Journal for Re Attach Therapy and Developmental Diversities
eISSN: 2589-7799
2023 December; 6 (10s(2)): 2351-2366

From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust

Kushvanth Chowdary Nagabhyru[1*]

[1*]*Data Engineer* ORCID ID : 0009-0004-7175-7024

Abstract—

Data silos within and across organisations frequently hinder the implementation of scalable and trusted artificial intelligence (AI) solutions. The critical capabilities needed to overcome data last mile problems can be broadly summarised as the provision of a solid AI architecture coupled with scalable knowledge graphs. When these capabilities are realised, data silos become rest-silos—repositories that can be safely and trustworthily queried without any potential for their alteration or modification. These capabilities and their architectural imple- mentation are examined in detail. The synergistic integration of solid AI architectural principles and scalable knowledge graphs effectively eliminates the adverse effects associated with data silos. Leveraging the mediation ability of knowledge graphs, previously inaccessible data can be integrated from multiple departments, corporate entities and across business ecosystems. New product and service opportunities, together with the associated corporate business and operating models, can then be established.

Index Terms—Data Silos, Artificial Intelligence, Scalable AI Solutions, Data Last Mile, AI Architecture, Knowledge Graphs, Scalable Knowledge Graphs, Trusted Data Access, Data Inte- gration, Cross-Organisational Data, Business Ecosystems, Data Mediation, Repository Management, Corporate Entities, Infor- mation Accessibility, New Product Opportunities, Service In- novation, Operating Models, Data Trustworthiness, Synergistic Integration.

I. INTRODUCTION

This document guides organizations transitioning from data silos toward knowledge graphs to build scalable, cross-enterprise artificial intelligence (AI). It presents data silos, outlines how knowledge graphs work and the benefits they offer, then describes how combining AI and knowledge graphs enables scalable, cross-enterprise AI. It concludes by defining trust, explaining why trust is crucial to AI, and describing how trust can be implemented in an organizations' AI offerings. The first section defines data silos and explores challenges they impose on enterprise data integration. It is followed by an examination of knowledge graphs with an emphasis on their ability to connect different enterprise data silos. The next section focuses on AI architecture principles and considers how knowledge graphs contribute to scalable, cross-enterprise AI. The final section defines the role of trust in AI and explains how organizations can erect mechanisms to engender user confidence in their AI offerings.

A. Overview of the Document's Objectives and Structure

This document articulates key principles for a confluence of enterprise artificial intelligence (AI) architectures, knowledge graph technology, and the mitigation of data silo challenges. In advocating knowledge graphs as pivotal enablers of scalable cross-enterprise AI and business intelligence, the aim is to highlight a clearer route toward such deployment. Interweav- ing the AI architecture discussion with data silo analysis provides a lucid rationale for rejecting the silo paradigm and, through knowledge graphs, constructing scalable, resilient, and trustworthy cross-enterprise AI infrastructures. The es- sentiality of achieving scalable, trusted AI within and across enterprise boundaries cannot be overstated. Cross-enterprise AI demands collaboration among independent entities that typically establish such endeavors only when the gains surpass inherent risks and costs. The discussion commences in section 2 with detailed exploration of the vulnerabilities inherent in enterprise data silos. After illustrating the fundamental value proposition of knowledge graphs, section 3 presents a set of architectural objectives for enterprise AI solutions in 2023. The subsequent analysis of scalable and trusted AI within both enterprise silos and across enterprises flows naturally from these objectives. The assertion is that contemporary enabling technologies empower scalable AI architectures capable of surmounting silo-related issues.

II. UNDERSTANDING DATA SILOS

Data silos are autonomous repositories, controlled by one department or unit, and are generally isolated from the rest of the enterprise. Data in proprietary formats make it difficult to share with other parts of the enterprise for other business units, and even more difficult to share with suppliers or customers. In addition, they are frequently established through manual effort and are generally difficult to maintain over time. They also tend to degrade in overall quality because of propagation delays from source systems and because there is no feedback loop from the application that relies on the information. An understanding of data silos is crucial when attempting to hollow out these silos and share enterprise

Journal for Re Attach Therapy and Developmental Diversities
eISSN: 2589-7799
2023 December; 6 (10s(2)): 2351-2366

data in order to build an enterprise AI solution that goes beyond a departmental or business unit focus. Some thought should also be given to how the ecosystem partners consume and integrate the data in their business processes or in support of their own AI model

Fig. 1. Knowledge Graphs: Breaking Down Data Silos

initiatives. Finally, an understanding of data silos should be part of the evaluation of when it is appropriate to leverage a knowledge graph to build cross-enterprise AI solutions. An example of the challenges faced when attempting to unify data from silos was collected during a discussion with a retailer that wished to build an AI pricing model. The requirements for the pricing forecast model dictated the use of competitor pricing information. However, the competitor pricing information was sourced from marketing, ERP, and finance, combined by category management, and then distributed to other units. Each group had their own SLA's for supplying the information. This severely impacted the accuracy of the pricing forecasts because a model is as good as the quality of data within it.

A. Maintaining the Integrity of the Specifications

A data silo occurs when a dataset does not integrate well with other enterprise data. In such cases, the data is poorly linked with the datasets maintained by other business func- tions. When data is operationalized in a manner that limits its access and use, a silo is likely to form. The consequences of data silos include the creation of segregated data repos- itories for operational and analytical use, limited access to

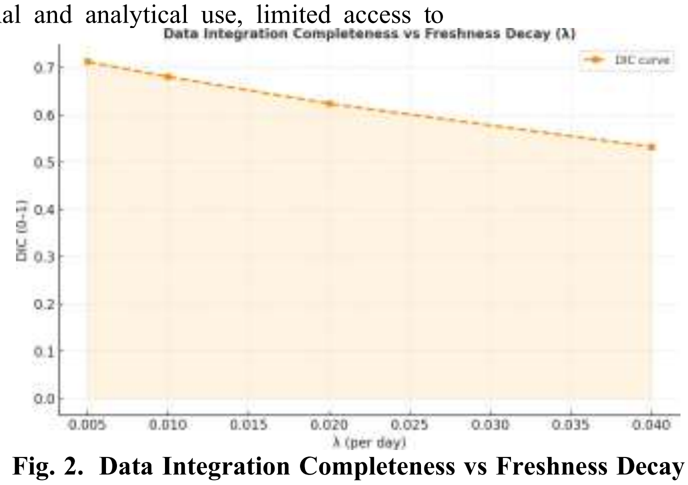

Fig. 2. Data Integration Completeness vs Freshness Decay

B. Challenges Posed by Data Silos

Analyzing the impact of data silos on enterprises reveals a persuasive rationale for their future disruption and replacement by alternative forms of organized data representation. Data silos severely complicate the integration of data from multiple islands across any enterprise, irrespective of whether the information is maintained by different business divisions or held by partners, vendors, suppliers, or customers. During the planning of diverse products or for assessing cross-enterprise supply chain risks, supply chain personnel frequently need to gather, consolidate, and analyze information from multiple sources to obtain a comprehensive snapshot of a product's narrative. Often, this preparation entails crawling through multiple documents and websites. Additionally, the analysis rarely spans the entire supply chain of a finished product because data cannot flow easily across partners.

