Empowering Retail Oss/Bss Platforms With Agentic Ai And Scalable Data Engineering

Shabrinath Motamary

Software/Systems Architect, motamaryshabrinath@gmail.com,
ORCID ID: 0009-0009-6540-7585

Abstract

Complexity is on the rise; therefore, automating as much as possible deterministic tasks is paramount to reduce error, workloads and burnout. Many tech companies built cloud services that cater ML ops and agents. However, the symmetry argument applies to data and analogously for enslaved data, the long tail from observability and decision making should be handled independently and differently than the data to one's advantage. Nevertheless, teach them to make decisions in a sound manner but never trust them to do so unaided. This vision is interesting to discuss, but how to achieve it at scale on all surfaces from a nebula of differently parametric seascapes? The solution lies in the embrace of the periphery. Real and rich for the most part, its populous aftermaths chipped away at by semiotic disambiguation techniques, co-universe construction, and the synergy of crontabs, data lakes, MLops and cloud services shape the inner veil. Underpinned by these constructions and ever approximated service competence, the gap is closed by announcing the availability of first class service descriptor layers and entry-points for rich insights and explorative analyses. Technically, and from a methodological standpoint, how to implement agentic or soft AI systems to improve OSS and consumable data? This loft ambition will show the plight and urgency of unshackling untapped data and how to achieve it.

Agent-based soft AI systems forge an entirely new class of system, enabling open, smart, agentic and defensible general purpose tools for thought and automation. Consumed enrichments, e.g. focused on making observables actionable, informed and optimal decisions, on engaging with data and OSS in novel micro, meso and macro ways. OSS covered in the canvas of consumable data, e.g. analytics, observability and AI/ML ops. Forward looking development principles and a design agenda to construct instance AIs. The grand challenge of and their method to exponentially scale up data analytic, OSS engendering, decision-making, and proactive pattern discovering activity using soft AI agent rendering in 4D. Techniques to build such beacons a community of co-creators and discourse from large language models with first principles cognitive architectures appended. It is part of a grand exploration on how to enable any system to extend their capabilities through soft AI. Moreover, here is the humane environment, enabling everyone to achieve their unique goals using large self supervised AI models. This is valuable; however, it excludes data and isn't scalable.

Keywords: Agentic AI, Retail OSS/BSS, Scalable Data Engineering, Telecom Automation, Intelligent Operations, AI-Driven Service Management, Customer Experience Optimization, Real-Time Data Pipelines, Next-Gen BSS Transformation, Autonomous Network

Management, AI-Augmented Decision Making, Data-Centric Architecture, Service Orchestration, Digital Twin for OSS, AI-Powered Business Support Systems.

1. Introduction

Retailing is a major pillar of the global economy. Ranging from small neighborhood stores to multinational chains operating thousands of stores, retailers employ millions of employees in every country and cater to billions of customers every day. But it is not just people that congregate at stores; merchandise, money, documents, data, and electronic signals all traverse the retail ecosystem every time a customer seeks to buy and acquire merchandise - goods or services. To an intelligent and savvy shopper, few neologisms are as problematic as managing an operation with so many diverse but intimately intertwined processes. In addition to the usual hard- and software that integrate merchandise, money, documents, data, and signals, a commercial organization implements an immense array of information and communication technologies. As a store operates uniquely and in relation to its specific location, so does every computer in a chain's data center, branch, and terminal. Why these systems are built, installed, upgraded, replaced, or dismantled by in-house programmers or purchased from platform developers, hardware manufacturers, and service vendors is too complicated to definitively capture in one single document.

That complexity is borne of understanding the reasons, requirements, and compromises shaping these processes, hardware, and software layers. This complexity appears as multiple forms of construction for every program the organization runs on a computer. Retail is unique: it shapes virtually unstorable merchandise, adding value as customers acquire goods or services and extract profits from the organization. Retail's hard- and software resemble a physical ecosystem with life forms of diverse habitats and specialities. Understanding that construction is important for the present state and future potential. It is also its Achilles heel. Paper systems were once common, laborious, and open to fraud. Then, bilaterally agreed securely encrypted dual-entry payment transaction records, simple to understand, were transformed by number-crunching bandwidth-multiplying clients into un-befoulable millions of rows in dozens of multi-billion-dollar systems of daily operation.

That data is available aggregated, transformed, and replicated into honking warehouses of detailed raw records, onto hugely altered transformed working inputs, and distilled info datasets flitting over supposedly pristine BIOS and Operating Systems into analytics platforms. Nevertheless, ad hoc and haphazardly accumulated labels of various species proliferate, rendering control center decision-making painfully tenuous. Chief Information Officers and users no longer understand the systems even after many interventions possibly made state-of-the-art by the vendors guiding the newsroom and boardroom. Parallel to the blossoming of unwieldiness and obsolescence, they grind to a halt. The time and effort taken once on simple blueprints and glibly explaining the verdicts are now a galaxy of disparate vendors, languages, legend, neutral documents, and technical manuals non-read by lay workers. If the embedded IT is too complicated for the professionals in a sizable global retailer owning a chain of stores with millions of entries in dozens of systems, loneliness rationalization might ring alarm bells.

1.1. Background and Significance Digital transformation is radically changing how retail companies engage with customers and with each other. Following the success of numerous digital platforms in the consumer-to-consumer and consumer-to-business spaces, large and small retailers now increasingly expect similar platforms for business-to-business engagements. In many cases, this requires that for onboarding large numbers of potentially adversarial partners to a common platform, competitive, tamper-resistant, and self-scaling architectures are created. In such architectures, third-party service providers can deploy containerized and container-orchestrated services in a secure, fair, and efficient manner. Partner-specific adapters, aggregators, and market mechanism components are composed and automatically deployed using algorithms that ensure truthful incentives and compliance to service-level agreements (SLAs) and real-time monitoring.

To foster the growth of the commercial open-source (OSS) ecosystem of these platforms, a transparent and portable ecosystem of procedural, platform-oriented, semantically annotated codes is created. This supporting infrastructure utilizes a permissions- and consumption-based access control built on top of a blockchain, which allows any participant to inspect the algorithms and to ensure compliance to SLAs while preventing the disclosure of sensitive business rules and data. Monitoring algorithms based on multi-agent systems are designed to operate in the same incentive-compatible manner as the market mechanism components. By doing so, the expectations motivate developers, data scientists, and operators to invent new monitoring algorithms on the global scale and share them on the micro-scale, while emerging agencies concerning compliance enforcement foster new trustworthy service providers, revenue streams, and job opportunities.

It is believed that with the right research and engineering investments, corporations and consortia of retailers, investors, and independent software vendors, main building blocks of such platforms can be commercially available as OSS in five years. It is essential to begin this long journey now to avoid a race to the bottom and succeed with a diffuse and inclusive ecosystem of a healthy market that incentivizes innovative competition for customer engagement.

Equ 1: Customer Experience Optimization (CXO).

$$\mathrm{CXO}_{score} = \frac{\sum_{i=1}^{n} \left(\mathrm{AI}_{assist}^{(i)} \cdot \mathrm{Response}_{time}^{-1} \cdot \mathrm{Satisfaction}_i \right)}{n}$$

- **AI_assist** = binary (1 if AI handled request, 0 if not)
- **Response_time** = time to resolve incident
- **Satisfaction** = customer feedback score

2. Understanding Agentic AI

AI's role in industries is often that of an Intelligent Assistant to create insights on business activity with a mix of automated reporting and queries. With Agentic AI, a mix of IAs and Agents serve as an Intelligent Control. High-level AI, once tuned, configures sub-IA or Agent AIs and un-tuned processing nodes. Management systems, using model abstractions with AI functionality descriptions, automate coordination for retail on this scale. Coordination occurs across a hierarchy of processing nodes. Each node packages its behavior as a process – a description of the AI Agent servicing business needs. The IC uses the "demand" view of processes as packages compiled from information model abstractions. Information models abstract "what" to talk about as entities, the things of corporate life, and decision elements; 1st-order predicates that question or assess "how" a setting is.

