International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

Integrating Generative AI and Agentic Systems into Enterprise Cloud Architecture: Frameworks and Governance Models

Gadepalli Sri Pratyak Aditya Swaprakash.

Sr Technical Director and Sr Solutions Architect, Independent Researcher, USA

PrakashaqA.GadepalliSP@outlook.com

ABSTRACT: Artificial intelligence (AI) is an incredibly fast-developing field, characterized by the creation of generative models and agentic systems that are able to independently complete complicated tasks, create content and aid in decision-making processes. Although a lot of recent talk about generative AI has picked up since 2020, governance principles, architectural patterns, and fundamental rules have been established beforehand. In this paper, a thorough model is proposed to deploy generative AI and agentic systems in enterprise cloud architecture, based on the research and industry practices prior to 2020. It provides an overview of architectural layers, data pipelines, orchestration patterns and governance paradigms required to deploy scalable and secure. The research also delves into issues of data privacy, model interpretability, operational risk, and compliance. This paper aims to integrate the key insights from previous studies in cloud computing systems, machine learning systems, and autonomous agents to offer a well-defined strategy for enterprises in developing AI-based cloud computing architectures, with a strong business control intention. In this paper, the authors attempt to combine the prominent findings of the aforementioned bodies of literature related to cloud computing systems, machine learning systems and autonomous agents to offer an structured strategy for the enterprises searching for methods adopting AI-based cloud architectures while keeping robust business control intention.

KEYWORDS: Generative AI, Agentic Systems, Enterprise Cloud Architecture, Governance Models, Machine Learning Systems, Cloud Computing, Autonomous Agents.

I. INTRODUCTION

Artificial Intelligence has been a gradual shift in enterprise systems, supported by the developments in machine learning, distributed computing and the engineering of data. In 2020, organizations have already been using early-stage predictive analytics models, recommendation systems, and predictive modeling techniques. At the same time, cloud computing had become a mainstream approach to scalable and elastic infrastructure, featuring elastic resource provisioning and distributed computing. The concepts behind agentic systems were preceded by earlier AI investigation (Iqbal and Saleh, 2020) with autonomous or semi-autonomous entities in the environment which are able to perceive, reason and act within that environment. Likewise, generative models like the Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) were the precursor for current generative AI (Thomas, 2020). This paper seeks to connect these core concepts to enterprise cloud architecture, offering a formal cloud framework to include generative and agentic capabilities. Emphasis is placed on pre-2020 literature to develop solid insights that are historically rooted.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

II. BACKGROUND AND LITERATURE REVIEW

2.1 Cloud Computing Foundations

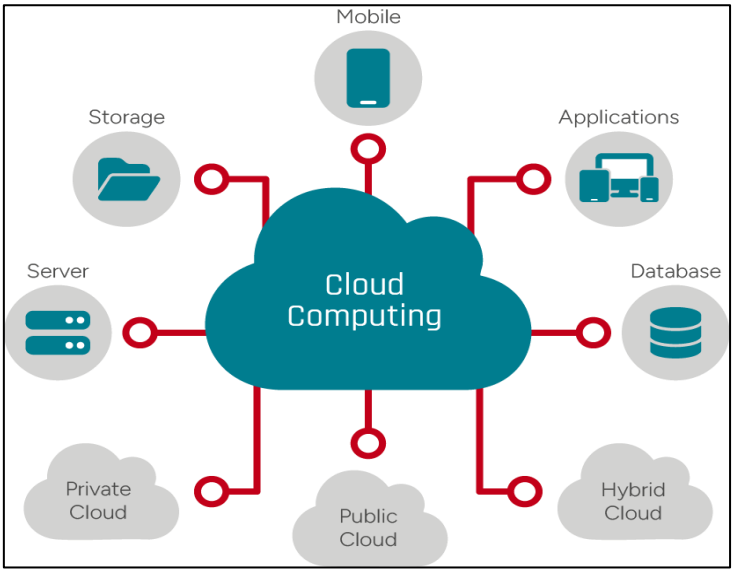

Figure 1: Cloud Computing
(Source: www.gfos.com)

Cloud Computing is the provision of computing services via non-exclusive infrastructure, platform, and software offerings on-demand. With enterprise adoption of cloud platforms, it is possible to take advantage of distributed data processing, high availability and cost. Further Kaur (2020) showed that Cloud Computing offered many economic and operational benefits, particularly with elasticity and scalability as an enabling feature for data-intensive applications. These properties are crucial for AI workloads, especially in the context of training and inference with large models.