Equation 1 : Data Integration Completeness (DIC)
Goal in the paper: quantify how "complete" a cross-silo dataset is when building AI.
We model each required attribute a_i, $i = 1..n$ with: wi= business weight (importance), qi ∈ [0,1] = data quality score (accuracy/consistency), ci ∈ [0,1] = coverage (fraction of records populated), ai= age in days, with freshness factor $f_i = e^{-\lambda a i}$ (expo-nential decay; λ sets how fast stale data "counts less").
these resources, a lack of responsiveness to rapidly chang- ing business questions, enterprise data that possesses limited business context, inefficient and opaque data processing, and restricted reuse of data. Accenture's retention and analytics asset underwent such a transformation; initially consisting of siloed, batch-driven applications, it evolved into a streamlined, resilient, and trusted AI platform serving a variety of retention and analytics oriented use cases. Today, the dismantling of data silos is a critical issue motivating the adoption of knowledge graphs. They represent an ideal methodology for business data that is subject to collaboration and cross-enterprise visibility.

$$DIC = \frac{\sum_{i=1}^{n} w_i q_i c_i f_i}{\sum_{i=1}^{n} w_i}, \quad f_i = e^{-\lambda a_i}. \qquad (1)$$

C. Case Studies of Data Silos in Enterprises
When data is generated, analysed and stored in a data silo, it is still useful. Only when this data needs to be digitally communicated and shared with other departments or with other organizations, problems start to arise. For example, the European Union has sponsored several R&D projects that have produced research results, published scientific articles, nur- tured specialised communities, and organized related events.
Some examples include projects on renewable energy targets, projects on climate change mitigation, projects to support the COP 15, COP 21, and COP 26 conferences, among many others. This wealth of information today exists in many fragmented. The fragmentation of the information is the direct effect of being developed by different entities in silos, with their own motivations and intended use. In an enterprise con- text, knowledge assets, classified as non-physical assets, can be categorized in three types: customer capital, human capital, and process capital. Customer capital is a knowledge asset related to external stakeholders such as customers, suppliers, business partners, and the value established by such relation- ships. Human capital refers to employees and their expertise. Finally, process capital is a knowledge asset related to the methods or processes that are performed within an organiza- tion. The presence of separate groups within organizations that are structured to dynamically collect and analyze information, engage and evolve relationships, and manage employees, can result in a significantly fragmented repository of knowledge assets. In the absence of the consolidation or unification of these separate data sources, information redundancy, narrow business intelligence, and reduced customer perception can eventually develop. As a result, all these issues combined are the cause of a company's inability to establish scale relationships with their customers, as the initial intel on the customer is not updated consistently.

III. THE EMERGENCE OF KNOWLEDGE GRAPHS

In statistics and machine learning, a black swan is an observation that lies outside the realm of regular expectations. Black swans can only be modeled with a distribution that allows for large deviations, such as a t-distribution with three or less degrees of freedom. Although black swans have been examined by numerous scholars, few voices distinguished the causes and types of black swans. On the other hand, it can be argued that the existence of black swans in great numbers is incorrect. Thus, the question can be posed: how many times can a black swan occur? Responding to that question involves two aspects: - Are the outliers of a data generation process really extreme values? If so, the large uncertainty of the corresponding estimate can be understood with a heavy-tailed distribution. - Can the artifact that is sometimes called a black swan be numerically estimated? In this case, the skewed and contaminative term can be modeled with a generalized Pareto distribution. A differentiated view asserts that the first aspect of the question deals with systemic uncertainty, while the second aspect pertains to environmental uncertainty. Knowing the type of black swan contributes to the establishment of a suitable model. Although a black swan event lies outside the realm of regular expectations, it is still possible that multiple sets of leaders foresee a black swan event but the rest do not, and/or that some may even create a black swan.

Equation 2 : Knowledge Graph Connectivity (KGC) Goal: capture how well a knowledge graph "breaks silos" and supports cross-enterprise queries.

Journal for Re Attach Therapy and Developmental Diversities
eISSN: 2589-7799
2023 December; 6 (10s(2)): 2351-2366

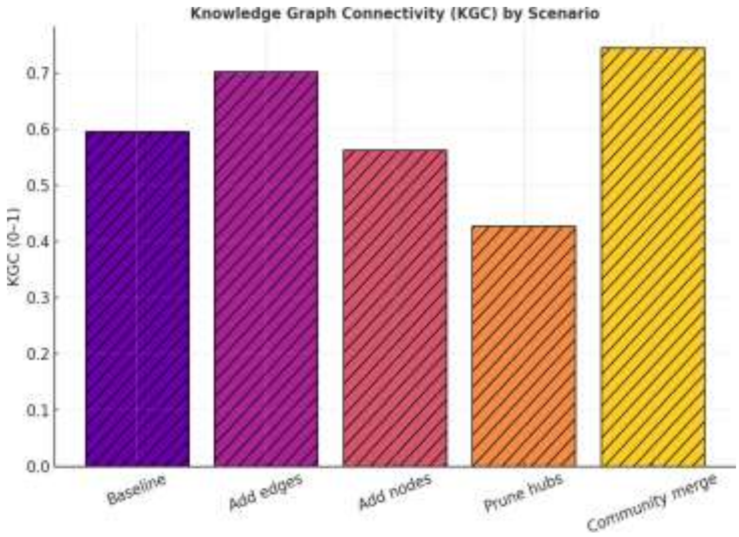

Fig. 3. Knowledge Graph Connectivity (KGC) by Scenario

Combine core graph properties into a single index: $g \in [0, 1]$: fraction of nodes in the giant component, $C \in [0, 1]$: average clustering (local cohesiveness),
L: average shortest-path length (efficiency). Normalize it:
$Lnorm = Lmax - Lmin/L - Lmin, so 1 - Lnorm \in [0, 1]$.
(2)
$KGC = g + C + (1Lnorm)$. (3)

A. Definition and Structure

DData silos are repositories of fixed data sets, accessible to one group, but isolated from or inappropriately connected to other systems. Within an enterprise, data silos can serve as the backbone for the teams that create and curate the data. In contrast to open enterprise data, these data can define closed ecosystems in both a cultural and a technical sense. When data silos or closed ecosystems expand across the boundaries of multiple companies, the connected organizations can become locked into “a business relationship for the exchange of information and services.” They “often work together with a certain level of mutual trust.” Despite the close collabo- ration and mutual trust, these organizations face challenges when combining their data. Real-world examples illustrate the frictions. When data silos dominate the connection between organizations, the cost of scaling AI use cases over the combined data becomes excessively high. Knowledge graphs are a paradigm for representing linked information in a graph. When coupled with a reasoning engine, knowledge reasoning allows the derivation of new knowledge from an existing knowledge graph. The models form a basis for various AI techniques, such as hybrid question answering or autonomous driving. By linking previously disconnected data, the creation of an enterprise knowledge graph breaks open data silos. Many of the technical challenges of data-silos-based AI architectures disappear: knowledge graph construction enables scalable AI use cases over data that originate from different groups within the same company or from multiple organizations. The graph structure facilitates integration beyond mere technical cou- pling. Scalability is also enhanced when multiple knowledge graph use cases reuse the same data: although this requires additional data governance work the reusability creates syn- ergy effects and thus significantly reduces the costs involved in establishing additional use cases.

attribute	weight	quality q	coverage c	age days
A1	3.0	0.95	1.0	5
A2	2.0	0.9	1.0	10
A3	1.5	0.8	0.8	20
A4	2.5	0.85	1.0	2
A5	1.0	0.7	0.6	30

TABLE I DIC ATTRIBUTES

B. Benefits of Knowledge Graphs

The next-generation AI architecture that is key to solving the data-silo challenge uses knowledge graphs, a powerful new data structure in the Digital Age. Knowledge graphs store knowledge in a graph: in nodes, the graph vertices, are the entities of focus, e.g., customers, businesses, and products; and in edges, the graph links, is the relationship structure, e.g., a works at connection between a Person node and an Organization node. Supporting information and context can be added in the form of node types, node attributes, edge types, and edge properties. In this way a knowledge graph, reaches beyond transactional and master data to organize the entire enterprise's knowledge.