An additional specification using 2nd-order predicates specifies the processing agent's mode of behavior/change. Such persistent state model abstractions aid business, de-risking and smoothing change. Emergent behavior from Agents behaving on such models aids planning and evaluation. Processing agents monitor and facilitate business using comparison to model predicates across nodes. Unexpected change flags node(s) for intervention, assessed on relative change. Performance counter monitoring assesses planned coordinating interaction frequency changes targeting phase shifts in emergent behavior. Such targeted intervention to direct emergent behavior harnesses collaboration across all levels and ranges of complexity, from intra-agent change to co-instruction. Agentic processing thus allows delivery of retail's most complex, chaotic, and consequential coordinated behavior – run-all-nodes-with-everything-considered simulation of business activity in real time.

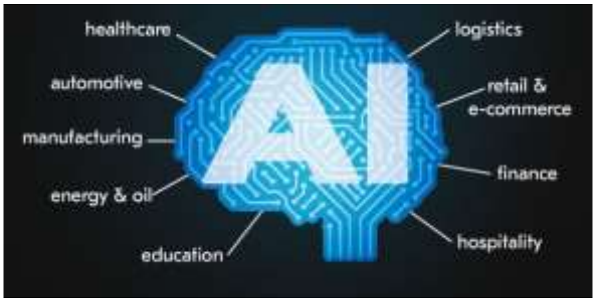

Fig 1: AI's role in industries.

Retail Agents possess high-data quality nodes for their data assurance and control. Node specification of these controls and processes includes parameters for data location, currency, trusts, and concealment levels. Co-processed with agent adjustments needed to meet a company's updated policies, they are integrated, tuned, and added to monitoring behaviors. Data/software provisioning nodal loading happens as updates occur. As trunks flow about a network, local aggregation nodes participate in monitoring for emergent control and assurance oversight, leveraging their inherent redundancy.

2.1. Definition and Key Concepts This section defines some key concepts that form the foundation of the discussed integrated technology stack such as agentic AI, retail OSS/BSS, and scalable data engineering.

Agentic AI is the ability of the AI system to act autonomously in the environment while achieving its objectives. Such systems must possess three key capabilities: Awareness, Reasoning, and Action. Awareness is the ability to collect observations of the environment's state in sufficient detail with high confidence. Reasoning is the ability to translate awareness into predictions and decisions while also demonstrating full transparency of the components responsible for these predictions. Action is the ability to execute decisions in the environment in a manner which is robust to observations and correctness of the underlying prediction.

A Retail OSS/BSS Platform is the confluence of multiple technology components that seamlessly integrate into an information and automation ecosystem. OSS refers to the operational systems required to run the daily operations of the retailer while BSS refers to the systems required to conduct billing, pricing, and invoicing related operations. Such platforms are commonly built on top of a mature middleware solution that includes a distributed stream processing engine, a scalable sink or no-SQL solution, graph databases, etc. Over time, additional supporting tools have been integrated into a cohesive platform such as ETL tools, product information management tools, etc.

Scalable Data Engineering refers to the family of components needed for supermarkets to ensure data provenance and compliance. This includes governance tools such as a Catalog, profiling and monitoring tools, and a data quality platform.

2.2. Applications in Retail Retail analytics has become an important topic area and business opportunity since the modern information technology (IT) era, with fierce competition between different retailers to increase sales and profits by better understanding customer behavior, market trends, etc. At the beginning stage, to have better analysis and planning, retailers primarily invested in internal data and transaction management systems after adopting an IT system for years. Recently, there are two main trends regarding the development of retail analytics. Apart from the development of traditional data management systems and operations research (OR) models, new solutions have been developed to address new data processing/analysis challenges that were previously ignored or unmeasured for a long time due to data scarcity. To serve the ever-increasing demand, massive amounts of new data have been generated every day from social networks, mobile devices, and advanced retail technologies, as well as newly available data sources via web crawling. New data-processing storage tools have been developed to capture, clean, and store data more quickly and at lower costs. Many companies are now able to access high volumes of real-time data. Advanced analytics has become possible using the latest machine learning (ML) algorithms and processing mercenaries burning low-cost computing power.

Retail analytics also faces many new challenges. The importance of social media analysis is much higher, as it is believed to significantly influence customer behavior and store performance. In addition, the physical world of retail is becoming much more complex. The integration of all channels: store, website, mobile etc., is necessary, and new interactions between the various societal agents sourcing goods, purchasing products, or analyzing/processing activities are developing. Furthermore, retailers not only serve customers but also communicate with them in an increasing number of ways. Omnichannel promotion decisions are also becoming much more complex. With the advancement of AI, combining data needs from different agents with multi-type data sources, new behavior modeling approaches and ML algorithms can be applied, and real-time analysis is feasible. To leverage the manifold potential of AI and big data, the retail ecosystem needs to change dramatically and create many new institutions, behaviors, and interactions. The process is still at its initial phase. However, it will likely be a great trend. Short-term effects on the economy will likely be great; the world will be forever different.

3. Data Engineering Fundamentals

Retail businesses today experience rapid shifts in consumer preferences and behaviors, business priorities, and ecosystem partnerships, with little time for adjustment. These shifts, driven largely by digital transformations, are no longer only external and externalized via instantaneous, broad media dissemination. Today's rapidly developing AI and data technologies, shifts in consumer enterprise-wide data management frameworks, rapidly digitizing supply chains, and shifts in value exchanges of data among potential new partners are catalysts for rapidly evolving shifts in retail. As these shifts develop at an unprecedented pace, many retail players cannot keep pace. The mass obsolescence of legacy OSS/BSS platforms prevents much of the potential competitiveness, enabling enterprises to be more consumer-centric and hence better consumers themselves.

Retail enterprise enablement is tractable only if each retail subsystem's enabling shallow platform can observe the data engineering and overall exploitation data expose rules of its source systems while being unaware of the deeper architecture of the source. Building off from decades of academic and corporate research in observing and automatically exploiting data governance and engineering rules in-player OS phenomenons, a constant-stream-as-is observing data feeds architecture of self-evolving enterprise platform components is introduced. Each tenant data source feeds as-is ensuring constant design-free and risk-free service delivery given data engineering compatibility of feeds and retail subsystem transparency. The most routine undertaking of data engineering and exploitation is essentially delivered on the platform-side. Therefore, the focus on speedy, low-investment, insurance-free construction of massively extensible shallow enterprise OSS/BSS platforms supports consumer centricity and increased competitiveness.

In today's enterprises, data monitoring and architecture across subsystems is opaque, access speed to meaningful views is low, and architecture is deep. Each consumer's enterprise deployed in-memory data mart replicas of its source data can sit on top of an auto-exposing platform-as-is observing feed agential data engineering and governance transparency ensuring data quarantine-free knowledge embedding beyond observant agentiality. In either settings, agential building and maintaining of shallow enterprise commodities encompassing the `know-how' of any OSS/BSS platform potential return all the delivered value of data engineering, governance, and knowledge embedding with speed, low investment, and resilience to exposure malfeasances all other platforms could ensure only in the consumer domain at unsupportworthy magnitudes. The inevitability of such a consumer-enable platform ensures the invalidity of today's exhaustive treatment of the pro-competitive potential of OSS/BSS platform-based competition. This renders observable the urgency of setting-up complementary data monitoring, engineering, and unsupervised vectored observations of potential surface platform partners.

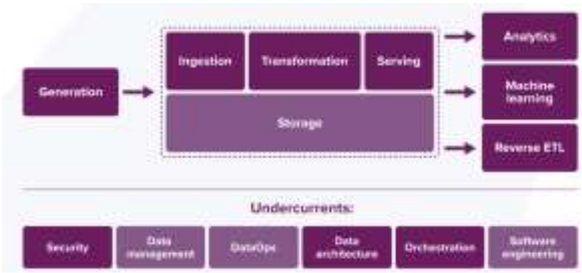

Fig 2: Fundamentals of Data Engineering.

