

Chapter 2: Architecting big data infrastructure for scalable insights in national food service and distribution networks

2.1. Introduction

The global food industry is estimated to reach \$8,994 billion by 2021. In developed countries, where the food industry contributes over 10% of the total GDP, a considerable amount of attention is devoted to ensuring safe and sustainable food. Furthermore, food safety and security vulnerabilities are raised as a transcultural issue that affects peoples' wealth and peace. Because of that, promoting safety and security of the food system has gained priority in discussing transnational governance in the future. As consumers' desire for high-value products increases, there is an increasing need for a systematic framework enabling efficient configuration, management, and coordination of improvement goals, entities, and collective actions in the food system. Reaching trusted quality in food service raises a number of complex technical challenges that concern transnational modeling of design, adaptation, and control of interactions in the food chain and markets. Irregular, unpredictable disturbances play substantial roles and particularly hard to address: transnationally distributed operational and managerial units implementing individual control actions based on local knowledge and data; proprietary, non-shared data and processes that render global knowledge mapping impractical; and intermittent, multi-resolution feedback signals arising from interlaced systems of multiple control loops. All of these result in a need for approximative methods that can simulate and analyze environmental performance of the larger system independently and in reasonable approximation.

There exist five main ways to leverage big data to benefit the food network. Firstly, creating transparency of data. Transparency of data drives transformation, increases

productivity and leads to informed decision making. For instance, providing transparency of food distributions could allow distributors to identify scale economies, competition effect, and even attributes of the enlisted food products. Food retailers and distributors could be brought together to exchange data on their service areas to engage in collaborative logistics initiatives. Secondly, enabling experimentation to identify anomalies, detect fraud and intelligence, and improve performance. Thirdly, micro-segmentation to customize actions. Fourthly, replacing/supporting decision making and data analyzing with automated algorithms. Finally, the exploitation of big data can lead to the innovating of a new business model, product, and service. The emergence of smart sensors and monitoring devices generates unprecedented amounts of data and illuminates new business opportunities for food networks.




Fig 2.1: Understanding Big Data Infrastructure

2.1.1. Background and Significance

The recent development of Big Data makes it possible for all food businesses to collect and analyze huge amounts of data about food product characteristics, merchandising handling, and food flows throughout the food network (Reuters, 2025; Revvence, 2025). In fact, there are unprecedented amounts of data about food networks available today. Different food supply chain actors have various forms of internal data on a wide range of topics (inventory, order intake, order fulfilment, point of sales, waste, etc.). In addition to that, multiple third-party data providers operate across the globe, offering a wide range of auxiliary data on food safety and quality, food tracing and logistics, price trajectories, and more. Researchers and government agencies have built several long-term databases on global food supply and demand, nutrition, and environmental statistics that are useful to examine trends and detect anomalies. For global applications, space satellites have recently become an important source of data on cropland, pastureland, and food production and prices. In the light of the recent outbreaks of output shocks and food

safety issues, the availability of so much data on food supply chains both complicates and facilitates decision-making.

2.2. Understanding Big Data

Big Data is often defined as data that is characterized by the 3 or 4 Vs, namely Volume, Variety, Velocity and Veracity. These characterize the types of data and data formats. However, it is important to note that this is a viewpoint considered from the data end. Big Data can also be expressed in another viewpoint where its properties are discussed in the context of value creation from data. This viewpoint allows common understanding among organizations and facilitates meaningful discussions regarding business needs and challenges. There are five main ways to leverage big data in the food network: Creating transparency: Transparency of data drives transformation, increases productivity and leads to informed decision making. Enabling experimentation to identify anomalies, detect fraud and improve performance: Big data needs to be collected, integrated and analyzed in real time to discover anomalies and fraud. Micro-segmentation to customize actions: Big data enables the definition of increasingly finer segments and precise marketing. Replacing/supporting decision making and data analyzing with automated algorithms: Big data analytics allows organisations to find unknown patterns in a time-efficient and cost-effective manner. Innovating new business models, products, and services: Using data provides new perspectives that can fuel innovation in food products and services.