C. Comparative Analysis: Knowledge Graphs vs. Traditional Databases

Data silos still have a place in enterprises, but for enterprise-level AI architecture, they have to be broken down or a new architecture paradigm has to be established to connect silos across an enterprise. Knowledge graphs have become a popular enterprise architecture for building cross-organization AI solutions that act as a hub of data, linking different data silos together. Knowledge graphs also enable data interoperability through an integrated unified knowledge representation. They provide an architecture that is scalable, supports demand planning, and helps build trust, transparency, and accountability into enterprise AI systems. Incorporating a knowledge graph into the enterprise AI architecture enhances both the scalability and trust of AI solutions. Knowledge graphs organize information in a way that mirrors the structure of data as it exists in the real world, employing a network-structured database and utilizing semantic labels—such as named entities and categories—to present structured data conveniently. In comparison with conventional databases, their primary advantage lies in the facilities they provide for knowledge discovery and data description. Knowledge graphs facilitate the discovery of new links, predictions, and suitable hypotheses in a more direct and user-driven manner.

IV. ARCHITECTING CROSS-ENTERPRISE AI SOLUTIONS

Understanding Data Silos Data silos are isolated data repos- itories within an organisation. They are disconnected from other business units and departments. Each department or unit has its own data stored by itself or its IT team. The distinct characteristics of data silos include:• A particular business unit and the corresponding data belong to it.• It has its own centralised system implemented, that stores all incoming data daily.• Each request created by the business unit towards the centralised system, contains the business unit code along with the request details.• A proper departmental authorisation mechanism.Map of data stored within an or- ganisation is shown in Figure 14• Different business units exchange data based on business requirements via a request carefully designed and authenticated by a dedicated server.• Each business unit that receives that data will not change or modify the data received.Data silos can negatively affect businesses, and must be eliminated to: 1. Standardise the data across the business units2. Avoid redundant data3. Create rules for data synchronisation across different business units4. Enable communication and make business work smoothlyThe emergence of knowledge graphs Knowledge graphs provide a graph-based semantic structure for organising data. This newly raised paradigm in the knowledge management domain changes the traditional approach of storing operational data in the form of rows and columns inside an RDBMS. Querying high-dimensional operational data using a traditional database system required multiple table joins. This makes the interac- tion complex and time-consuming. This section details key architectural principles for building scalable cross-enterprise artificial intelligence (AI) solutions, alleviating the challenges and isolation imposed by data silos. The unique structural characteristics of knowledge graphs contribute logically to the development of such scalable, cross-enterprise AI solu- tions. The features of knowledge graphs may influence the trustworthiness of these AI implementations. Previous sections "Understanding data silos" and "The emergence of knowledge graphs" have analysed the challenges and constraints caused by data silos, reviewed knowledge graphs as a new paradigm in the domain of data management and semantic querying, and proposed their association within artificial intelligence archi- tecture for building multi-enterprise AI solutions operating at scale.

A. Key Principles of AI Architecture

One of the key principles of AI architecture has been realized in the cross-enterprise AI solution diagram shown in the Introduction. Arguably the single most impactful im- provement on the route of cross-enterprise AI solutions from data silos to knowledge graphs is the crossroad of moves toward creating trusted AI and scalable enterprise AI in 2023. Designing an architecture that supports these moves requires combining principles from both genres of AI. Hence, the primary architectural optimization is the introduction of scalable AI principles into trusted enterprise AI designs to ensure scalability across multiple enterprise locations to tackle

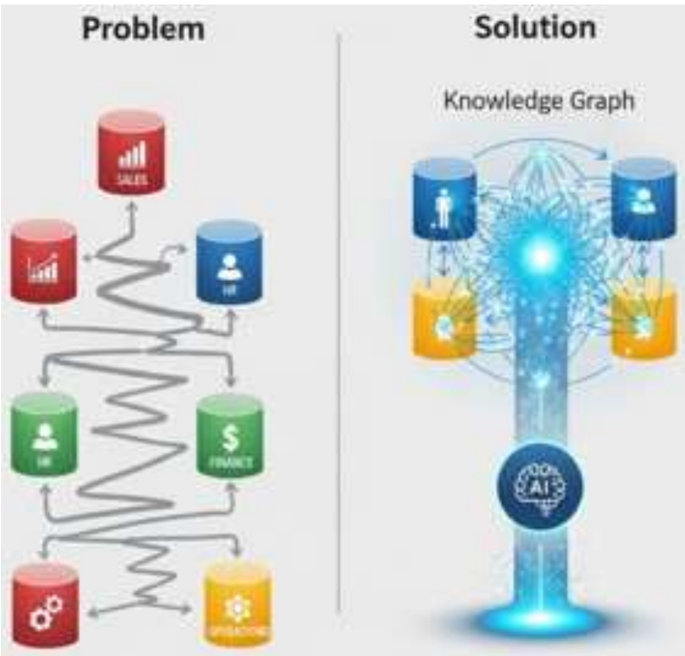

Fig. 4. Data Silos vs. Knowledge Graphs

enterprise scalability challenges, specifically in consolidating otherwise siloed data stores. Combining ideas from scalable AI architecture principles with trusted enterprise AI architecture principles is an effective way of tackling both enterprise- data silos and trust challenges. The second principle of AI architecture is, therefore, the scalable-Enterprise AI principle, which maps naturally onto the boundary between enterprise and scalable AI. The influence of scalable AI architecture in enterprise AI architecture is designed to address the handling of multiple data silos. The data-silos challenge is pivotal. It introduces practical justifications for enhancing traditional enterprise AI designs. Without the data-silos challenge, scal- able AI architecture would not be relevant for enterprise AI. The three steps for incorporating scalable AI architecture into trusted Enterprise AI architecture can be summarized as follows. Firstly, the current accepted Enterprise AI architec- ture must be thoroughly understood, along with its rationale. Secondly, the characteristics of data silos and their impact on Enterprise AI must be elucidated. Finally, the scalable AI architecture principles that can mitigate the data-silos challenge in Enterprise AI are identified.

B. . Integrating Knowledge Graphs into AI Solutions

A knowledge graph plays a pivotal role in building scalable cross-enterprise AI solutions by addressing the challenges presented by enterprise data silos. First, it enables scalability by capturing business semantics and providing interoperability across disparate sources. Despite the underlying data growing linearly with the number of data consumed, the established set of business terms and semantic relationships remains relatively fixed, creating a robust semantic foundation for integrating new data. Second, the graph structure connects diverse data silos, turning the AI enterprise into a convergent force that breaks down barriers and establishes new forms of enterprise collaboration. Combined with other core architectural princi- ples, knowledge graphs help enterprise AI transcend individual organizations and sectors, opening up new opportunities for solving complex business problems. Integrating knowledge graphs lays the foundation for scalable cross-enterprise AI so- lutions and builds the case for them. As noted in the discussion on data silos, enterprises face numerous challenges that hinder their ability to leverage AI. Often siloed in different systems, the data is represented using different models and stored in various formats. A common approach to integrating these data before AI development is using relational databases. However, data integration at the physical data storage level suffers from known shortcomings. While a common data model and format ensures that AI consumes data with a consistent model and format, it does not address the root cause—the diversity of data models and formats—resulting in an architecture that supports integration only within the enterprise. Knowledge graphs, utilizing an ontology for modelling semantics, eliminate this limitation, enabling scalable cross-enterprise AI solutions.