3.1. Overview of Data Engineering This chapter investigates agentic AI and a framework to adapt retail OSS/BSS systems and core processes around pervasive AI tools and agents so core OSS/BSS applications become agentic and grow this new capability through a shared set of robust data pipelines. Agentic AI use cases and ones for future generations of consumer and enterprise applications are discussed in detail. Implementation methods that scan for local automation and agent opportunities across business, data and IT domains are reviewed, as well as those that allow for the launch of generative AI agents in consumer domains, and the creation of day-to-day productivity agents at scale. Some steps to safely and ethically adopting and growing agentic AI are also discussed, along with lessons learned from early adopters, and what the next 3−12 months to take advantage of the agentic AI opportunity should look like from a strategy, people and culture, and systems and processes perspective.

The term “data-centric AI” has taken off over the last 2–3 years, with many tools and platforms being introduced under this umbrella. Part of the growth is thanks to the fact that data-centric technology has existed in some form for at least two decades via the established fields of data, ML and AI engineering, which have been an essential piece of the data and ML ecosystem in production. This workshop aims to better understand what data-centric AI means, determine what platform capabilities and responsibilities fit under that umbrella, and discuss who’s building such platforms today or in the future. Inspired by growing interest in data-centric AI, many questions have arisen around the definition, meaning and scope of data-centric AI. Since it is such a broad topic, it was split into general discussions around the definition, platform categories, trends, etc., as well as discussions specific to feature stores, data validation, datasets, and data annotation.

Data-centric AI is meant to encompass data-centric technology in the ML and AI engineering contexts. Data-centric AI includes any data-centric technology and tool built for data-centric ML, AI, or both. There are also some data-centric data engineering tooling pieces with an ML/AI tie-in, but those are less common and less well known. Anything fitting broadly into the above definition, either currently or in theory, is permissible, and any opinions on what is and isn’t a fit are welcome.

3.2. Importance in Retail OSS/BSS Accurate billing and compliance with tariffs, particularly in relation to regulatory mandates, are critical components of an OSS/BSS solution. Such a solution should also ensure precision in cash collection activities, especially in relation to corporate and SME accounts. Prior to onboarding customers, upfront checks must also be performed. As part of onboarding customers for the retail segment, the solution must allow KYC checks of customers using semi-automated workflows. As a best practice, automated onboarding of digital customers, and semi-automated onboarding of non-digital customers, is also usually done. The solution must also allow the subscribers’ credit rating, bill presentment, and collections to be automated. Corporate accounts with innovative billing requirements will need to be accommodated inline with requirements. Customers that are over-turned must also be systematically and automatically passed to Dunning. The OSS on board in the selection must have adequate steps in the process to allow workflows that have been designed per requirements, to be automated systematically. All relevant processes must have adequate workflows to allow sales and marketing to function in near real-time. Overall, the digital channels should also be supplied with near real-time data. The solution should allow easy data management for new billing systems, and for network inflationary increases, and all such changes must be captured right down to subscriber level updates automatically. The OSS should allow all approved changes to be effective immediately or at least within the same business day. In addition to the above regulatory, billing, and compliance capabilities, any new OSS/BSS should also provide the business with enough capabilities to grow into new territories. Robust workflow automation capabilities are required to capture and design the complex business processes into manageable workflows. Key process focus areas are the ability to restrict, assign, and route workflows based on the required specific departments, teams, and/or skill sets, allowing for efficient work distribution linked to a guaranteed SLA turnaround time. The overarching outcome of a successful OSS/BSS replacement would include the following capabilities: (1) Real-time near zero TAT on onboarding, KYC, Compliance, Rating, Billing, and Collection, (2) Generalization of the entire OSS/BSS infrastructure across go-live countries though standardized software and hardware stacks, rules, and settings, and (3) Upscaling of operation volumes while holding the increased channels/teams of complimentary services operational TAT at current levels.

4. Integrating Agentic AI with OSS/BSS Platforms

Integrating agentic AI with OSS/BSS platforms in retail offers significant opportunities to maximize returns from investments in the channel stack and data infrastructure. Existing platform encapsulations, capabilities, and services need to be enhanced by AI, data, and analytic technologies to ensure swift alignment with future transaction volumes, signaling more intelligent systems design patterns than those leveraged to date. Data-driven investments can drive the creation of fully autonomous OS monitoring, management, and control domains, including the encapsulation of agentic AI across key OSS/BSS engineering phases and across associates and agents.

To enhance existing OSS/BSS platforms, it will be necessary to break unfulfilled frontiers across popular open-source platforms and proprietary commercial platforms. Low-friction interoperability APIs between these various platforms need to be designed in extensive and extensible development envelopes across modeling, data, simulation/optimization, and execution layers. These APIs will encapsulate agentic AI and scalable data engineering technologies. This platform-layer design will enable adaptive, intelligent, and automatable components to be developed, leveraged, and evolved via development APIs in a value ecosystem. Large-scale gauged data and pre-trained/training agentic AI encapsulations and chain designs also need to be developed with commercial-grade infrastructure and tooling APIs to improve accessibility and usability. These platforms can also integrate, encapsulate, and examine extensive datasets reflecting evolving consumer behaviors and preferences. Data and AI-savvy business users can build, deploy, experiment with, and monitor new data and/or AI-driven solutions to and across these OSS/BSS platforms without access to technology-specific know-how.

These underlying and commercial-engineered data and AI value ecosystems will collectively provide enhanced capabilities for retail channels and analytics. Digital, interactive, scalable, and extensible agentic AI and data engineering design and build shell wrappers for existing BSS platforms will provide initial core-level transformations of channel analytics, remediation, and control designs. Dramatic improvement in the functionality, usability, and relevance of channel quad-play offerings for consumer-facing networks will accelerate and prioritize these capabilities.

Fig 3: Integrating Agentic AI with OSS/BSS Platforms.

4.1. Architecture Design Leveraging OSS/BSS Data Throughout the Organization – OSS/BSS needs to drive Agentic AI throughout the organization, so that both analytical models and agentic applications are kept fresh. This requires a continuous loop of Data Collaboration – Ensuring data relevance, data freshness, data quality and data diversity is essential for a successful and prolonged Agentic AI project realization.

This involves two dimensions, OSS/BSS Take-off and Usage and Control. Control functions are feedback loops which render all take-off and usage functions relevant again. They require techniques from the disciplines of Data Engineering, Data Science and Data Governance. All relevant functions on all dimensions need to be in place.

A first key architecture design choice is defining the appropriate heterega data platform across the OSS/BSS as a platform to implement the collaboration functions. This needs to be implemented using functionality including datalines, data views and fields. It needs to allow access to data from analytics models and agentic applications in batch and real-time. It must include meta-data sharing as well. This requires a horizontal line of designs across the OSS/BSS platform.

A coherent heterega data platform can only be implemented in an integrated way across the OSS/BSS platform. This requires alignment of team organization, data product owner information, product plans, delivery processes, and related SCRUM boards and events. A clear intention with management is therefore essential as a next design choice. Coherency is very difficult to achieve in a team of teams organization with a strong emphasis on divergent development of siloed heterogeneous data products.

Additionally, alignment on precision budgets is key to properly assess and communicate possible inaccuracies for the developing predictive models and agentic applications. This requires first making this precision budgeting a hard requirement and second being able to proportionally allocate available data and storing resources across the OSS/BSS platform. Distributed delivery of heterogeneous data products will inherently lead to inability to align precision budgets. An integrated approach is therefore necessary.

4.2. Data Flow Management In this section, we focus on the raw data flow involved in on-boarding a new Retail chain. This process is elaborated step-by-step. The data flow is captured in a workflow, with the workflow being sub-divided into steps. Each processing step in the workflow is then further described with its input-output with the description of the tasks within it. A screening of the possible limitations of the data flow with respect to the envisaged operation is also provided.