Big data poses unique challenges to organizations (Ambilio, 2025; Intelligent Core, 2025; Real Business AI, 2025). These challenges can be categorized according to the data properties of the systems or environment in which they arise. The nuances of big data pose additional challenges that are not necessarily evident in traditional data analysis. These challenges can be considered as a variation and amplification of existing problems, although they may pose more difficult problems than those faced when traditional, small data, are used. In broad terms, there can be data challenges, processing challenges, algorithm challenges, architecture challenges and human challenges.

2.2.1. Definition and Characteristics

The imperatives of food safety, food security, and sustainability strongly shape the evolution of food service and distribution networks. These imperatives are driving the networking, centralization, and digitization of national food service and distribution networks (FSDNs), which are expected to integrate all food-related stakeholders and processes by adopting multi-regional food networks. Data-driven technologies (DDTs) are leveraged by the Federal Government for strategic governance of FSDNs through

analyzing, calibrating, simulating, and controlling all stakeholder /process behaviors. However, data collection, management, and analytics poses significant Grand Challenges (GCs) of integrated big data infrastructure, quality data, understanding and knowledge mining of complex data and knowledge representations in FSDNs. Addressing the GCs requires theoretical and methodological advances, with an emphasis on the broad types of data streams, modalities, and uncertainty resolving best practices in food service and distribution networks across spatial-temporal scales.

FSDNs are characterized by multi-regionality, spanning multiple states or countries, and by complex food process-to-process stakeholder relationships that span multiple supply chains in each region. With the integration of farmland, processing, transporting, storage, and food service technologies, all food-related processes are converted to digital Data Streams (DSs). Each individual DS type presents a significant DGC such as integrated infrastructure and services, continuously evolving at multiple evolution modes, and cross-DSs collaborative data-driven analytics and stakes. Multimodal (MM) data streams from DDTs are generated at millions of locations via heterogeneous integrated platforms with evolving formats, structures, and devices. Presenting vast GCs for data acquisition, management, and cost-effective analytics, due to their high volume, high variety, and high velocity (V3), and due to Milliseconds (ms) cross-variety/variety/stream processing needs. Furthermore, due to their inherent Missing Data (MD), Outlier (OL), Degree-of-uncertainty (DU) issues presented by spontaneous, naïve, and autonomous real-world sensor networks.

2.2.2. Types of Big Data

Big data can be categorized in various ways. Several experts have presented their categorizations in literature as a contribution towards better understanding of this concept. They are:

- **Structured Big Data:** A structured database is one which is arranged in a well defined manner allowing swift access.
- **Semi-Structured Big Data:** Data which does not follow a predefined structure falls under this category and is more flexible in nature. Semistructured files can be transformed to structured databases with additional processing.
- **Unstructured Big Data:** Data which has random arrangement and form falls under this category presenting challenges for extraction of meaningful insights. It may be present in textual, pictorial or video formats. Even on a single framework it can be in diverse languages and formats such as .pdf or .txt files.

On the other hand, they can be broken down into:

- Data from sensor and network devices.
- Streaming data.
- Unstructured or multi-format data.
- Offline bulk data.

There are diverse technologies and platforms available to handle big data. Some of them are:

- Hadoop ecosystem.
- MongoDB.
- Spark.
- Apache Kafka.
- Cassandra.
- R and

Weka for statistical and machine learning computations. • Neo4j and Gephi for graphical and network visualization along with analysis of scientifically formatted data. • Apache Flink and Mattermost can process and visualize streaming data.

2.3. Importance of Big Data in Food Service

The Global Food Safety Initiative (GFSI) holds immense significance for the food industry and the public. As key players in crucial sectors such as handling, production, processing, and retailing, the food industry's progress and evolution significantly impact both the quality of food products consumed worldwide and the business operating stories. Moving into an era of big data, the food industry is wizedened with intelligence from various operational levels compared to conventional years of experience. Specifically, a real-time digitized data information platform driven by big data sources offers an opportunity for players in the food industry to tap into valuable insights, enhancing overall speed and quality across ever-expanding food networks. Analysis and uncovering intelligence from these big food data sources then pave the way for a smarter and safer food supply chain. Serving large-scale data and providing comprehensive visualization, it could contribute to the bonding of food supply-chain players into an industry-standard solution space for smarter food safety intelligence, regulatory operations, and internal quality management.

The food supply chain (FSC) is important for ensuring food safety as it involves various processes from production to consumption. Food leaves local suppliers to move through large structured distributions across states, and various third-party logistics (3PLs) are involved in moving food products. With the current global situation and the Covid-19 pandemic, emphasis is focused on greater visibility and transparency across complicated food distribution networks to reduce the risks of food quality and safety concerns.