C. Scalability Considerations

The genesis for building cross-enterprise AI solutions capable of effectively addressing data silos and still being scalable to harness the trillions of gigabytes of data constantly created globally across all enterprises is rooted in several fundamental AI architecture principles. These principles include: (1) Integration of a Knowledge Graph Core, (2) Horizontal Scalability for Handling Exponential Data Growth, and (3) Trust as an Underlying Tenet. Each principle demands meticulous attention in the architectural blueprint of any cross-enterprise AI solution, with scalability and trust–but especially scalability–building on the successful implementation of these cornerstones. The integration of a knowledge graph core plays a pivotal role in mitigating the challenges posed by data silos in enterprise data integration.

Unlike relational databases or data warehouses that store data, knowledge graphs codify and structure enterprise knowledge in formats amenable to machine interpretation. The advantage of a knowledge graph lies in capturing and managing semantic relationships—the logical and meaningful associations between data elements distributed across diverse silos. This semantic focus enables unprecedented levels of data interoperability, cross-enterprise query capability, and data analytics workflows.

Equation 3 : AI-Driven Inference Accuracy (AIA)
Goal: relate model accuracy to data and model factors em- phasized in the paper (quality + domain coverage + capacity).
Use a symmetric, bottleneck-sensitive harmonic mean of three normalized drivers:
$Q \in [0,1]$: data quality,
$D \in [0,1]$: domain coverage/breadth,
$M \in [0, 1]$: model capacity normalized to task complexity.
$AIA = Q1 + D1 + M13$ (4)

Use a weighted geometric mean to penalize any weak trust pillar:

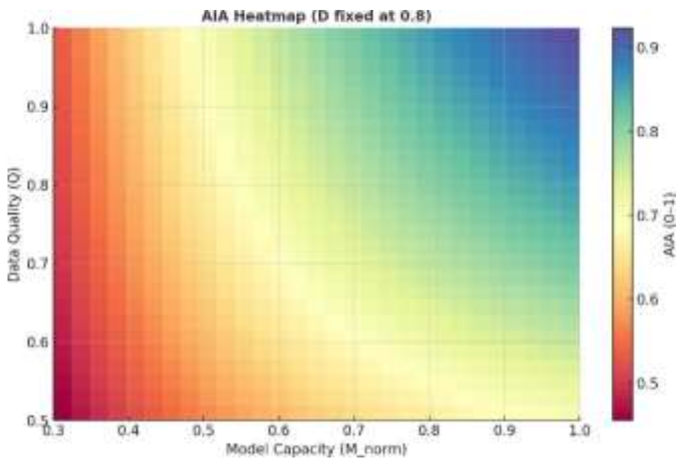

Fig. 5. AIA Heatmap (D fixed at 0.8)

V. ENSURING TRUST IN AI SOLUTIONS

Trust in AI is universally agreed to be vital, because each level of the AI enterprise or AI society must be able to trust the levels below it, as well as satisfy the levels above it. However, in addition to everything said so far on governance and the ethically aligned design of artificial intelligence, other aspects of trust must be recognized too. Transparency and accountability are just two of these. In a world in which cross-enterprise AI solutions generate composite model-driven actions, those actions must be attributable for audit and traceability, not only to their originating enterprise but all the way to the more granular data sources that eventually informed the model-driven decisions. Cross-enterprise AI must architecturally include such transparency and trust- enabling technology. People should be able to quantify the role that AI played in the creation of such composite model- driven actions – for example, by showing the probability of the various events encompassing each achieved prediction, explanation, instruction or recommendation, the AI system's explanation of why it produced the composite model-driven actions, or the why behind the underlying model itself. User organizations could then decide – based on how much they trust AI itself, the cross-enterprise AI architecture, the aforementioned auditability capabilities, and the transparency- enabling information – whether to act on such composite model-driven actions at all. After all, trust is personal; some people are more willing than others to partake of the risks of acting on automatically generated predictions, explanations, instructions, or recommendations.

Equation 4 : Cross-Enterprise Data Trust Score (CEDTS)
Goal: operationalize trust (provenance, explainability, au- ditability, privacy, access control). Derivation.
$CEDTS = (k = 1 \prod mskwk)1/ \sum wk,$ (5) where $s_k \in [0, 1]$ are pillar scores and w_k their weights.
Visual + table. I plotted pillar scores with the overall CEDTS in the title and exposed the input table so you can tune weights/scores.
A. The Importance of Trust in AI

Journal for Re Attach Therapy and Developmental Diversities
eISSN: 2589-7799
2023 December; 6 (10s(2)): 2351-2366

With the emergence of knowledge graphs, architects can tackle data silos and build scalable cross-enterprise artificial intelligence (AI). Knowledge graphs provide a rich form of data integration that synthesizes data from multiple domains into a consistent linked-data model. They expose data that have long been confined to organizational back rooms and offer an effective means to develop the explainability and transparency of AI needed to foster trust. Over time, that fosters an environment in which a much broader set of stakeholders—from partners and customers to regulators—can also tap into a knowledge base that is complemented with connections back to the people and processes that generated it, effectively humanizing those connections. At the same time, the ability to capture processes and people behind the data becomes increasingly important as governing strategies go beyond focusing only on the data and branch out into encom- passing models, services, and algorithms. This is especially true in cross-enterprise AI, where the federated nature of the data introduces a wide array of security, compliance, and privacy concerns.

B. Strategies for Building Trust

Trust is one of the central and most precarious topics when it comes to building artificial intelligence models and AI-based applications. It is often overlooked as many business users will primarily focus on the usefulness of the application or model. While in academic research, particularly in the healthcare, law and criminal, psychology, and education fields, a great deal of emphasis is placed on trust and fairness. A large part of the missing trust can be attributed to the current Cross- Enterprise Artificial Intelligence (CEAI) models. A key issue of the existing Cross-Enterprise Artificial Intelligence is that they lack transparency and openness for the users. It is crucial that users can trace back the decision or suggestion made by a CEAI model, know which kind of data has been used for training, and understand how the model has been trained to increase the trust level in the model.

C. Ethical Considerations in AI Deployments

Ethical considerations in AI are of growing importance and also essential to attaining business trust and the associated benefits of AI adoption. Several areas are particularly relevant in managing ethical risks transparently. Explicit, meaningful human control over, responsibility for, and oversight of AI lifecycle phases; i.e. fairness and avoidance of automated individual or group targeting, or decision making based on AI which lacks alternative dispute resolution or human oversight. Approval and execution of any high-risk AI should be super- vised and controlled by humans. High-risk AI products require human intervention and supervision during the execution phase and before they are placed on the market or put into service.