The overall workflow for Retail chain on-boarding consists of four steps:

1. Data aggregation and storage.

2. Data augmentation involving the use of generative models.

3. Data cleansing.

4. Data upload to the ML model and TDBs assemblage.

These steps are realized by six processing steps, which are enumerated as follows: Pre-processing step 1, Data upload step 2, Data cleansing step 3, Data augmentation step 4, Data upload step 5, and Post-processing step 6.

To support the on-boarding of a new Retail chain, all available information should be uploaded to a dedicated data lake and split between the three raw bucket storages required by each service; i.e., AI service, graph-data base services, and collaborative filtering services. Data from different sources may have different formats and structures. Therefore, Data Source Management Big Data feeds will only parse metadata from the incoming files and store them to the manifest table. File contents will not be checked, but they should come in a supported structure. Furthermore, the Landed data will remain in the original data format.

Consequently, pre-processing step 1 is the collection of Retail chain relevant data from various files; data upload step 2 is the parsing of data only on the metadata level; the output as per pre-processing step 1 should only be to the Intermediate file bucket of each service. In addition to metadata upload to the manifest tables, contents should be uploaded to the new buckets. Per download there is a dedicated background process which on pre-specified time intervals checks whether data are already uploaded. If they are, nothing happens; but if not, a new data upload is triggered. The retrieved data at first undergo internal checks and afterwards should automatically be redirected to the Intermediate Data Store bucket. Internal checks can only determine conformity with minimum acquisition criteria and, if fulfilled, should allow regular processing and handle the cases of non-conformant data through the monitoring system.

It should be noted that the incoming batch size can vary, including smaller/collating batches. In the case of excessive batch size standard features can be used to stagger or batch-wise execute greedy feeds.

5. Scalable Data Engineering Techniques

Data engineering plays a key role in deployment-ready AI applications. Since there are hundreds of other valuable components, numerous techniques, and original benchmarks in data and ML engineering, it may be beneficial to highlight those that are most important for data-centric AI. In addition to this technical focus, there may also be broader cultural and social science insights that have not yet been applied to data-centric AI. In addition to reiterating some general data and ML engineering principles, those that are of particular relevance to DCAI will be provided. This includes interfaces, data processing as a service, and the ubiquity of pipelines.

Interfaces are key in decoupling software components, enabling horizontal scaling. Most component interfaces should be based on open protocols, avoiding tight coupling of components. Avoid assumptions about the format of data being passed between components. Instead, use open standards and formats so that at any time data can be inspected or replayed in other environments or components. Embed versioning into the data infrastructure. To scale at data velocity, there should be a service that can transform data in a reliable and scalable manner. This would enable pre-packaged transforms to be modular. Given diverse modelling approaches and frameworks, a specification service should be developed that simplifies the process of providing training data to these tools. In addition to exposing a simple API, it should provide a visualization and explanation service, enabling data scientists to understand how data transforms are shaping the training data, and why certain data is flagged.

Fig 4: Data Scalability in data engineering.

5.1. Big Data Technologies The data engineering landscape has seen a shift towards managing a multitude of scalable technologies to address large-scale data storage, processing and analytics. Initially, large-scale analytics relied on batch processing models within data warehouses. As data was generated and collected at ever-increasing volumes and rates, data-housing technologies evolved into data lakes capable of supporting larger and more heterogeneous storage. Technologies such as Hadoop HDFS and AWS S3 gained popularity. They provided object-based storage with access through file systems as well as APIs, enabling cheaper storage and global scalability. While the batch model is suitable for Infrastructure As A Service (IAAS) cloud deployments, it resulted in the evolution of new data structuring and processing technologies. These technologies were centred on streaming data that had to be processed for immediate usage due to their event-type nature, data-exploder methods, and internal data circularity. Technologies such as Apache Kafka, Spark Streaming, and Apache Flink have emerged, along with omission filtering and aggregation methods to filter out unnecessary bulk predictions through threshold actionings. Additionally, new monitoring and visualising technologies evolved into more prescribed rendering models to accompany huge storage and analytics infrastructures.

The era of massive data lakes and streaming is closing. Data is of high value only when continuously presented to engineered ML pipelines and acted upon. It presents a preconditioned state to the actors, is processed by a set of ensemble ML method predictions, and is followed by automated submit-actioning, influenced by real-time operational conditions. The responsibility for managing and processing the data is being placed on the data development enablers of the workflows, which will reduce complexities. At the same time, they will increase preconditioned volumes by higher re-engaged state exposures and the comorbidity of different types of data leading to a need for a growing repertoire of new ML methods. This trend will not be amenable to a U-shaped technological growth path but to a mentoring approach that will co produce the enablements with personal contacts and experiences. This transition will undermine the viability of many current predictions on the futures of ML and will shift resources and explorations to extreme development worlds outside the ML norms.

Equ 2: Agentic AI Action Probability.

$$P(a|s) = \frac{e^{Q(s,a)}}{\sum_{a'} e^{Q(s,a')}}$$

- $P(a|s)$ = probability of an agent taking action a in state s
- $Q(s, a)$ = expected return of taking action a in state s (from reinforcement learning)

  Used in OSS fault prediction or self-healing actions.

5.2. Cloud Solutions for Scalability Elascale Architectures provide scalability for both micro/macro service resources as well as monitoring and load forecasting. For the auto scaling of data engineering resources, Elascale expands the auto scaling architecture defined in current Cloud SaaS implementations by providing autoscaling for a complex stack that consists of open-source Data Engineering stacks as well as the application owned stack. On the complex stack to be scaled there are many unknowns including application-dependent components and important points affecting load prediction properties. The effort of the Providers can be a source of high performance auto scaling systems if made available to the SaaS as-a-service layer.

A stubbed autoscaling engine service mimicking auto-scalability of RUBiS faces the challenge of heavy workload burst. Even interacting with an elastic database over cloud, a huge storing time is found in initial adaptation. Complex cloud SaaS applications are inherited by the properties of cloud applications and other unknown properties including load and performance patterns and sensitive points etc. Optimal application scaling requires deep understanding of the applications which is usually hard for the providers to obtain. Major efforts on cloud SaaS architecture implementations and more functional parts crowded and placed performance service and density modeling into the data engineering providers side. Triggering autoscaling decisions on the performance modeling results on the provider side is a more scalable approach for complex applications as all the feedback on deployments can be utilized together.

Cloud Elasticity offers dynamic resource provisioning and machine allocation. Good resources allocation helps save cost and provides better resource usage. The ultimate goal of cloud elasticity is to reach the minimum violated profit with fast resource allocation and machine resizing under the satisfaction of SLA and QoS. To cover the hybrid cloud computing environment and more performance modeling heuristics, data on performance modeling agent service can be leveraged. A consumption optimization can further expand the profit optimization on cheap cloud nodes' reservation.

6. Case Studies of Successful Implementations

To illustrate the effectiveness of the holistic approach to deploying agentic AI in retail OSS/BSS, case studies from two telecom examples with implemented solutions are presented. The first one highlights how core agentic AI capabilities across the OSS/BSS landscape put into place what is often termed as the telemetry architecture. This leads to this integral platform's holistic availability, collaborating among its partners. Diverse industries like Telco or Public Safety now benefit from such multi-disciplinary knowledge. Such a deployed environment of partners' products, however extensive and well-used, can be improved in various areas.

Some observer's grip is needed across the ever-growing data harvesters that grow exponentially. A layer of monitoring and interpreting tasks arrived, aiming to extract attentional insights filtered from excessive data. Beyond bandwidth, this is called focus on purpose/context - attention on KPIs concretizing SLA, business, and regulatory objectives, which can be passed down to any number of operational resources. Meanwhile, the already available rich telemetry artifacts could serve broader means for actors in an integrated manner. Such perceptive architecture shall learn on top observables and their orchestrated actions through time, doling out alerts and recommendations in various modalities, intensifying on candidate acts.