2.2 Machine Learning Systems in Production

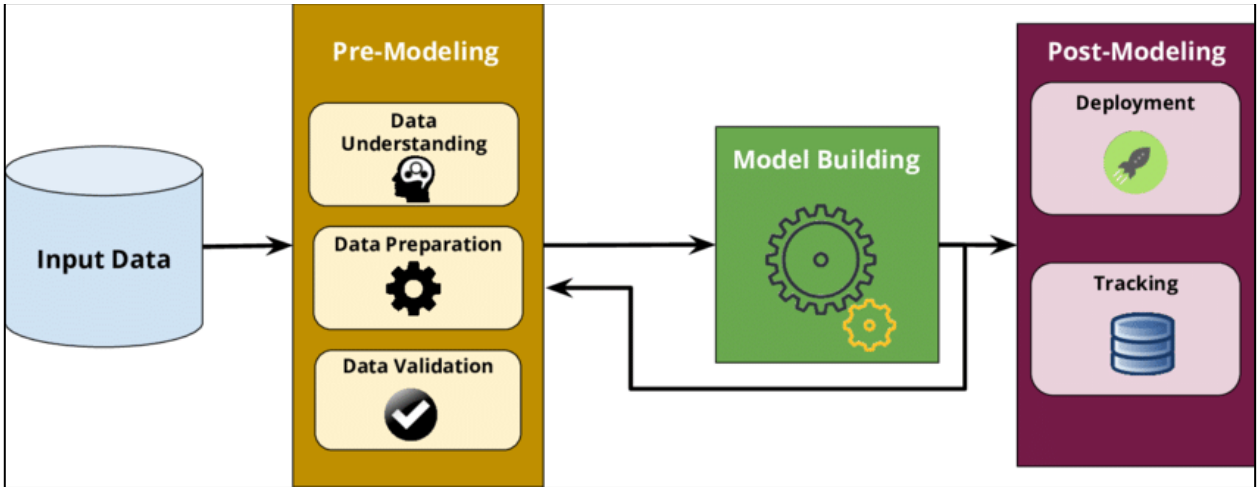

Figure 2: Machine Learning Systems in Production
(Source: www.researchgate.net)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

Much of the work before generative AI was big on deploying machine learning systems in production. Balaganski (2015) introduced the concept of 'technical debt' in the context of ML systems to highlight the challenges with managing data dependencies, developing models, and integrating systems. Arora (2017) presented an ML pipeline platform called TensorFlow Extended (TFX), which focuses on the need for data validation, model training and evaluation, and deployment in scalable architectures.

2.3 Generative Models

Generative models are designed to be trained to model the underlying distribution of data, in order to produce new data samples. Early approaches include:

- Variational Autoencoders (VAEs)
- Generative Adversarial Networks (GANs)

These models proved to be capable of generating images, text, and other types of data the seminal concepts of future generative AI systems.