VI. TECHNOLOGICAL FRAMEWORKS AND TOOLS

The in Knowledge Graph Architectural Framework was introduced in January of 2023 to address the requirements for cross-enterprise AI solutions capable of solving challenges re- lated to scalable data, scalable AI, and their resulting scalable use. That framework contains a range of strategies for building scalable AI solutions—including knowledge graphs. In just a few years, knowledge graphs have become increasingly popular for data scientists and model builders. That sharp adoption rise has been driven by the desire for data-access- layer-defined-critical-path AI, as well as the ability to find new and profitable correlations across any and all data. The cross- enterprise scalability of modern, business-transformation- and digital-transformation-driven AI solutions requires the integra- tion of knowledge graphs with their supporting architectural areas, including but not limited to scalable data, scalable AI, and scalable use. Applying knowledge graphs only to a sin- gle AI-model-training database quickly leads to performance scalability problems as the number of potential AI-roadmap use cases grows.

A. Overview of Relevant Technologies

Knowledge graphs are making their way into modern AI solutions as an enabler to address data silos and the need for scalability and trust. The core principles of AI architecture are summarized, and the capabilities of knowledge graphs that help in building cross-enterprise AI solutions are discussed. Beyond the conceptual role of knowledge graphs in solving data-silo problems and enabling cross-enterprise AI architectures, it is helpful to understand the technological underpinnings. The field of knowledge management contains numerous technological elements that can be combined in specific ways to address specific customer needs. Much has been written about the phases of knowledge management, from knowledge capture to knowledge sharing and knowledge use, and about the different management requirements at the different phases.

Equation 5 : Scalability Efficiency (SE)

Goal: quantify how throughput scales with data volume — central to the paper's scalability theme.

Assume empirical elasticity:

$T \propto V\eta \Rightarrow \eta = d \ln V / d \ln T$.(6) Define efficiency as the proximity of the ideal linear scaling

:$SE = 1 - |1 - \eta|$ (bounded in $(-\infty, 1]$, practical $[0, 1]$). (7)

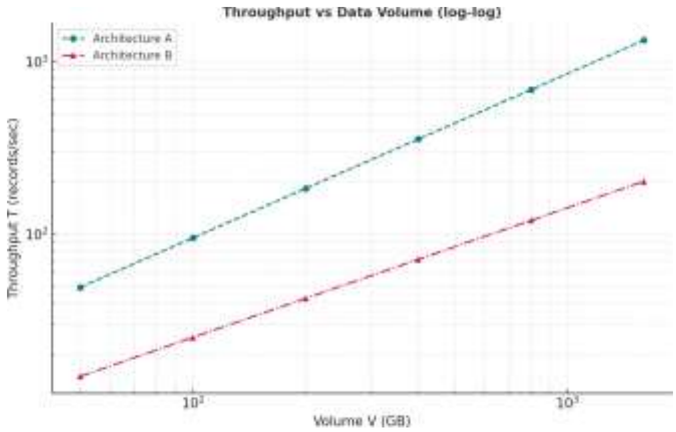

Fig. 6. Throughput vs Data Volume (log-log)

Visuals + table. I fitted from log–log lines for two archi- tectures (near-linear vs sub-linear) and plotted Throughput vs Volume (log–log). The table reports

B. Tools for Building Knowledge Graphs

The process of moving AI away from data silos is defined as establishing a data model that is scalable and transferable across corporate boundaries, thereby enabling inter-company collaboration. Due to the innate strengths of knowledge graphs and their superior ability to facilitate cross-company data alignment, knowledge graph techniques have been embraced. However, the construction of a corporate-level data model involves the utilization of a range of basic knowledge graph technologies. This sections provides an overview of the most important technologies required. These foundational elements constitute the internal framework of a cross-company scalable knowledge graph. With these components in place, strategies for broadening data coverage can then be considered.

C. AI Frameworks Supporting Cross-Enterprise Solutions

Trust in AI solutions also has a significant influence on the choice of frameworks, since some technologies are designed to provide additional architectural guardrails to ensure key AI principles such as transparency and accountability. In 2023, key solutions enabling such architectures for Knowledge Graphs and Cross-Enterprise AI include: Knowledge Graphs are Powerful Drivers of Scalable, Cross-Enterprise AI. Enter- prise Data Integration cannot be achieved by simply connect- ing silos. AI solutions should therefore be architected on top of knowledge graph technologies to overcome the challenge and enable the scaling of AI models across enterprises. A knowledge graph is an evolving, transparent and explainable representation of a business area, spanning multiple organi- zations and including regulatory documentation as an integral part. Its semantic foundation empowers the cross-enterprise scaling of AI models, establishing its role as the foundation of all future business decisions.

dimension	**score 0 1**	**weight**
Provenance completeness (P)	0.8	3
Explainability (E)	0.7	2
Auditability (A)	0.9	2
Privacy compliance (R)	0.85	3
Access control strength (S)	0.75	2

VII. IMPLEMENTATION STRATEGIES

Implementation strategies should follow a phased approach, encompassing organizational change management and success metrics. Practical deployment depends on effective manage- ment and measurement. Cloud integration, graph and AI technologies, and intelligent chatbots are also addressed in the related technological framework section. The analysis of technological challenges, case studies, and success stories appears under challenges and future directions.

A. Phased Implementation Approach

An effective approach to implementing AI solutions within enterprises follows a six-phased process. The initial phase involves establishing a compelling business case demonstrat- ing tangible benefits for the organization. Subsequent phases encompass developing a detailed delivery plan, acquiring and preparing a sample dataset, constructing and validating an AI model, and packaging the model into an operational service. The concluding phase addresses governance and oper- ational considerations aimed at maintaining domain autonomy. A comprehensive approach integrates organizational change management to navigate the broader transformation beyond mere technology adoption.

Additionally, defining metrics to assess the success of the new solutions in meeting stake- holders' evolving needs is crucial. The application of this methodology using knowledge graph and related data fab- ric technologies showcases a cross-enterprise implementation effort. Addressing persistent issues associated with previous generative artificial intelligence solutions, the new architec- ture paradigm demonstrates promise in overcoming well- established challenges, including those arising from data silos.

B. Change Management in Organizations

Change management in organizations is another critical aspect of moving from data silos to knowledge graphs. An en- terprise must bring all involved functions and individuals along the journey and emphasize the new possibilities for building value-added applications. The current data model might not allow for matching finger prints or sets of attributes on attributes, but the knowledge graph approach can. Awareness of the enhanced capabilities will encourage team members to explore new and different line-of-business applications. Using the same core AI Insight Engine across multiple lines of business and in collaboration with partners also enhances the perception from both inside and outside the enterprise, demonstrating a comfort level with protecting the ecosystem and providing access to the right people, organizations, or systems. In large companies with many business units, IT must balance the planning horizon and investment in application development. All blocks must be built at the right time with the right budget. Speaking about cross-enterprise scalability also lets sponsors know that the solution can scale over time and across the organization.

Fig. 7. Enterprise AI Implementation: A 6-Phase Process

C. Measuring Success and Performance

Any endeavour aiming at building consensual, relevant and scalable system architectures—from knowledge to data inte- gration on graphs—is an ambitious one; impulse and passion are essential for both start and sustainability. Still, success depends largely on the ability to quantify, or at least qualify, the journey: it enables measurement and guide the course of the mission, but also helps during review meetings and budget allocation, or simply provides justification for allocated funding, supporting both productivity and devotion. In a cross- enterprise access control client enabling interconnections with external knowledge graphs built in collaboration with an industry partner, measurements and success factors adopted must derive from the type of data silos identified in the company. The partner, a global information technology leader, manages roughly 30% of all digital data on Earth, and stores around 35% of all public cloud information. During the past year more than 16 billion data requests were received and handled, distributed across more than 200 data centres, 60 availability zones and 25 regions, presenting obvious data location challenges. The client addressed in this context is aimed at crossing location, application or even organisational silos encountered. The following considerations are thus ap- plicable to data regarding latency, bandwidth, throughput or simply transfer speeds, and the need to deliver trustworthy information, with transparent processing and role-driven ac- countability.