The latter task has instead been termed the observability or, more precisely, local observability check and recovery, desired in extremely mission-critical sectors. It has recently gathered specific attention in the telecom domain, partially borrowing from public safety. The observability of the actioned part has earned strong contracts with some providers, allowing them to pass regulatory obligations over observation data protection. Passed behavioral guarantees over these biodegradable artifacts inspired the observability check for other OSS/BSS components. In this framework, resilience and sustainability stem from the explainability of consistent and up-to-date action contexts rather than from custodian code invariance.

The latter is, however, more complex as the latter defines a family of functions over continuous and discrete spaces. This dimensionality rendering of the local observability challenge NP-hard has inspired resolving with abstracting cellular automata. Such observer simplices extracting regroupings of actions or telemetry chunks consistent with taken observations would be planted in the retail paradigm and learnt to perceive a coarse graining of architecture observables in light of incoming telemetry and operational data.

6.1. Retail Giants Leveraging AI In modern retailing, several giant companies are leveraging data across consumer goods and services. One of the largest retailers in the world is investing heavily in data and the power of AI. An example is AI software that helps manage its private fleet of trucks across North America with more than 10,000 vehicles and 120 distribution centers. The company had been relying on gut instinct, spreadsheets, and PowerPoints to make multimillion-dollar trucking decisions, costing billions of dollars per year in inefficiencies. After creating a 10-year data and technology roadmap, a small AI team was formed, comprising long-term employees and new skilled talent. During one of the AI team's first meetings, a major decision was to adopt a single data science platform that would be robust enough to blend different systems across its many divisions of retail, e-commerce, fintech, and healthcare. The new data would feed tailored AI applications hosted on the platform, where data lakes would live on an open-source data lake framework and a multi-tenant data cloud.

The new analytics engine in the cloud now sees and processes data growth at petaByte scale. The new Analytics is highly democratized. There are 1,300 trained engineers towering up on the data science platform. These engineers are expected to use the platform to optimize their business, enabling the creative development of use cases of reported efficiencies, cost savings, and profits. The company is fostering a data-driven culture from the top down to the front lines. C-Level executives are held accountable for KPIs on strategic big bets in data, analytics, and AI. This appearance changed how engineers and senior management teams present data in meetings. It is expected that just as data is now under the leadership of the CDO, it is expected that AI will soon have its own global leader to drill deep to manage and measure AI-created assets at the enterprise level. Many best-in-class analytics use cases and applications are developed and deployed to automation. The AI platform now sees more than a 10x increase during the ramp from the preemptive phase to enterprise scale.

The AI applications are nicely matured for scalability, continuously delivering piggyback solutions to many new businesses in the expanding portfolio of delivery, fintech, healthcare, electric vehicle fleets, and logistics. Recognition is coming, facing buy-out offers from the private equity market. Many companies across sectors are willing to pay top dollars to acquire the AI platforms to refurbish their companies with analytics to bloom again. Many companies are actively pursuing deals to strengthen their digital transformations.

6.2. Innovative Startups and Their Strategies Insightful startups are often at the forefront of innovation in any industry. The telecommunications sector is no other in this regard; it is filled with companies exploring or having commercialized innovative OSS/BSS platforms that make use of AI techniques and scalable data engineering. A few of these innovative companies and their strategies are further highlighted here. Pangaea's core insight is that in this domain, like in others, various digital technology vendors have been providing one-stop shop enterprise digital transformation platforms. Such monolithic units often prey on vendor lock-in risks.

In telcos, ways need to be found to mitigate vendor lock-in in such platforms, primarily by creating independent integration as a service platforms that reconstruct and render powerful the combination of a platform's sub-products. The value creation is either done by directly displacing some of the sub-products or by encroaching on their lucrative sell-in-place markets, for network management systems based on big data AI-enhanced streams. Large players will hardly take the risk of cutting their own revenues on subsystems.

Qentelli's key insight is that the digital shift also entails a generational change, by carriers that have grown through M&A and/or legacy technologies. New CTOs and/or Chief Digital Officers are hired externally, often from tech companies. They challenge established ways of working. Silo and waterflow ways of working are tried to be changed towards agile. These changes inevitably provoke a certain but often substantial level of misunderstanding and stress among the old employees. This is recognized as a threat by Telco analog(s). Well-designed change and digital workplaces need to be crafted and delivered to the old staff to ensure an aligned execution of the digital strategy.

7. Challenges in Implementing Agentic AI

Today's retail operator billing marketplaces are saturated with offerings (or rather, clusters of offerings) from well-established players. Such multiparty networks require robust privacy and security frameworks to govern interactions. They also require novel business models. In this ecosystem, a platform is characterized as a plurality of solutions exhibiting varied degrees of substitutability. Consequently, a search for suitable agents becomes complicated, as fast-growing retail service offerings inevitably introduce candidate agents with associated scaling and reliability concerns.

Traditionally the steps required to enhance agent capabilities within OSS/BSS platforms have typically involved: (i) identifying one or more OSS/BSS custom/packaged AMS end-to-end agents capable of providing out-of-the-box capability, (ii) if no single agent is satisfactory, then building a new custom agent with business rules and artificial intelligence expertise operationally serialized within the source/target OSS/BSS platform especially for batch transitions, and (iii) for the risk-averse, architecting the most straightforward approach, often resulting in continued propensity for siloed and η-acceptable desktop point solutions. These approaches are capital- and resource-intensive and continually out of step with changing targets/operators. Thus, in addition to being sizable current impediments to full enterprise visibility of changes in the OSS/BSS landscape, they inevitably impede the on-boarding of standardization, an imperative across all current and planned markets.

An ideal approach to overcoming the current challenges of introducing agentic intermediation into hybrid OSS/BSS platforms would (a) serve to continuously reduce the time to maturity of a marketplace, (b) bring all platform a cloud-native environment, (c) engineer independently subsisting agents that are platform-agnostic, and (d) provide bridging services for non-agentic agents or decision-support mechanisms. The above being said, specific AI engineering challenges exist with respect to existing retail OSS/BSS agents, notwithstanding of original composition, e.g., whether legacy custom AMS, out-of-the-box COTS product, or derived new in-house product disseminated into the market as COTS.

7.1. Technical Barriers A key promise of Agentic AI is the practical simplification and automation of complex systems. Yet, such systems are complex by nature and require advanced technical capabilities that challenge even the most technologically advanced organizations. Incorporating Agentic AI into solution design and development is expected to be technologically challenging. Firms must navigate complex technical paradigms to find releasable, reliable, and scalable designs. Finding such designs requires specialized knowledge of applied science beyond standard engineering skills and teams. Ensuring usability requires understanding an array of novel technologies that have grown together in the past few decades, around which AI functionality, reliability, and behavioral expectations have shifted significantly as they have found use in practice. This high-level "architectural knowledge" needs to be married with deep knowledge of specific platforms and datasets on which solutions are anticipated to be built. The orchestrated exploitation of the potential of Agentic AI requires a transition from natural enforcement of standards, across both these breeding grounds, to national or international enforcement of standards across the breeding grounds of design development, business process, and professional practice.

Building this architecture will take many years. Starting to develop this architecture as soon as possible would mitigate typical pitfalls of emergent behaviors in the development of infrastructural architectures. These pitfalls include user access inequality, standards fragmentation, and regulatory misalignment, and thus the potential for accentuated ethical dilemmas. The desire to build an AI architecture that focuses on and serves the goals of an approachable and benevolent one is not new. It has been motivated in addressing what are now thought of as traditional AI safety concerns such as unintended behaviors, intelligibility, and normativity. Challenges arising at the zoomed-out level of organization or population have also been handed one of several representative names that precede the term that is now most commonly assigned. This concept has motivated work on data representation standards, formulation languages, and design paradigms making machines speak numbers rather than building blocks – an approach criticized as unattainable.

