap. Prioritization and resolving trade-offs between the four vendor platforms should be the responsibility of a cross-functional steering committee that includes information technology, supply chain, procurement, and food safety/quality. Again, as discussed in organizational-resilience section 5.5, training and change management should be thought of as a work stream of equal status to the architecture, and not an add-on. Having to coordinate contracts, security, costs, and other issues with four different vendors creates an overhead burden which must be proactively managed, not simply ignored. Requirements for explainability and audit trails should be proportionate to the consequences of providing the information to inform that decision – a food-safety-relevant recommendation or an automatic purchase order to a supplier should be materially more closely examined than a dashboard metric. Finally, there should be a strong focus on costs: implementing such a multi-vendor stack has some costs involved both in terms of integration and licensing, and each step should be evaluated based on the improvement in the resilience KPIs outlined in section 7 rather than on the technology for its own sake.

IX. LIMITATIONS AND FUTURE RESEARCH

This paper makes a conceptual and technical-architecture contribution and not an empirically validated deployment study. The figures used in Sections 4, 5, and 7 are based on vendor documentation and industry case studies, and are not rigorous or peer-reviewed assessments, and may reflect some selection bias to projects where implementation has been successful and well resourced, and the reader should not interpret them as an effect size for any project.

Architecture is also vendor-specific, with a focus on SAP and Google Cloud, Vertex AI, and Palantir Foundry. While this design pattern is designed to be generalized beyond these specific vendors, future research would be needed to evaluate this design pattern against other technology stacks, such as Databricks-based, Microsoft Fabric-based, or AWS-based. Likewise, this paper does not thoroughly test generalizability across CPG subsectors: fresh/perishable products, business-to-business channels, and highly regulated products in the nutrition-adjacent pharmaceutical sector all put different pressures on the buffer-stock and substitution strategies outlined in Sections 3 and 6.

Finally, empirical pilot studies with well-quantified before-and-after comparisons of the KPIs proposed in Table 4, ideally



using a control group or a staggered deployment design, will be worth pursuing—as well as using a formal cost-benefit, or total-cost-of-ownership, analysis for the KPIs and four-layer stack proposed here, compared with single vendor or narrower alternative architectures.

X. CONCLUSION

Best known for their high public stakes, thin operating spreads, and multi-tier, global sourcing, packaged goods and food supply chains are regularly hit by demand, supply, logistics, and macro-environmental shocks that generally hit their customers—read: the public—before in most industries. This paper has introduced a four-layer reference architecture: SAP as transactional backplane, Google Cloud as data unification layer, Vertex AI as predictive and prescriptive intelligence layer, and Palantir Foundry as the ontology-driven decision interface layer, which operationalizes existing theory of resilience, including the principles of Christopher and Peck, Pettit, Fiksel, and Croxton's vulnerability-capability framework [2], [5]. Time-to-recover and time-to-survive metrics of Simchi-Levi et al. and ripple-effect and digital-twin concepts of Ivanov, and is a deployable technology pattern that has been documented across a variety of food, beverage, and broader consumer-product deployments [7], [10].

At every turn, the paper has stressed that this technology is not intended to supplant organizational judgment or culture but is designed to shorten the time cycle from sense to act to determine effective time-to-recover and to bring preemptive stress-test insights to the surface to extend effective time-to-survive, rather than replacing the human planner in consequential decision-making. The CPG technology and supply-chain leaders who are considering investing in a technology-based approach to resilience will find the following architecture, evaluation framework, and roadmap to be a practically grounded starting point and one that brings decades of research in resilience and literature to a point where it directly addresses the technology choices that supply-chain leaders are faced with today.