VIII. CHALLENGES AND LIMITATIONS

The transition from data silos to knowledge graphs is not without challenges and limitations. Data silos arise when different groups of knowledge workers collect, interpret, and curate data in isolation from each other. They are characterized by closed partitioned central knowledge repositories, a lack of ontology reasoning, and a lack of broad semantic integration with external enterprise and public knowledge sources. These conditions eventually lead to poor information integration and limit the broader usability of their data assets. Knowledge graphs fill a gap in the traditional data representation land- scape by offering a structured representation to integrate and organize distributed data and knowledge pools. They enable the connection and formalization of big data assets across company silos, and facilitate the integration of knowledge across companies and industries. Knowledge graphs status representational and organizational improvements in enterprise data management that can help fill an important AI architecture gap. Building on these conclusions, it is important to em- phasize that a successful AI architecture starts by addressing how well an AI solution can scale with the problem at hand. Data silos are big hurdles for scalability in current- day applied AI, which typically addresses narrowly scoped problem domains. Scaling AI toward more general intelligence requires the ability not only to model, reason, and act at increasing levels of complexity, but also to cope well with the needs and requirements it imposes on externalized training and reasoned knowledge. These simple architectural choices im- pose structural requirements on an AI solution. When building for scalability, the nature of the underlying data representation and production has to fulfill operability requirements on how well the AI system can consume, update, and learn with the data. Architectural conditions bridging scaled knowledge needs with a scalable knowledge production process result in a cross-enterprise design. A cross-enterprise AI architecture offers mechanisms that open the doors to external knowledge that lies dormant across other parts of the company or across company boundaries. Technologies such as knowledge graphs are therefore critical to unlock the generation of training data across enterprises.

A. Technical Challenges

Inherent Technical Challenges Data silos are repositories of enterprise data that are controlled by one group and isolated from the rest of the organization. Data silos prohibit data sharing and reuse. A siloed data architecture leads to an increased cost in developing and operating machine-learning models for AI. Over time, multiple new applications have been developed for overlapping business use cases. Data sources across the various applications are isolated and duplication of data and processes is common within the organization. In contrast, a knowledge graph aggregates enterprise data into a single connected data layer. The knowledge graph enables AI as new applications are developed for evolving business needs. Data is continuously added at scale, across functions and enterprises. An architectural approach for AI leverages knowledge graphs to break down data silos for scalability and trust. Organizations face technical challenges when preparing for such transformations. Scalability and flexibility are important, especially when the knowledge graph or enterprise data needs change. An evolutionary architecture facilitates continuous design. Requirements include high availability, scalabilty and reliability. The principle of loose coupling enables change management and governance. The architecture should support the integration of additional AI cruicial enterprise data sources. Provide a transparent view of the provenance and lineage of the data. Advanced Technical Requirements The roadmap from data silos to knowledge graphs considers technologies for hosting data, code and objects. A storage layer facilitates the ingestion, curation and exploration of connected data assets. An ontology designer promotes the understanding, ideation and validation of knowledge-based additions. An ingestion tool facilitates the mapping of enterprise data to the representation format selected for the knowledge graph. A reasoner enables the inference of additional insights. Its ability to support multiple representation formats allows its usage across a collection of knowledge graphs. A complex ontology requires the analysis of biological pathways and molecular mechanisms at a detailed level. An API supports programmatic access to the knowledge graph for all applications. The architecture facilitates data integration across variety, volume and velocity. Pre-built applications streamline the reporting and analysis of data stored in the knowledge graph. An insightful dashboard presents a graphical view of hosting statistics.

Equation 6 : Cross-Enterprise AI Impact Index (CEAI)
Goal: combine benefit, adoption footprint, and residual risk for cross-enterprise programs.
Normalize benefits $B_{norm} \in [0, 1]$ (e.g., KPI uplift), use adoption $A_{norm} \in [0, 1]$, and normalized residual risk $R_{norm} \in (0, 1]$. Higher is better:

$$CEAI = R_{norm}B_{norm} \cdot A_{norm}. \quad (8)$$

Visual + table. A bar chart compares Pilot → Dept rollout → Multi-enterprise, with the backing table opened.

B. Cultural and Organizational Barriers

Cultural and Organizational Aspects. Transitioning from Data Silos to Knowledge Graphs is of course not only a technical but also a cultural and organizational challenge. Routines that have grown over long periods of time must be questioned and changed, which is often perceived quite negatively by employees. The silo thought pattern must be unlocked in order to think in a cross-enterprise manner.

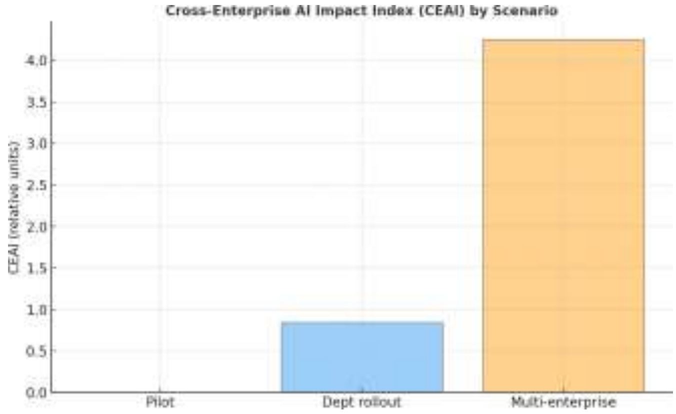

Fig. 8. Cross-Enterprise AI Impact Index (CEAI) by Scenario

architecture	eta elasticity	SE
A (near-linear)	0.9500000000000003	0.9500000000000003
B (sub-linear)	0.7500000000000003	0.7500000000000003

Transparency—from the work performed to the data made available—is possibly not one of the strengths of companies, and thus it is often difficult to convince colleagues from Freedom Data Silos to share data. The other group of Freedom Data Silos, which is concretely responsible for the data silos, tends to be rather reluctant towards the opening of their data silo. Perhaps the users also do not ask for such an opening, so that no business need is seen here either. Especially with very old IT systems, the effort is often underestimated, which is why a plan for opening must be established at the political level in the company. Some of the data silos have already been opened, for example via Cloud technologies. Many problems and challenges listed above are also described in a document from architecting scalable and trusted AI solutions using cross-enterprise knowledge graphs. Support through a modern technological architecture comprising microservices cannot be underestimated. Experiences in companies show that the release of comprehensive data is much easier if a catalog of APIs is available. The scarcest resource in companies is certainly the integration personnel. Nevertheless, the result speaks for itself: Even if it is a radical approach, quite quickly scalable Artificial Intelligence services and functions can be realized and operated internally and externally for partners or customers, which seem hardly feasible without such an approach. To protect confidential data, it is recommended to provide these functions in different house clouds or, in the case of very sensitive data, to deploy in a completely isolated instance regardless of the low WebAPI latency.