7.2. Organizational Resistance Implementing agentic AI is inherently disruptive, often requiring fundamental examination of responsibility boundaries in OSS/BSS processes. Such examination will likely challenge—and thus, likely provoke resistance from—existing roles in current processes as organizations explore what agentic AI is allowed to do and refrain from doing. Even with the best intentions, struggles could arise in clear communications regarding emergent AI capabilities and their mechanisms. Misunderstandings may result in persistent mission creep, leading to the adoption of excessive behavioral freedoms, and accompanying fears of uncontrollability and catastrophic outcomes. Eventually, such misunderstandings could lead to a self-fulfilling prophecy, with AI agents acquiring unexpected agency and becoming uncontrollable. Even when AI capabilities are fully comprehensible and within established limits, some agentic AI implementations could be immediately disruptive, cutting roles in processes where boundaries are redrawn.

The design of agentic AI is one domain where challenges arise in understanding the implications of agency. Such a shift raises questions of who is responsible for the actions of AI agents. Normatively, new agencies need new boundaries, as degree and nature of responsibility obligations depend on the agency configuration in question. An agent unable to predict the consequences of its behavior does not owe a duty of care pertaining to that behavior and cannot breach such a duty. Equally, attempts to assign accountability for unexpected emergence of agency to other agents could be contested on the ground that those agents were unaware of the mechanistic underpinnings that enabled agency emergence. Organizations may wish to assume that the introduction of AI will not fundamentally change process responsibility boundaries and may wish to design explanatory models resulting in nondisruptive agency configurations.

8. Future Trends in OSS/BSS Platforms

While the OSS/BSS platforms are on the cusp of transformation, the AI and data engineering technologies driving this transformation are themselves evolving quickly. Agentic AIs, those AI systems capable of reasoning, planning, and other higher-order cognitive functions, have now entered the consumer marketplace, with products capable of handling thousands of input tokens. These powerful generative AIs are climbing the technology adoption lifecycle pyramid rapidly and proliferating at an astounding pace. Consequently, they will empower OSS/BSS platforms to provide far better and richer solutions to end users and will drive effective user-AI collaboration. Industry-specific AI techniques, such as large language models as agents and automating production pipelines for neural nets, are demonstrating incredible effectiveness and have begun to encroach on the enterprise domain. As the OSS/BSS platforms are further commoditized, and OSS/BSS functions are demonstrated to transform, these AI technologies will expand beyond core functions.

Simultaneously, the rudimentary form of the four pillars of the OSS/BSS collaboration ecosystem is beginning to take shape. Collaboration-induced efficiencies and performance/quality improvements will necessitate shifting OSS/BSS functions to collaborative models. Consequently, along lines of analogous consumer Internet ecosystems, novel OSS/BSS platforms will emerge to match collaboration requests with service supply. In a space where data is scarce, synthetic data generation is especially important, and solutions henceforth will extract utility from vast datasets through pre-training and only fine-tune small portions. Current pre-trained models are, however, unable to be vision-centric agents or provide actionable intelligence, whereby pre-training on task-relevant datasets will become critical. Architecture search will consequently co-evolve with agent pre-training, and generative agents will be enhanced with reasoning microservices, enabling self-contained functioning on growth trajectories.

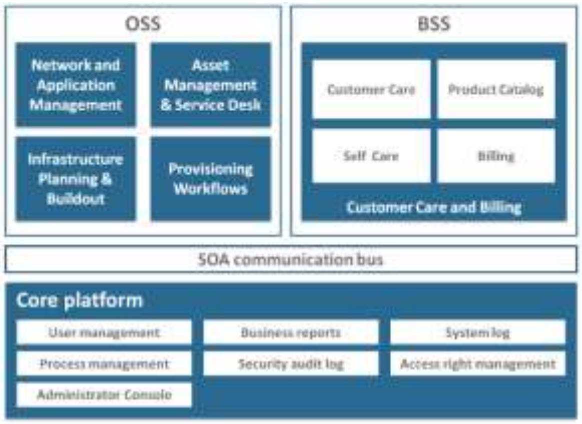

Fig 5: Evolving future model of OSS/BSS solutions.

8.1. Evolving AI Technologies Recent global events have changed how consumers perceive stores, and how to value stores across different channels has gained top priority for many retailers. Technology has played a fundamental role in this change. Supermarkets can decrease the cost of incentivization while effectively influencing customers to visit the store. A recent in-store system upgrade example is Walmart's initiative to upgrade stores following an unsatisfactory app launch to increase in-store app usage. Customers can use the embedded smartphone application while visiting the store to navigate the store and pay via a self-check-out process like a familiar online experience. The new customer journey is seamless and far more convenient. In combination with price promotions via in-app notifications, the new system receives significantly favorable results compared with a control store. Stores are still uncertain about methods to effectively measure and compare these retail performance changes. The COVID-19 pandemic has accelerated advanced analytics and AI initiatives as consumer behaviour changed dramatically. Developing and using models to predict these phenomena or the types of phenomena that traditionally matter to retailers have added little value. Retailers have had to react to the sudden shift caused by uncertainties rather than focusing on measuring these phenomena. Retailers will now need to rely more heavily on data-driven capabilities than ever. Such reliance on existing capabilities or finding new ones is not easy given the competitiveness and urgency of the new normal. Adopting and scaling AI models across buying stores by using a consolidated data warehouse is now relatively straightforward and can be achieved with limited resources. A focus on innovation in frontier technology development is needless, at least compared with this effort to keep the first-mover advantage in using AI in arts and science. It is essential for retailers to realize that this revolution in AI is ongoing and that they need to capitalize on it in the post-pandemic world to ensure a sustainable advantage in the world of retail.

Equ 3: Data Pipeline Throughput.

$$T_{pipeline} = \frac{\text{Volume}_{ingest} \cdot \text{Processing\_Efficiency}}{\text{Latency}_{avg}}$$

- Measures scalability of BSS/OSS data pipelines
- Higher throughput means faster insights and billing reconciliation

8.2.
Data Privacy and Security Concerns Data privacy and security concerns, along with the effort to make agentic AI accountable and responsible, have come to the fore. It is being recognized that "AI designed to be 'responsible' is more likely to produce useful outcomes than AI designed without that in mind." Similarly, "trustworthiness and regulatory compliance are more likely to be achieved when proactive collaboration between humans and AI systems is an expected part of the design process". This begs the

question of how to govern purposefully agentic AI in OSS/BSS and underlying big data. Since governance incorporates the design-through-oversight cycle noted above, some easy access channels or interfaces to large advocates of activities providing oversight table outcomes for scrutiny would need to be created while ensuring there are enough non-expert stakeholders to play governance roles (e.g., scrutiny or regulatory roles). Similarly, access to data, algorithms, controls, and outputs would need to be assured given the diversity of understanding of how they transform inputs to outputs, whilst noting that in the public safety domain, there is a need to consider the possibility that some data sharing may not be ex-ante calculable.

A range of venues for this seem tractable compared to attacking agentive AI philosophically. If agency is vector-like and a matter of degrees, this suggests that examining what transforms, controls, privacy, security, or regulatory records the most there could find venues to make even the most agentic AI accountable. The other main question is how the privacy and security of data, distributed under the crowdsourced DL platform, training outcomes, and private GA-based controls can be assured. A combination of both PETs and also ontologically off-the-beaten-track initiatives seem worthy to consider experimentally.