There are many participants in the food network, and considerable amounts of big data are generated for food safety monitoring, such as food source and distribution information. The current situation involving black-box food channels from states to restaurants provides challenges for regulators in maintaining the overall safety of food supply. There are visible demands and interests in developing and adopting a big food transparency system capable of sourcing diverse data across food networks to provide an integrated representation.

The most common analytical technologies are basic parameterization and reporting, but very few use highly predictive analytical tools or machine learning. It is also shared that there has been a significant increase in the use of BI within reporting tools and additional exploratory and complex analytical capabilities. Restaurant operators are currently

spending \$225 billion and projected to grow to 493 billion in 2030 on food consumption related to about 300 billion transactions.




Fig 2.2: Big Data in Food Industry

2.3.1. Market Trends

The estimate of the big data technology market in 2019 introduces food consumption. They can share machine learning and AI in it. Also explain how sensors, wearables, and IoT are Top-rated Technologies. In 2019, they also shared the projected growth of consumption related to food. The start of development includes information related to cash flow and customers. These data are investigated to develop data for food consumption, food delivery, and food waste. The report also shares government-recommended programs for food safety. Last, the proposal of big data in Food SYNC will be presented. In the era of big data in the food service and distribution domain, information related to food consumption, safety, technological development, waste, and money transactions is being created in big volume. However, the complete and deep utilization of such data is still found in most food service and distribution networks, which consist of multi-users and its information is unknown to the others. To visualize and scale insights into such information, a comprehensive and insight-configurable big data ecosystem called food SYNC is proposed. According to this market report, the big data technology market has already exceeded \$15.7 million in 2019 and it is expected to grow to \$37.2 million and \$67.5 million within the next 5 and 10 years. A part of this report shared best practices of big data in food service and distribution networks. They

shared the usage of BI and analytics by U.S. foodservice distributors. The results show that most common BI technologies are basic tables and reports.

2.3.2. Consumer Behavior Analysis

With the rapid advancement of Internet technology and the increasing popularity of social media networks, data is growing exponentially. For consumer insights-oriented consumer behavior analysis, aggregation, archiving, and examination/dissemination of consumer behavior data is of paramount importance, for which the social media platforms are most valuable, given their vast networked membership, comprehensive coverage of various social traits, styles, and themes, and ubiquitous dissemination of messages. This growing body of big data, unprecedented in size, complexity, and availability, has far-reaching applications for business intelligence, market capitalization, and social computing.

In the context of food service and distribution networks, several types of consumer behavior data are to be aggregated, over an integrated big data infrastructure and network, so as to generate consumer insights to be shared with the food service and distribution companies. First, fundamental social signals from social media platforms need to be archived. This includes people's online reviews and ratings, check-ins, and comments about food products or outlets in a spatio-temporal fashion. These are highly valuable for aesthetics-oriented consumer behavior analysis. Next, time-meta structured social signals from search engines and social media platforms need to be collected for query-based event-context consumer behavior analysis. The feeds provide satisfactory fine-grained consumptions, but around only a limited number of events. The simmering consumer appetites are exhibited by the online queries, which trigger avalanche accumulative social replies in the form of feeds.

The first batch of consumer insights comprises aesthetic-based consumer behavior analysis results of consumer characteristics, spatio-temporal exploratory analysis of food outlets, analysis of online reviews, detection of the rapidly decreasing food venues, and recommendations based on habit dollars. The second batch of insights concerns query-based event-context consumer behavior analysis results of evaluative consumer behavior analysis of event context on competing food outlets, sentiment analysis results, event-oriented consumer characteristic analysis results, and event context topic mining analysis results.

2.4. Current Challenges in Food Distribution Networks

Food distribution plays an important role in the cost efficiency and food safety of food supply chains. A food distribution network typically consists of multiple wholesale food distributors; each wholesaler purchases food directly from manufacturers or producers and sells it to a number of retailers in the surrounding areas. The retail stores place orders to their assigned distributors, which send trucks to deliver food orders to them. Although current distribution networks have achieved a high level of cost efficiency in a hierarchical way, they are nevertheless unable to provide up-to-date data and insights for new business scenarios due to historical implementation hurdles, manual data entry and aggregation, as well as technological limitations. Scalable and up-to-date analytics on food distribution networks beyond basic key performance indicators is a crucial enabling technology for building up a smart food supply chain. Despite the complexity of food distribution networks, existing food network studies assume static graphs and process-oriented analysis.