2.4 Agentic Systems and Autonomous Agents

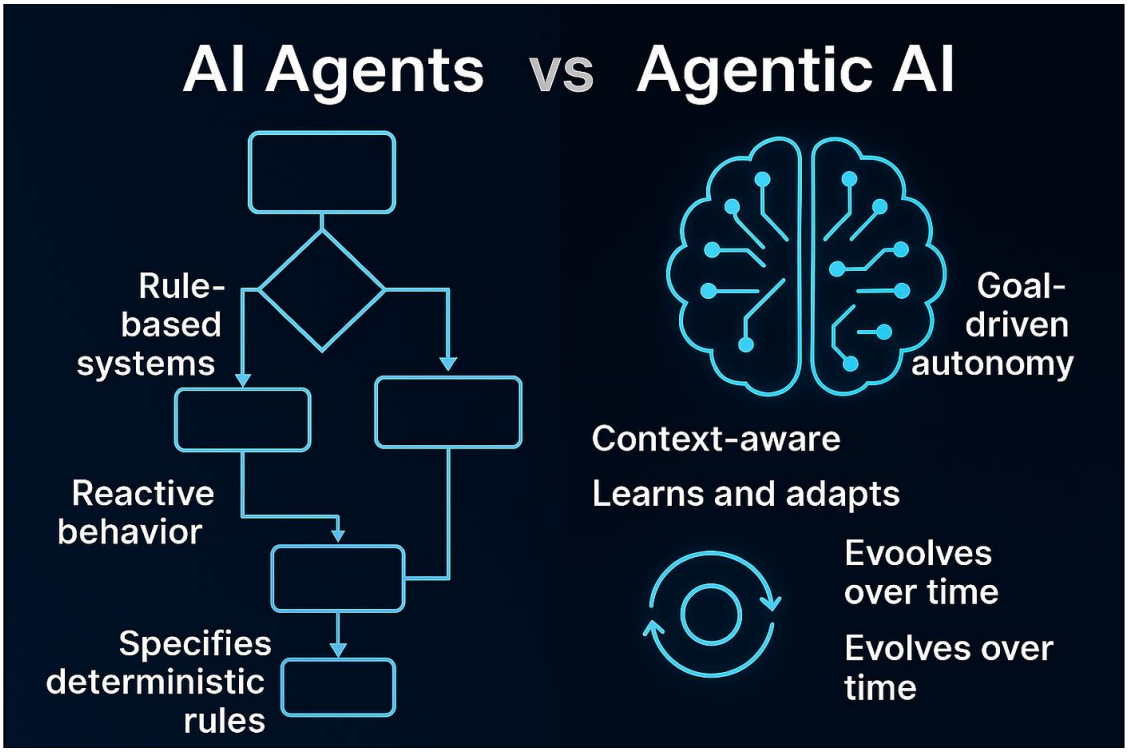

Figure 3: Agentic Systems and Autonomous Agents
(Source: medium.com)

The idea of agent-based systems has been widely researched in AI. According to Benlian et al., (2018) agents are entities that can act independently to accomplish certain goals. Multi-agent systems (MAS) allow a problem to be tackled in a distributed fashion by cooperating and coordinating multiple agents. Nippatla (2018) group Agents into the following classes: simple reflex agents, model based agents, goal oriented agents, and utility based agents. These groupings can guide the design of the architectures of today.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

III. ENTERPRISE CLOUD ARCHITECTURE FOR AI INTEGRATION

3.1 Architectural Layers

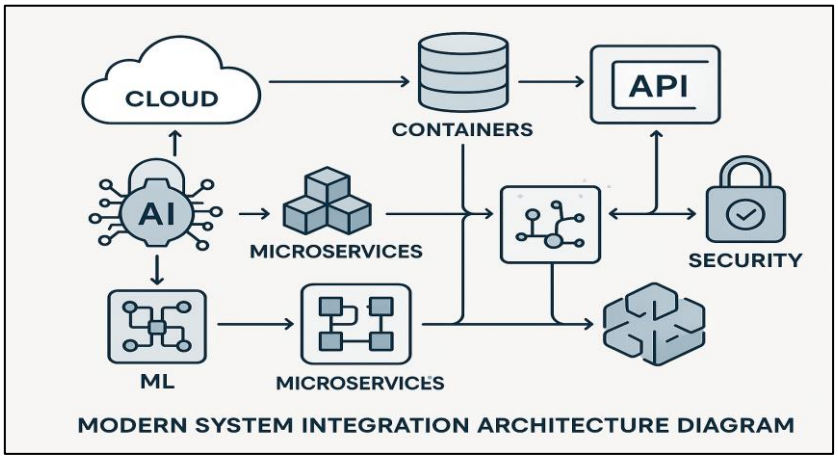

Figure 4: Enterprise Cloud Architecture for AI Integration
(Source: towardsdev.com)