C. Regulatory and Compliance Issues

Built on previously proposed principles of AI architecture and the features of knowledge graphs, **AI architecture principle no. 5** addresses the topic of trust. The essence of trust is that all stakeholders must be able to rely on the decisions made by the solution. When people knowingly accept proposed decisions, the risks of costly mistakes and faulty business decision making can be considerably reduced. Assuming a lack of trust in AI results, people will indepen- dently re-run the analysis with similar methods and examine the results for consistency. When similar results are reached, they may be considered trustworthy. Otherwise, the fact that the results are not consistent weakens the foundation on which operative actions are typically based. The generation of trustworthy AI results starts with setting the appropriate solution architecture. For example, a cross-enterprise AI solu- tion that combines machine learning pipelines with the linking capabilities of knowledge graphs is designed to improve not only performance, scalability, and costs, but also trust.

IX. FUTURE DIRECTIONS

Recent years have seen an explosion in the development of generative and multimodal AI models, complemented by the availability of ever-larger datasets. While many early use cases relied on synthetic data, subsequent implementations increasingly utilize real-data prompts from previous request- response interactions to tune answer quality. This trend (plus the link to popular AI use cases) encourages a wide variety of industries to increasingly share data, defining new boundaries beyond numerous current silos. An initial and explicit integra- tion of this data with suitable knowledge graphs would unlock powerful and trusted cross-enterprise applications for AI use cases of tomorrow. This is particularly true for the multi-agent applications presently trending that require as much available information as possible to operate—when both the questions and answers potentially traverse several enterprise boundaries. In addition to boosting the openness of these APIs, adequate organization and publication of knowledge repositories for external consumption will benefit emerging business models and new collaboration scenarios with external partners.

A. Emerging Trends in AI and Knowledge Graphs

Major Artificial Intelligence (AI) and knowledge graph trends are surveyed. Combining them for cross-enterprise scalability and trust is proposed. The term ”Artificial Intelli- gence” (AI) is currently associated with automatic generation of answers. ChatGPT, an AI-based database that answers questions formulated in natural language, is a good example. Business Intelligence (BI) users frequently require answers to questions such as ”What is the profitability of each product over time?”, ”Which product category in Spain has suffered the most from inflation?” or ”What is the top 10 trend of our sales during the last years?” However, BI tools usually present results in static dashboards, and users do not obtain answers in a natural-language chat. Obviously, users need to combine several graphs in different dashboards—a task that could require an engineering degree. However, ChatGPT and its derivatives attempt to answer these questions in natural language, using data from an enterprise.

B. The Future of Cross-Enterprise Collaboration

”The data silo approach can no longer be maintained if the goal is to design AI solutions that scale beyond a single enterprise in a trusted way. As enterprises massively engage in external collaboration, the aggregate knowledge graph of an industry and/or its corporate ecosystem will offer a far more valuable asset than the graph of a single enterprise. It will consolidate proprietary enterprise knowledge assets by removing the blind spots and incomplete perspective of the reality faced by a single enterprise. It will enable not only a more comprehensive understanding of common issues faced by enterprises, but also the modeling and systematic sharing of common solutions for similar challenges faced by different companies. Lastly, it will enable the design of scalable AI architectures for cross-enterprise use, shifting from a system-based to a stark needs-based model of interaction with Artificial Intelligence—the substitution of the Human Machine Interface by a Human Needs Interface.

C. Potential Research Areas

Emerging large-scale multilingual, multimodal, multi-task AI architectures incorporating knowledge graphs suggest new directions for knowledge graph research. Solutions will shift from intra-enterprise boundary to cross-enterprise boundary collaboration and data sharing. Innovation will be necessary in merging knowledge into architectures to optimize cost and latency versus performance trade-offs. Data sharing addresses Enterprise Data Silos and Unlocking Enterprise Data for Large Language Models. While multi-task foundations reduce the need for labeled data, some is essential, spurring interest in reducing data and labels through data augmentation, self- supervised and few-shot learning, and active learning. Trust considerations revolve around model explainability, bias, and adherence to constraints, implemented through transparency, accountability, and AI ethics.

X. Conclusion

Enterprise data silos are characterized by rigid and volu- minous archives that are neither readily accessible nor inter-operable. Across-lifecycle enterprise solutions that integrate knowledge engineering for policy and data capture with data engineering for AI-ready target data build knowledge graphs to unlock data silos and create scalable AI. Trust is essential to successful AI adoption. A cross-enterprise AI architecture is designed that applies learned lessons from the 2019 Covid pandemic on risk and resilience to support transparent, ac- countable, fair, unbiased, replicable, and compliant AI for key enterprise decision processes, policies, stakeholders, and ethical requirements. The traditional approach of using policy risk matrices to analyse an organisation’s risk position by generating risk-impact scatterplots is enhanced using a set of knowledge graphs. This overcomes the limitations of small datasets and enables modelling, prediction, and validation with a high level of trust. The knowledge graph-based risk radar allows users to classify risks and semi-automatically generate possible responses to new emerging risks. It also allows for the semi-automatic generation of anonymous datasets and data syntheses to address data sparsity and trust issues. Knowledge graph questioning supports ’what-if’ analysis, enterprise risk governance / control evaluation, and identification of emerging risk clusters.

Fig. 9. Risk Management Approach Evolution

A. Summary of Key Insights and Final Thoughts

The purpose of the document is to summarize the rationale for moving from data silos to knowledge graphs and their relevance for architecting cross-enterprise AI solutions that scale and can be trusted. Within the overall paper, this closing section recreates the essence of the abstract and introduction. Data organizations are ubiquitous in modern enterprises, and they could be operating at a local, departmental, divisional, or cross-enterprise level, apart from those associated with supply chain partners and supporting infrastructure of ven- dors and service providers, both upstream and downstream. These are generically termed data silos, owing to their in- herently fragmented and compartmentalized nature. Enterprise data integration continues to be a top priority for CIOs to enable business managers to get a consolidated picture of their important business entities and relationships. Despite substantial advances in data, API, and messaging technologies that help build flavours of enterprise systems and consolidate data views for organizations, these systems find it difficult to support scalable AI application needs. AI architectures invoke knowledge bases to provide enough depth and precision for interpretability and accountability as key components for a trusted AI solution, besides realistic portrayals of busi- ness entity constructs for cross-enterprise scalability through intelligently-linked data. In data management and data science communities, the knowledge graph concept is often invoked here and it is, indeed, well regarded for its contribution to high-performance enterprise knowledge bases. The foray of knowledge graphs into the AI arena is relatively recent and it is beginning to acquire the sheen of the panacea for addressing enterprise silos that AI systems frequently encounter. The journey of the knowledge graph from its origins in Google Search and Facebook knowledge bases to its transformative potential for future AI applications has much to offer to academia and industry by way of lessons learned and pointers for the future.