In contrast to these studies, an approach that builds up a Big Data infrastructure and analyses the dynamic, multi-dimensional, and multi-layer food distribution network for scalable insights is proposed. The data architecture and processing methods are designed and implemented. Typical analytic use cases are proposed and analysed. The design is applied to the daily food distribution operation of a national wholesale food distributor, and the associated network operation data is collected and presented. The current challenges for a Smart Food Distribution Network are also presented. Enabling technologies on both the Big Data infrastructure and the scalable insights are presented. Existing analytics on food distribution networks are based on heuristically designed structural approaches. However, due to the limited availability of a comprehensive picture of food supply chains or the lack of up-to-date data, the analysis methods, when implemented, cannot generalise to other use cases in practice.

2.4.1. Data Silos

Most large-scale websites are evolving to recognize the need and value of exposing their data for use by third parties. Until recently, sites divided their data across siloed databases and proprietary internal tools while relying on non-standardized web technologies and tactics to share data with other applications. Moreover, archived examples of siloed data environments are abundant in many organizations, including universities, governments, and even multinational corporations. The pressure to mine new value from existing data stores has never been greater. In post-ingestion months or years later, data products are often less valuable to the organization than to an outside partner with proprietary analytics skills. Siloed data makes it challenging to provide evidence of compliance with regulations across heavily regulated industries such as

finance, energy, transportation, trade, and accounting, some of which carry millions of dollars of fines for non-compliance, or within government for answering congressional inquiries and public scrutiny.

In virtually all organizations, multiple levels of risk must be considered whenever data moves out of company control from data owners to data scientists, resulting in a subset of labeled data. Data silos tend to fragment the database into the least amount of data necessary to support a particular process while protecting or disregarding data not needed for that process. Data providers must work closely with data scientists to select which attributes to expose to data science teams and prevent variables they guarantee not to allow for public inclusion. Data proliferation guarantees a pre-emptive filtering approach is needed. As an industry taskforce involving senior government officials, business leaders, and data scientists suggested, “the industry needs to improve its understanding of and ability to share and expose data rights, trade offs,” and responsibilities to support new data products and services. Stakeholder agreement of licensing terms is key. Data owners cannot exploit external data products to build their analytics on if owners don’t control data rights or that data is not licensed in an appropriate form.

2.4.2. Inefficient Resource Allocation

Challenges arising from inefficient resource allocation can be observed across a variety of sectors, including health care, logistics, financial services, marketing/advertising, etc. Resource allocation problems can broadly be divided into CAPACITY PLANNING which is concerned with resource dimensioning, and RESOURCE ALLOCATION which determines appropriate resource assignment across users/tasks/flows/applications. Typical performance metrics include end-to-end throughput and end-to-end response time.

Food waste has drawn increasing attention as a global issue prompting academic and practitioner efforts alike to route-out causes and reveal mitigation strategies. Avoiding food waste has consequently risen as a priority for most governments. Not only is this an agitative sustainability and climate change issue because of greenhouse gas emissions and loss of carbon sink, it is also a pressing macro-economy issue due to operation costs and government funding for distributing food security. For understanding the reasons behind food waste and developing sound policies for avoidance and waste reduction, it is fundamental to uncover the nature of on-going food waste on top of data collection. Appropriate and timely data collection helps get a better understanding of the drivers of food waste/reduction, but appropriate and timely data processing streamlines these strategies further.

Mistaken resource allocation is ubiquitous in fledgling food security and distribution systems (FSDS) with safety stock either on excess supply-side or excess demand-side, thus putting pressure on governments. Nonetheless, resource allocation strategy changes over food marketplace transitions. A dynamic yet holistic big data systems architecture is therefore necessary for FSDS. Data sources sharing and reuse keep beyond query reformulation paradigms. Query-centric extraction is yielded upon comprehensive transaction exploration. Attributes discovery across data sources provides learn-to-rank transformation for models in an FSDS. Multi-criteria tree representation guides greedy off-line sample and control display strategies.