Generative AI and agentic systems can be integrated into a robust enterprise architecture, which can be modeled in several logical layers that work seamlessly together, allowing for scalability, flexibility and maintainability. Data ingestion, storage, and data preprocess in the data layer, and this is usually performed through a distributed file system and a data lake that can be given the structured and unstructured data. On top of that, above the model there is the training, validation and versioning of AI models where artifacts of the model are made reproducible and controlled for its evolution (Ferreira et al., 2017). The service layer typically offers consistent interfaces, in the form of APIs and microservices, which allow for easy interaction with models that are trained and ready to use for actual inference. The orchestration layer manages processes, pipelines and relationships between independent agents, facilitating the scheduling of actions in a coherent and efficient way. The application layer wraps around the end-users and enterprise systems, bringing AI capabilities into business processes and user faces. It is layered architecture works with the microservices principles, which allows to scale and evolve the parts in an stand alone manner.

3.2 Data Pipelines and Infrastructure

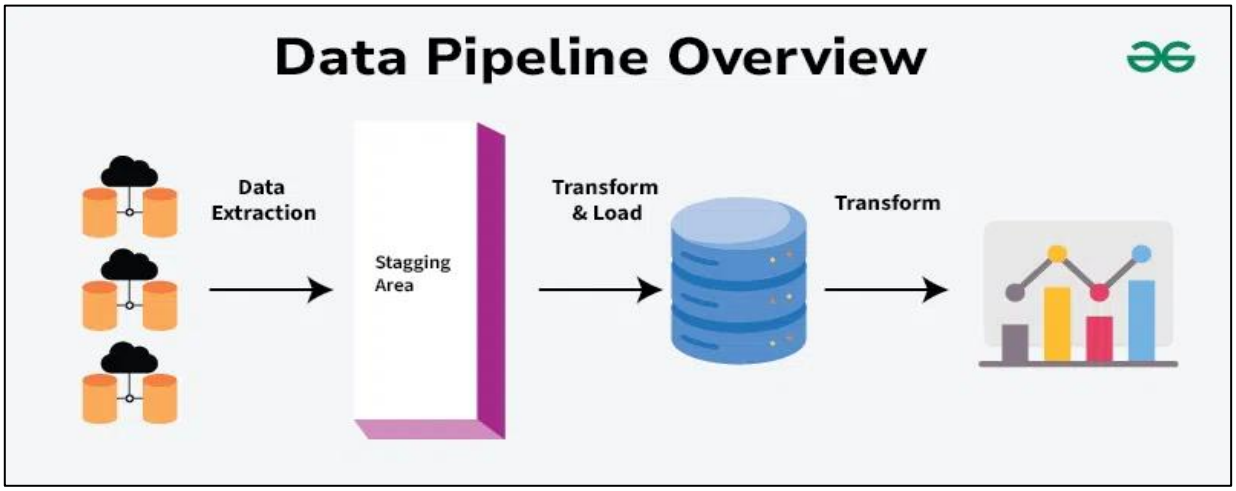

Figure 5: Data Pipelines and Infrastructure
(Source: www.geeksforgeeks.org)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

Generative models and agentic systems rely heavily on data pipelines for seamless and continuous data flow throughout the system. Before 2020, the research and industry were focused on Import/Export/Transform processes to tackle the machine learning problem & deliver ready data. The main objective of these pipelines was to process high volume of data with maintaining the quality of data by implementing validations and data lineage tracking. Real-time data streaming grew in significance when apps that need low-latency responses required it, and in addition to batch processing, it is now crucial for some applications. Distributed streaming platforms like those developed by Wareham et al., (2014), Apache Kafka, allowed organizations to take in and method data real-time, which profited AI applications that addressed the requirement for dynamism. Infrastructure components like these can all help to ensure that generative and agentic systems are running on information that is current and accurate, maintaining performance and reliability.