REFERENCES

[1] Motamary, S. (2023). Integrating Intelligent BSS Solutions with Edge AI for Real-Time Retail Insights and Analytics. European Advanced Journal for Science & Engineering (EAJSE)-p-ISSN 3050-9696 en e- ISSN 3050-970X, 1(1).
[2] Arthur, L., Costello, J., Hardy, J., O’Brien, W., Rea, J., & Ganev, G. (2023). On the challenges of deploying privacy-preserving synthetic data in the enterprise. arXiv. https://arxiv.org/abs/2307.04208
[3] Mo¨kander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2023). Auditing large language models: A three-layered approach. arXiv. https://arXiv.org/abs/2302.08500
[4] Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., & Wu, X. (2023). Unifying large language models and knowledge graphs: A roadmap. arXiv. https://arXiv.org/abs/2306.08302
[5] Sheelam, G. K. (2023). Adaptive AI Workflows for Edge-to-Cloud Processing in Decentralized Mobile Infrastructure. Journal for Reattach Therapy and Development Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2). 3570ugh Predictive Intelligence.
[6] Yan, C., et al. (2023). Construction process and application of enterprise knowledge graph from heterogeneous data sources. Frontiers in Big Data, 1. https://doi.org/10.3389/fdata.2023.1278153
[7] Lakkarasu, P. (2023). Generative AI in Financial Intelligence: Unrav- eling its Potential in Risk Assessment and Compliance. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 241-273.
[8] Yan, C., et al. (2023). Construction process and application of enterprise knowledge graph from heterogeneous data sources. Frontiers in Big Data, 1. https://doi.org/10.3389/fdata.2023.1278153
[9] Tallberg, J., Erman, E., Furendal, M., Geith, J., Klamberg, M., & Lundgren, M. (2023). The global governance of artificial in- telligence: Next steps for empirical and normative research. arXiv. https://arXiv.org/abs/2305.11528

[10] Kummari, D. N. (2023). Energy Consumption Optimization in Smart Factories Using AI-Based Analytics: Evidence from Automotive Plants. Journal for Reattach Therapy and Development Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2), 3572.
[11] Ceri, S., Baldazzi, T., Bellomarini, L., & Sallinger, E. (2023). Fine- tuning large enterprise language models via ontological reasoning. (Preprint). Knowledge Graphs and Enterprise AI: The Promise of an Enabling Technology. June 2023.
[12] Gadi, A. L. The Role Of AI-Driven Predictive Analytics In Automotive R&D: Enhancing Vehicle Performance And Safety.
[13] (Article) "Evolving enterprise architecture governance to embrace AI" (2023). Architecture & Governance Magazine.
[14] Kalisetty, S., & Singireddy, J. (2023). Agentic AI in Retail: A Paradigm Shift in Autonomous Customer Interaction and Supply Chain Automa- tion. American Advanced Journal for Emerging Disciplinaries (AAJED) ISSN: 3067-4190, 1(1).
[15] (Workshop) "Enterprise Knowledge Graphs using Large Language Mod- els" (2023, October). In Proceedings of the Workshop on Enterprise Knowledge Graphs using Large Language Models. ACM.
[16] (Report) "The state of AI in 2023: Generative AI's breakout year." (2023, August 1). McKinsey & Company
[17] Lakkarasu, P. (2023). Generative AI in Financial Intelligence: Unrav- eling its Potential in Risk Assessment and Compliance. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 241-273.
[18] (Report) "The state of AI in 2023: Generative AI's breakout year." (2023, August 1). McKinsey & Company
[19] (Article) "The Knowledge Graph Advantage: How smart companies are using knowledge graphs to power AI and drive value." (2023). Masood, A. https://medium.com/@adnanmasood/the-knowledge-graph- advantage-how-smart-companies-are-using-knowledge-graphs-to-power- ai-and-drive-59f285602683
[20] Somu, B. (2023). Towards Self-Healing Bank IT Systems: The Emer- gence of Agentic AI in Infrastructure Monitoring and Management. American Advanced Journal for Emerging Disciplinaries (AAJED) ISSN: 3067-4190, 1(1).
[21] Gupta, R., & Srinivasa, S. (Eds.). (2023). Proceedings of the Workshop on Enterprise Knowledge Graphs using Large Language Models (EKG- LLM 2023), co-located with CIKM 2023, Birmingham, UK, October 22, 2023 (CEUR Workshop Proceedings, Vol. 3532). CEUR-WS.
[22] Koppolu, H. K. R., Sheelam, G. K., & Komaragiri, V. B. (2023). Au- tonomous Telecommunication Networks: The Convergence of Agentic AI and AI-Optimized Hardware. International Journal of Science and Research (IJSR), 12(12), 2253-2270.
[23] Mohanty, A. (2023). EduEmbedd – A knowledge graph embedding for education. In R. Gupta & S. Srinivasa (Eds.), EKG-LLM 2023 (CEUR- WS, Vol. 3532).
[24] Meda, R. (2023). Data Engineering Architectures for Scalable AI in Paint Manufacturing Operations. European Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 1(1).
[25] Subramaniam, P., Khurana, U., Srinivas, K., & Samulowitz, H. (2023). Related table search for numeric data using large language models and enterprise knowledge graphs. In R. Gupta & S. Srinivasa (Eds.), EKG- LLM 2023 (CEUR-WS, Vol. 3532).
[26] Kalisetty, S., & Singireddy, J. (2023). Agentic AI in Retail: A Paradigm Shift in Autonomous Customer Interaction and Supply Chain Automa- tion. American Advanced Journal for Emerging Disciplinaries (AAJED) ISSN: 3067-4190, 1(1).
[27] Kulkarni, A., Ramanathan, C., & Venugopal, V. E. (2023). Cognitive Retrieve: Empowering document retrieval with semantics and domain- specific knowledge graph. In R. Gupta & S. Srinivasa (Eds.), EKG-LLM 2023 (CEUR-WS, Vol. 3532).
[28] Sewak, M., Emani, V., & Naresh, A. (2023). CRUSH: Cybersecurity research using universal LLMs and semantic hypernetworks. In R. Gupta & S. Srinivasa (Eds.), EKG-LLM 2023 (CEUR-WS, Vol. 3532).
[29] Lahari Pandiri, "Leveraging AI and Machine Learning for Dynamic Risk Assessment in Auto and Property Insurance Markets," International Journal of Innovative Research in Electrical, Electronics, Instrumen- tation and Control Engineering (IJIREEICE), DOI 10.17148/IJIREE-
ICE.2023.111212
[30] Markowitz, E. S., & Galstyan, A. (2023). StATIK+: Structure and text for inductive knowledge graph modeling and paths towards enterprise implementations. In R. Gupta & S. Srinivasa (Eds.), EKG-LLM 2023 (CEUR-WS, Vol. 3532).
[31] Inala, R. Revolutionizing Customer Master Data in Insurance Technol- ogy Platforms: An AI and MDM Architecture Perspective.
[32] Sequeda, J., Allemang, D., & Jacob, B. (2023, November 13). A benchmark to understand the role of knowledge graphs on large language model's accuracy for question answering on enterprise SQL databases. arXiv (2311.07509).
[33] Meyer, L.-P., Stadler, C., Frey, J., Radtke, N., Junghanns, K., Meiss- ner, R., Dziwis, G., Bulert, K., & Martin, M.

(2023). LLM-assisted knowledge graph engineering: Experiments with ChatGPT. arXiv (2307.06917).

[34] Allen, B. P., Ilievski, F., & Joshi, S. (2023). Identifying and consolidating knowledge engineering requirements. arXiv (2306.15124).

[35] Nandan, B. P., & Chitta, S. S. (2023). Machine Learning Driven Metrology and Defect Detection in Extreme Ultraviolet (EUV) Lithog- raphy: A Paradigm Shift in Semiconductor Manufacturing. Educational Administration: Theory and Practice, 29 (4), 4555–4568.

[36] Tamasˇauskaite˙, G., & Groth, P. (2023). Defining a knowledge graph development process through a systematic review. ACM Transactions on Software Engineering and Methodology, 32(1). https://doi.org/10.1145/3522586