2.5. Architectural Framework for Big Data Infrastructure

The increasing volume and velocity of data generated by IoT sensors, combined with the need for collaborative decision-making backed by insights and knowledge creation, require a paradigm shift in the way data, information, and knowledge are collected and processed for scalable insights. To shift from localized decision-making to a national food service and distribution network level analysis approach, many ongoing challenges are to be addressed on the capabilities of the national food service and distribution networks (NFSDN) in terms of requested data, cloud storage, and analytical capabilities. To address these issues, the architecture of national big data infrastructure for collaborative data processing and analytical tasks is described to support scalable insights in the NFSDN resilience assessment and knowledge generation tasks.

The proposed architectural framework is composed of systems and tools for cloud-based temporal data storage, deployment of distributed parallel analytic tasks, and automated analytical results visualizations for ingesting, processing, and generating knowledge from and for the national level food service and distribution networks as an integrated national big data infrastructure. The national big data infrastructure is described based on its functionalities and capabilities for generating knowledge, streamlining food service and distribution operations, food service and distribution aid support, and facilitating tracking, monitoring, and validation tasks. The technical feasibility of the proposed architecture is assessed in six dimensions of processing and analytical capabilities for the deployment of cloud-based data processing and analytics frameworks and an iterative top-down and bottom-up approach for generating additional knowledge and insights from the results.

Scalability, availability, storage, fault-tolerance, and privacy architecture are proposed to support the national food service and distribution network resilience assessment and foster the generation of additional knowledge and insights from the produced analytics as an integrated infrastructure perspective. The anticipated enhancements of the

proposed architecture are outlined for the near future upon the advancements in the food service and distribution information and knowledge domains.




Fig : Blockchain Technology in the Food Industry

2.5.1. Data Ingestion Layer

The data ingestion layer is designed to receive data from multiple data sources in real time or via batch operation. In this layer, the data source can be either streaming or batch data and stored in any formats like text files, JSON format, or even in a database. Once the data is ingested to the Hadoop File System (HDFS), an Apache Pulsar queue is used to balance the ingestion and parsing load, then make sure the incoming data stream is optimized for further analysis. Similar to the streaming ingestion, the batch ingestion layer fetches the data from HDFS and, again through a Pulsar queue, parses the raw files and saves it into intermediate AVRO files. The pre-processed data files are sent to a Flink job for further parsing into a Hive table. Notice that the batch ingestion is an operationally complicated task. Its sophistication comes from the assumption that the food service network is already functioning and the data files are generated starting from such time. Hence there are related schema files that record the parsing offsets and the IDs of the last parsed file. The respective staging process first checks whether the harvest schedule was missed on the day of ingestion. During the staging, the parsing progress will be accordingly updated in the schema files. Each increment in the progress will trigger the Flink parsing layer. The real-time ingestion mechanism designed to use Apache Pulsar which is a distributed and efficient publish-and-subscribe message queue. There are partitioning topics for each parsed level. For the top events tier, the most recent files in HDFS are unordered data streams. Each such file is assigned one or more partitions and is streamed in a free-stand variable number of workers. Each minute, the

workers will look at their assigned partitions to compute the next file and postpone it until the time is filled.

2.5.2. Data Storage Solutions

Coloring of the maps is based on the total annual value of food purchased by each supermarket chain-of-custody facility. While not represented in this mapping, secondary supplier distribution locations and retail store locations are integrated as points, size-scaled to total annual food purchased values. Significant portions of the nation's spending on domestic food purchases is handled by a relatively small number of connected firms.

Using a sample of 12 billion transaction records for over 19,000 stores in over 100 chains across the U.S., storage devices can provide hundreds of terabytes of input transaction records in an efficient and fault-tolerant format for analysis. Input text records can be represented with 4, 8, or 16-byte floating-point numbers, integer, or text sequences, depending upon the representation employed by the target algorithms. The compatible storage layer can generate multi-partition files, files, or flat text files with integrated new references. can directly use files generated by many existing tools.

While supporting multi-node clusters, it can be used to construct analytic views over all or parts of data to enable fast SQL- or DDL-based input data selection built on a SQL-based data analytics layer. This will be specifically designed to support nested data representations, various index formats, and sophisticated algorithms for narrow to wide analytics. The goal of the data analytics layer is two-fold: generate scalable and efficient algorithms that take advantage of distributed data processing and storage to analyze and query data; and package such algorithms with APIs in a way that can be called by processes written in different languages.