3.3 Model Deployment and Serving

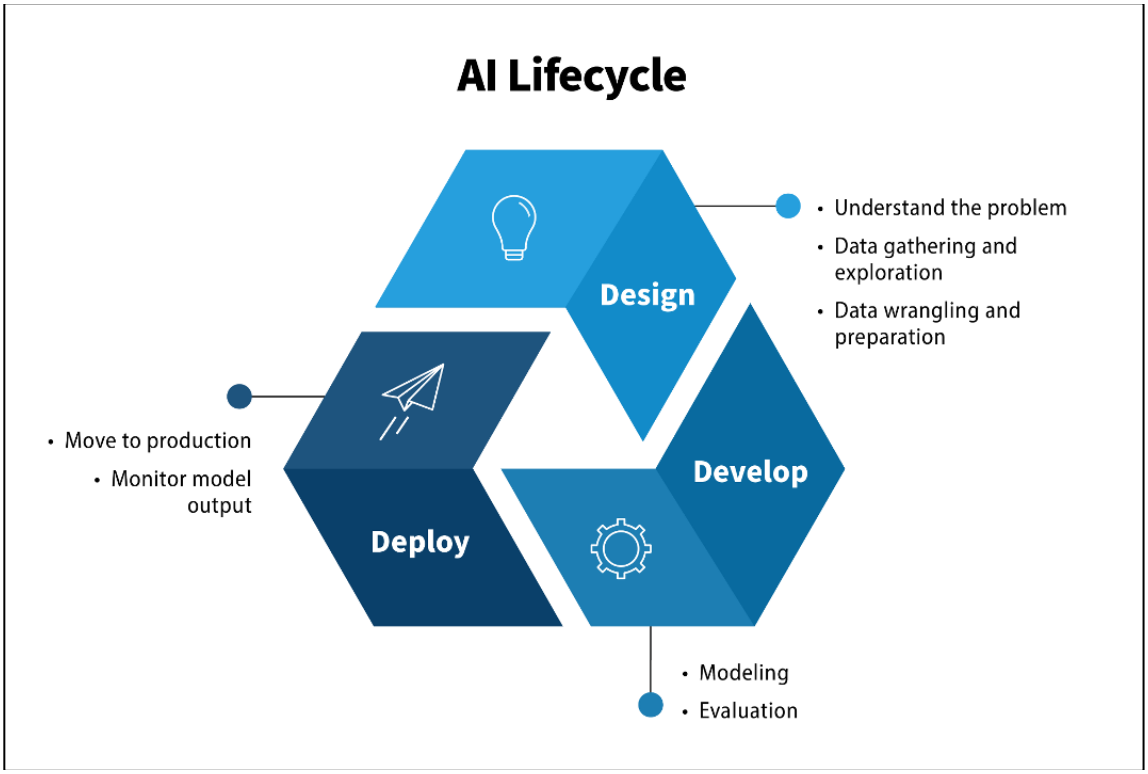

Figure 6: Lifecycle of AI systems
(Source: coe.gsa.gov)

The deployment and serving of models are an integral part of the lifecycle of AI systems, and are considered critical points when the trained models are made available for use in the real world. Normally this entails wrapping models in containers, so that they look the same from the development and production sides. The models can be packaged along with their dependencies using containerization technologies like Docker, for portability and reproducibility. Orchestration platforms, like Kubernetes, offer investment management features like load balancing, scaling and fault tolerance to help handle the applications within a containerized surroundings at scale (Pantazis and Gerber, 2018). These features are crucial for enterprise-level applications, especially those operating in various scenarios with fluctuating workloads and ensuring both high availability and low latency. APIs enable organizations to connect with models, embedding AI into other systems and apps and enriching their capabilities and user experiences.