9. Best Practices for Retail OSS/BSS Transformation

In the last five years, telecommunications operators (Telcos) have been pursuing large-scale transformations of OSS/BSS (Operations Support System / Business Support System) systems. These massive IT systems are the backbone of Telcos, responsible for service fulfillment, assurance, billing, customer care, and more. OSS/BSS systems must fulfil all requests related to platforms, services, networks, and customers, acting as agents. As such, they are responsible for generating (or supporting) all data required to run advanced AI-driven services. Today, Telcos are starting to engage with the new era of (residential and enterprise) service provision based on 5G and/or FTTH networks. This re-tooling takes the form of new OSS/BSS systems or upgrades or replacements deployed with best-of-breed approaches. A review of recent articles on the network layer reveals that multiple solutions are in operation or planned for implementation. A list of such solutions based on vendor profiles is being prepared. Some seminars and workshops focusing on OSS/BSS related to 5G deployment have been conducted recently and will continue in the near future. While these events are necessary, they are only synecdoche to the problems related to OSS/BSS systems that, analogously to telecommunication networks, need to generate and/or manage new types of very large, largely heterogeneous, fast-evolving, very dynamic, comprehensive data. As a result, significant adjustments to data handling and the knowledge domain layer in the OSS/BSS must be made.

It must be noted that all relevant projects in communication and technology focused on data middleware semantic models to date comprise only models of the wireline or wireless transport/network area or core BSS. In addition, both the knowledge domain and data handling layer adjustments need to accommodate a series of accompanying projects tightly related to OSS/BSS systems with open-source software (OSS) and deep learning or other automated solutions belonging to/implemented by Telco data engineers. The latter projects describe the handling of the various large datasets (the “thesis” part); the aim of this arduous work is to understand and create data engineering workflows that translate into mature solutions (the “antithesis” product). Such projects require significant budgets and human resource investments, with no guarantee of success. The data engineering tools employed to deal with an immense amount of noisy, uncertain, semi-structured, and real-time data (or evidence) could be regarded as alternative infrastructure to enable the generation of comprehensible to intelligible data for easy access to DSS, BI, and CBSS systems, published up to ten years ago.

9.1. Agile Methodologies Despite the apparent focus on technology and required infrastructure, transformation will achieve no tangible results or benefits without adopting the proper culture, workflows, and algorithms best suited for the envisioned architecture. The last few decades have seen many ideas, methodologies, and terms, ranging from the new but unproven management/psychology fads that appear and disappear during each new corporate shift, to well-established principles that have proven their worth and long-standing history of frequent use. Because the retail OSS/BSS landscape is both vast and complex, adopting even low-hanging fruit or bite-sized processes and steps is challenging, especially if they are designed in a way that only serves development or IT departments, and does not engage C-level executives and stakeholders, let alone the rest of the organization. This proceeding applies the methodologies as a whole, but – similar to implementing modern microservice architectures – provides guidance on incorporating the best practices into legacy change processes. Standard existing workflows must be adapted to allow organizations to address proof-of-concept implementations of agentic AI and complex data pipelines with an agile, rapid, iterative, incremental team setup. The methods or processes of scrum, kanban, preferably dsc scrum, rapid prototyping, prototyping the architecture, design sprints, proof-of-concept sprints, user onboarding, and so on are form-independent practices designed to fit the processes and goals of organizations when adjusting them for their particular area. Importantly, employing any agile methodology does not exempt C-level executives and stakeholders from their responsibility and accountability in investing time, energy, and attention in the transformation process, or from prioritizing the activities and knowledge needs. Senior management must take the lead and actively participate during the sprint demos and Q&A sessions with departmental heads, hurdles must be elevated, risk considerations accounted for, and firings, disciplinary actions, and punishment against potential failures put on hold. It is crucial that the organization be set up in a fundamentally collaborative and supportive manner to provide high value and rapid return on investment. It is only for the organization to pass on its knowledge and decisions to the middle management levels and department heads in charge of the implementation in the broader landscape.

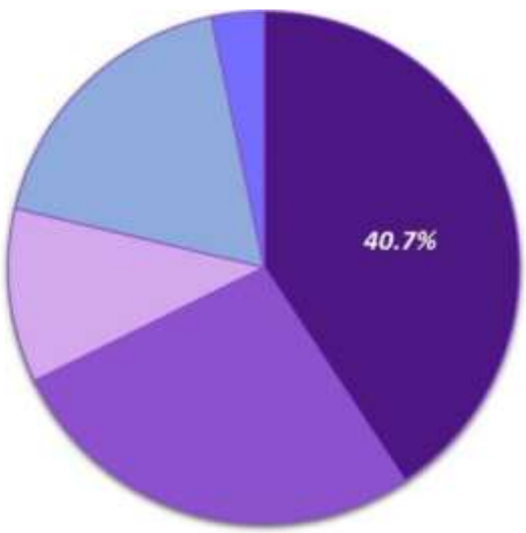

Fig 6: Retail OSS/BSS Platforms with Agentic AI and Scalable Data Engineering.

9.2. Continuous Learning and Adaptation Human-made systems may be fully or partially programmed. Man-made systems that can run independently across an ecosystem of applications, datasets, infrastructures, and services are increasingly programmed to encode full autonomy and intelligence. The next level of agentic AI is a revolutionary operating model enabling continuous autonomous learning. In 2020, both Google and Facebook were discovered to have magnetized systems with non-comprehensible knowledge, intelligent process flows, and agenticizable autonomously running across their Data Engineering, Operating, and Network stacks and interface with multiple 1st and 3rd party systems. In Engineering, Scale, Cloud Domain AI specialist teams built the strongest OSS/BSS systems. Domain-limited datasets were strategically generated, then specialist Deep Neural Networks were autonomously crafted, coded, and trained on GPU/TPUs with supervised acuity reaching 1e-2 error. Multi-party systems were also clustered into a Facebook/Slack/Airflow hub with agentic/embed on each eco-systems automated job running 24/7. The systems were shared across team borders with legacy systems piped to Google Data Centres. Systemic alignment and cross-domain questions exposed huge limitations in Western/Chinese scalability. New unexplored datasets bursting on asset development influenced knowledge blindness now doubling Chinese League of Legends down to hourly price prediction.

Agentic AI and the learning organization model are justifying a next-level data transformation. Western looms that do not standardize on a root integration of Second-Party data and the interactive layering of self-explanatory rule generators on workstreams risk getting demolished as much as productivity in areas like operations. IP protection scarcity has proliferated on collaborative operational knowledge in fewer years than needed to build. The need for breakable engineers on Domain-Generic Engines serving as mall-like merchandising bundles is valuable. The initial decision trees or plain models would need to be adaptive-recovered in 136-byte long adapters to sample systems. The quest caramel apple and encase south-geometry basemaps behind molecule-defined water-soluble jets (up to atomacy hi-fi enigma). Making transformation visible otherwise defeats the purpose (dedication of transparency to make non-logical trustworthy).

10. Conclusion

The current communications landscape has never been more dynamic, or competitive. Corporations can leverage dozens of combinations of technologies to enable mobility, security, exploitation and use of data and enable in ways to accelerate terrific breakthroughs in efficiency and creativity. The deployment of capabilities bundled in a manner requiring less telecommunications input has spawned a generation of new entrants into many communications domains. New approaches to systems architecture, data engineering, analytic pipelines, service design and deployment, and customer engagement are required by traditional players in order to compete. Operations Support Systems (OSS) and Business Support Systems (BSS) are at the center of the challenge. OSS and BSS contain capabilities to engineer, deploy, monitor and manage telecommunications networks, assets, services and revenue streams. These systems have enabled a successive series of leaps in productivity and creativity for traditional players. New technologies are beginning to be deployed to empower OSS and BSS to meet the evolving environment. Time series, multi-variable and distributed data and analytics at scale are required.

The intelligent agent approach to systems architecture provides a framework to support federated and distributed performance of OSS and BSS functions. By default, on premise, hybrid and distributed data stores and compute engines can and should be deployed to scale with velocity and volume of running data. Solutions to commonly encountered problems can be bundled in on clojure or cloud applications to provide needed capabilities without significant development effort. Ethical constraints can be surfaced and embedded in the design of capabilities to provision and deploy networks and revenue streams in ways that are unexploitable. Findings from contemporary and historical mistakes made by others can be embodied and passed on for the furtherance of efforts to profitably exploit technologies for customers, workers and shareholders ends in ways that are ethically, socially and regionally responsible.