2.6. Conclusion

Reliable data management and governance strategies are crucial to the architecture of big data systems. Reliable data sources and management techniques are vital to a successful big data infrastructure. However, bad data leads to inaccurate results, unfounded decisions, failed implementation, and lost business opportunities. Thus, the blind date approach to big data may not work due to data reliability problems. A software automation technology called Reliable Data Services is proposed to turbo-boost the architecture of big data systems along the data management life cycle, including data sources, data ingestion, data validation, data processing, data governance, and data usage. Various data scientific techniques, tools, and the power of the computation cloud

are discussed. The solutions are driven by cost-efficient and user-friendly software automation designed for end-users rather than data scientists and aimed for effective and reliable poultry big data, especially for repurposed manufacturing data – untamed data directly exported from automations. Architecture and a prototype of RDS are presented focusing on data sources, data ingestion, data cleansing, data integration, and data governance with a pilot study on data. Enterprise big data solutions are inadequate for poultry networks due to small data teams, high user demand, and low data reliability.

Industry-based services for the infrastructure of big data are expected in that they are cost-effective and user-friendly. As an important supplement, a public service has been born to turbo-boost the insight discovery process of the poultry industry. All data scientific services are designed for chicken and egg networks to know their big data but not building big data systems. During nearly a year of operation, it has been found to be effective and popular. It is necessary to enable extensibility to cover many more domains and to build a community on applied big data together with professional organizations. The blue ocean of food networks is expected to be rapidly developed as various issues rely on big data and many more professional organizations will follow e-commerce and big data. While joining the win-win process, it is insisted to focus on what the poultry and food industry really want rather than a traditional scientific approach to ensure sustainability of services.

2.6.1. Emerging Trends

The National Food Service Distribution Network has emerged as one of the largest data generators in the quick service restaurant and food delivery industries. The substantial volume of data continually generated by retailers, manufacturers, and distributors presents an opportunity for scrutinizing efficacies in this supply chain. Exploring these data affords profound insights into food distribution applications, but it also raises serious challenges. Due to rapid growth in availability, data complexity, and emerging technologies, there has been an unprecedented shift toward utilizing big data analytics and ad-hoc approaches. However, advancements in computational functionality and data accessibility challenge decision-makers to decide the appropriate technological architecture configuration.

Big data shared across the supply chain has great potential. However, configuring an appropriate architecture that allows for scalability, accessibility, and data-driven decision-making remains an unsolved scientific challenge. This research aims to bridge this gap for the national food distribution industry by exploring methods to efficiently store, process, and analyze its distribution and ordering data. Emerging data strategies are identified, and then mapped architecture personas representing decisions taken by organizations in these industries are synthesized. A multistage approach is undertaken

to identify key drivers of performance for the architecture, which are modelled in the context of insights generated. Dynamic programming is applied to explore the design space and identify scalable architecture configurations based on the a priori prescribed utility of the decisions taken regarding data strategy and architecture persona. Machine learning approaches are utilized to explore key decision settings, providing insights into the impact of technological parameters, with a focus on parameter setting. A framework for devising an architectural configuration is finally presented.

Insights gained from analyzing big data from the national food network in the quick service restaurant sector are capable of facilitating better distribution practices, thereby saving substantial costs and improving customer satisfaction. Allowing these efficiencies to be realized, the ETL process developed concordantly with the pragmatic concerns of food distributors is crucial. By incorporating a framework for the “three Vs” of big data, the integration of processing capabilities is made compatible with the volume, variety, and velocity of data generated.

References

- Ambilio. (2025). Leveraging Agentic AI in Supply Chain: Top Successful Implementations. Retrieved from ambilio.com
- Intelligent Core. (2025). Agentic AI Transforms Banking Industry. Retrieved from intelligentcore.io
- Real Business AI. (2025). Building Resilient Supply Chains: Agentic AI for Real-Time Visibility and Agility. Retrieved from realbusiness.ai
- Reuters. (2025). AI agents have a clear mission, a hazy business model. Retrieved from reuters.com
- Revvance. (2025). Introduction: The Transformational Potential of Agentic AI in Banking. Retrieved from revvance.com