3.4 Agent Orchestration

In agentic systems, they need sophisticated orchestration mechanisms, in order to orchestrate multiple autonomous agents or semi-autonomous agents. These mechanisms guarantee that jobs are accomplished effectively and all facets of the system are held together. Typically, workflow engines, or rule-based systems, are used to control the dependency

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

and sequence of agent actions. Depending on the context of the application a different coordination strategy can be selected. Centralized orchestration uses one controller to control the agents' behavior, giving a Controlling Agent a high level of control but there is a potential single point of failure. Contrast to the decentralized forms, which enable agents to work more independently, relying on communication protocols and common aims for achieving coordination Lauterbach, (2019). Another approach that can improve scalability and adaptability is by the use of market-based mechanisms (market places), where agents ask each other about their tasks. The orchestration strategies used here are essential in allowing a complex behavior in the enterprise scale agentic systems.

IV. FRAMEWORK FOR INTEGRATION

4.1 Design Principles

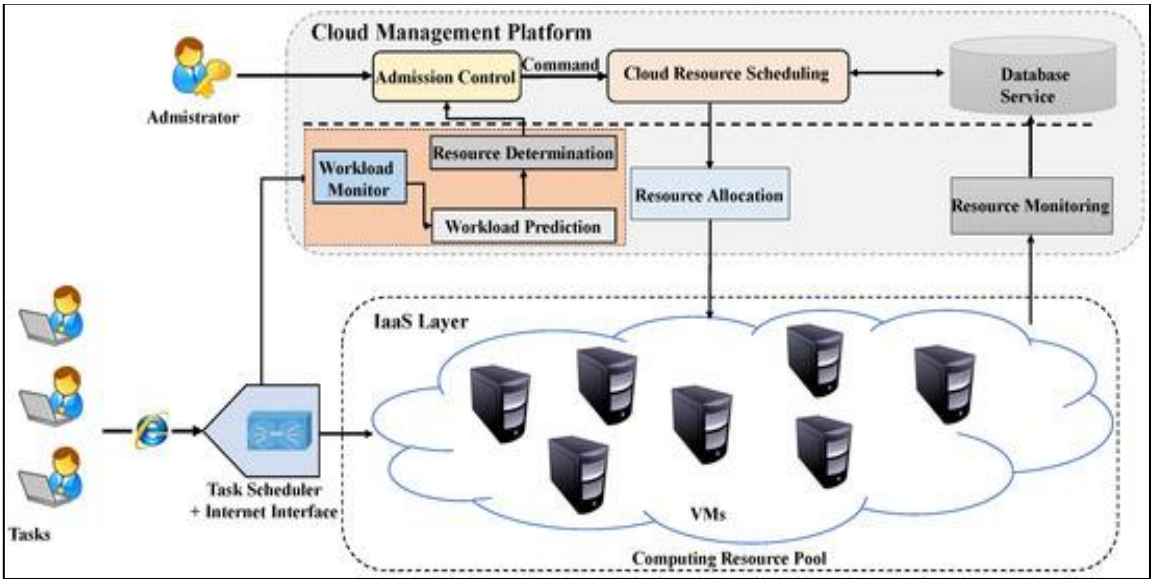

Figure 7: Elastic nature of a cloud-based platform
(Source: www.mdpi.com)

The integration framework follows principles to make enterprise systems that use AI work well and be sustainable. Scientific computing involves using the elastic nature of a cloud-based platform to dynamically allocate the resources according to the workload. The user needs are met by scaling out (and down during times of rest) cloud-based systems. Modularity is achieved by adopting loosely coupled organization which allows independent development, deployment and maintenance of loosely coupled modules (Gelev, 2017). Interoperability is promoted by using common interfaces and protocols, allowing different parts of the architecture to communicate easily. Specific design and recovery mechanisms ensure resiliency in the system as required to be reliable despite failures. All of these principles offer a framework for the creation of strong and flexible AI systems in enterprise cloud architectures.