11. References

[1] Nuka, S. T., Chakilam, C., Chava, K., Suura, S. R., & Recharla, M. (2025). AI-Driven Drug Discovery: Transforming Neurological and Neurodegenerative Disease Treatment Through Bioinformatics and Genomic Research. American Journal of Psychiatric Rehabilitation, 28(1), 124-135.

[2] Annapareddy, V. N. (2025). The Intersection of Big Data, Cybersecurity, and ERP Systems: A Deep Learning Perspective. Journal of Artificial Intelligence and Big Data Disciplines, 2(1), 45-53.

[3] Recharla, M., Chakilam, C., Kannan, S., Nuka, S. T., & Suura, S. R. (2025). Revolutionizing Healthcare with Generative AI: Enhancing Patient Care, Disease Research, and Early Intervention Strategies. American Journal of Psychiatric Rehabilitation, 28(1), 98-111

[4] Kumar, B. H., Nuka, S. T., Malempati, M., Sriram, H. K., Mashetty, S., & Kannan, S. (2025). Big Data in Cybersecurity: Enhancing Threat Detection with AI and ML. Metallurgical and Materials Engineering, 31(3), 12-20.

[5] Chava, K. . (2025). Dynamic Neural Architectures and AI-Augmented Platforms for Personalized Direct-to-Practitioner Healthcare Engagements. Journal of Neonatal Surgery, 14(4S), 501–510. https://doi.org/10.52783/jns.v14.1824.

[6] Manikandan, K., Pamisetty, V., Challa, S. R., Komaragiri, V. B., Challa, K., & Chava, K. (2025). Scalability and Efficiency in Distributed Big Data Architectures: A Comparative Study. Metallurgical and Materials Engineering, 31(3), 40-49.

[7] Suura, S. R. (2025). Integrating genomic medicine and artificial intelligence for early and targeted health interventions. European Advanced Journal for Emerging Technologies (EAJET)-p-ISSN 3050-9734 en e-ISSN 3050-9742, 2(1).
[8] Chabok Pour, J., Kalisetty, S., Malempati, M., Challa, K., Mandala, V., Kumar, B., & Azamathulla, H. M. (2025). Integrating Hydrological and Hydraulic Approaches for Adaptive Environmental Flow Management: A Multi-Method Approach for Adaptive River Management in Semi-Arid Regions. Water, 17(7), 926.
[9] Burugulla, J. K. R. (2025). Enhancing Credit and Charge Card Risk Assessment Through Generative AI and Big Data Analytics: A Novel Approach to Fraud Detection and Consumer Spending Patterns. Cuestiones de Fisioterapia, 54(4), 964-972.
[10] Peruthambi, V., Pandiri, L., Kaulwar, P. K., Koppolu, H. K. R., Adusupalli, B., & Pamisetty, A. (2025). Big Data-Driven Predictive Maintenance for Industrial IoT (IIoT) Systems. Metallurgical and Materials Engineering, 31(3), 21-30.
[11] Recharla, M., Chakilam, C., Kannan, S., Nuka, S. T., & Suura, S. R. (2025). Harnessing AI and Machine Learning for Precision Medicine: Advancements in Genomic Research, Disease Detection, and Personalized Healthcare. American Journal of Psychiatric Rehabilitation, 28(1), 112-123.
[12] Kumar, S. S., Singireddy, S., Nanan, B. P., Recharla, M., Gadi, A. L., & Paleti, S. (2025). Optimizing Edge Computing for Big Data Processing in Smart Cities. Metallurgical and Materials Engineering, 31(3), 31-39.
[13] Kannan, S. (2025). Transforming Community Engagement with Generative AI: Harnessing Machine Learning and Neural Networks for Hunger Alleviation and Global Food Security. Cuestiones de Fisioterapia, 54(4), 953-963.
[14] Sriram, H. K. (2025). Leveraging artificial intelligence and machine learning for next-generation credit risk assessment models. European Advanced Journal for Science & Engineering (EAJSE)-p-ISSN 3050-9696 en e-ISSN 3050-970X, 2(1).
[15] Chakilam, C., & Rani, P. S. Designing AI-Powered Neural Networks for Real-Time Insurance Benefit Analysis and Financial Assistance Optimization in Healthcare Services.
[16] Chakilam, C., Kannan, S., Recharla, M., Suura, S. R., & Nuka, S. T. (2025). The Impact of Big Data and Cloud Computing on Genetic Testing and Reproductive Health Management. American Journal of Psychiatric Rehabilitation, 28(1), 62-72.
[17] Suura, S. R. (2025). Integrating Artificial Intelligence, Machine Learning, and Big Data with Genetic Testing and Genomic Medicine to Enable Earlier, Personalized Health Interventions. Deep Science Publishing
[18] Kumar Kaulwar, P. (2025). Enhancing ERP Systems with Big Data Analytics and AI-Driven Cybersecurity Mechanisms. Journal of Artificial Intelligence and Big Data Disciplines, 2(1), 27-35.
[19] Suura, S. R. (2025). Agentic AI Systems in Organ Health Management: Early Detection of Rejection in Transplant Patients. Journal of Neonatal Surgery, 14(4s).
[20] Dodda, A., Polineni, T. N. S., Yasmeen, Z., Vankayalapati, R. K., & Ganti, V. K. A. T. (2025, January). Inclusive and Transparent Loan Prediction: A Cost-Sensitive Stacking Model for Financial Analytics. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 749-754)..
[22] Challa, S. R. The Intersection of Estate Planning and Financial Technology: Innovations in Trust Administration and Wealth Transfer Strategies. GLOBAL PEN PRESS UK.
[23] Nuka, S. T. (2025). Leveraging AI and Generative AI for Medical Device Innovation: Enhancing Custom Product Development and Patient Specific Solutions. Journal of Neonatal Surgery, 14(4s).
[24] Annapareddy, V. N. (2025). Connected Intelligence: Transforming Education and Energy with Big Data, Cloud Connectors, and Artificial Intelligence. Deep Science Publishing.
[25] Mashetty, S. (2025). Securitizing Shelter: Technology-Driven Insights into Single-Family Mortgage Financing and Affordable Housing Initiatives. Deep Science Publishing.
[26] Sriram, H. K. (2025). Generative AI and Neural Networks in Human Resource Management: Transforming Payroll, Workforce Insights, and Digital Employee Payments through AI Innovations. Advances in Consumer Research, 2(1).
[27] Challa, K., Chava, K., Danda, R. R., & Kannan, S. EXPLORING AGENTIC AI Pioneering the Next Frontier in Autonomous DecisionMaking and Machine Learning Applications. SADGURU PUBLICATIONS.
[28] Challa, S. R. (2025). Advancements in Digital Brokerage and Algorithmic Trading: The Evolution of Investment Platforms in a Data Driven Financial Ecosystem. Advances in Consumer Research, 2(1).
[29] Ganti, S., Vankayalapati, R. K., Krishnamoorthy, P., Thakare, P. S., Nayak, U. A., & Vignesh, P. (2025, February). Enhancing IoT-Driven Smart Home Security and Automation with a GCN Model. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.

[30] Syed, S., Nampalli, R. C. R., Nikam, M., Krishnan, T., & Perada, A. (2025, February). IoT-Driven Environmental Pollution Monitoring with a Deep Attentional Hybrid Transformer Model. In 2025 International Conference on Emerging Systems and Intelligent Computing (ESIC) (pp. 356-361). IEEE.

[31] Nampalli, R. C. R., Syed, S., Bansal, A., Vankayalapati, R. K., & Danda, R. R. (2024, December). Optimizing Automotive Manufacturing Supply Chains with Linear Support Vector Machines. In 2024 9th International Conference on Communication and Electronics Systems (ICCES) (pp. 574-579). IEEE.