XI. REFERENCE

- [1]. Burgos, D. and Ivanov, D. (2021). Food Retail Supply Chain Resilience and the COVID-19 Pandemic: A Digital Twin-Based Impact Analysis and Improvement Directions, (p.102412).
- [2]. Christopher, M. and Peck, H. (2004). Building the Resilient Supply Chain, doi: 10.1108/09574090410700275.
- [3]. Sheffi, Y. (2005). The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage.
- [4]. Ponomarev, S.Y. and Holcomb, M.C. (2009). Understanding the Concept of Supply Chain Resilience, doi: 10.1108/09574090910954873.
- [5]. Pettit, T.J., Fiksel, J., and Croxton, K.L. (2010). Ensuring Supply Chain Resilience: Development of a Conceptual Framework, doi: 10.1002/j.2158-1592.2010.tb00125.x.
- [6]. Pettit, T.J., Croxton, K.L., and Fiksel, J. (2013). Ensuring Supply Chain Resilience: Development and Implementation of an Assessment Tool, doi: 10.1111/jbl.12009.
- [7]. Simchi-Levi, D., Schmidt, W., and Wei, Y. (2014.). From Superstorms to Factory Fires: Managing Unpredictable Supply-Chain Disruptions.
- [8]. Simchi-Levi, D. and Simchi-Levi, E. (2020). We Need a Stress Test for Critical Supply Chains.
- [9]. Wieland, A. and Durach, C.F. (2021). Two Perspectives on Supply Chain Resilience, doi: 10.1111/jbl.12271.
- [10]. Ivanov, D., Sokolov, B., and Dolgui, A. (2014). The Ripple Effect in Supply Chains: Trade-off "Efficiency-Flexibility-Resilience" in Disruption Management, doi: 10.1080/00207543.2013.858836.
- [11]. Ivanov, D. and Dolgui, A. (2019). New Disruption Risk Management Perspectives in Supply Chains: Digital Twins, the Ripple Effect, and Resilience, (pp.337–342).
- [12]. Hobbs, J.E. (2021). Food Supply Chain Resilience and the COVID-19 Pandemic: What Have We Learned?, (pp.189–196).
- [13]. Xu, Z., Elomri, A., Kerbache, L., and El Omri, A. (2020). Impacts of COVID-19 on Global Supply Chains: Facts and Perspectives, doi: 10.1109/EMR.2020.3018420.
- [14]. Shah, S. and Bungert, J. (2026). SAP and Google Cloud Launch BDC Connect for BigQuery.
- [15]. Khan, I. and Gutmans, A. (2025). SAP and Google Cloud Redefine Enterprise AI With Zero-Copy Data for Google BigQuery.
- [16]. Martin, A. (2023). Vertex AI Forecasting.
- [17]. Palantir Technologies. (2023). Winning the Retail Battle: How Palantir Foundry Transforms Supply Chains.
- [18]. Palantir Technologies. (2026). Use case patterns: Operational process coordination.
- [19]. Aramark. (2024). Managing \$20B+ Purchasing & End-to-End Supply Chain
- [20]. Hamda, S.B. (2026). Palantir Foundry Vs AIP: Which AI Platform Actually Works in 2026?
- [21]. Johnston, L. (2024). General Mills Expanding Generative AI in Procurement Amid Shift to Always-On Supply Chain.
- [22]. Simchi-Levi, D. et al. (2015). Identifying Risks and Mitigating Disruptions in the Automotive Supply Chain, doi: 10.1287/inte.2015.0804.

IJEAST

INTERNATIONAL JOURNAL OF ENGINEERING APPLIED SCIENCE AND TECHNOLOGY



AUTHORS

 Asha Kiran Ganesna



ABOUT IJEAST

International journal of Engineering Applied Science and Technology (IJEAST) is a peer-reviewed, open access journal that publishes high - quality research paper in the field of Engineering, Applied Science and Technology.

IJEAST aims to provide a platform for researchers, academicians, and professionals to share their innovative ideas, research findings, and practical experiences with the global scientific community.

JOURNAL METRICS



H-INDEX

33



i10-INDEX

154



IJEAST
CITATIONS

9,647

(IJEAST CITATIONS)



PEER-REVIEWED
& OPEN ACCESS

PUBLISHED
MONTHLY



For more information, visit our website

www.ijeast.com



PEER REVIEWED

All submission are rigorously peer reviewed to ensure quality.



AUTHORS ARE OUR PRIORITY

We value our authors and recognize their contribution to the advancement of research and knowledge.



OPEN ACCESS

Free and unrestricted access to research for all.



GLOBAL REACH

Connecting researchers and professionals worldwide.



TIMELY PUBLICATION

We ensure a swift and efficient publication process.



editor@ijeast.com



www.ijeast.com



India



2455-2143