4.2 Reference Architecture

A reference architecture includes:

- Data ingestion pipelines
- Model training environments
- Model registry and version control
- API gateway for inference
- Agent orchestration engine

4.3 Workflow Integration

Workflows combine data processing, model application and agent actions. Apache Airflow (pre-2020) is a tool that can be used to schedule and monitor complex pipelines.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

V. GOVERNANCE MODELS

5.1 Data Governance

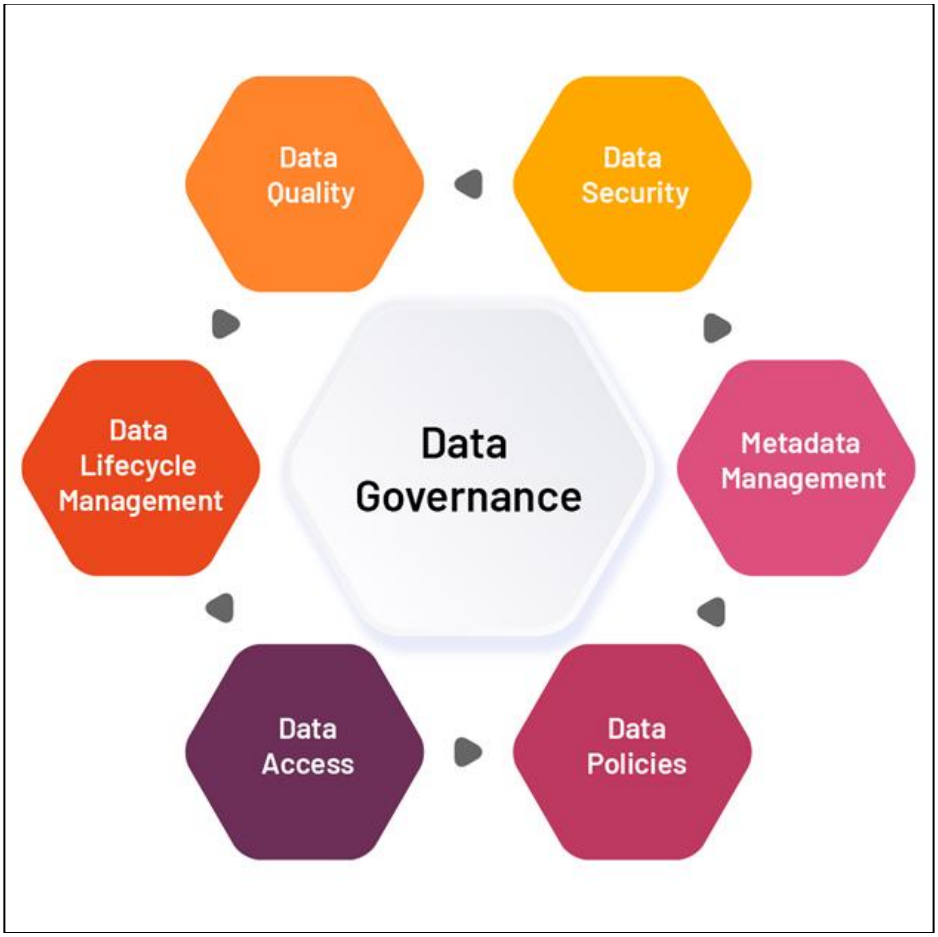

Figure 8: Data Governance in AI Models
(Source: www.factspan.com)

Data governance is key to making sure that data analyzed in AI systems is accurate, secure and regulatory compliant. It can include clear policies on ownership and stewardship of the data, responsibility for data management and oversight processes to monitor data quality. Access control and encryption are key tools in data governance's arsenal, ensuring that sensitive data is kept safe from the prying eyes and hands of the untrustworthy (Kop, 2019). The compliance with regulatory frameworks further guarantees that the organizations abide under the law and ethical norms in the utilization of data. Good data governance is the basis for trustworthy and reliable AI systems.

5.2 Model Governance

Model Governance is a way of managing the lifecycle of AI models, from their development, deployment, to continuous monitoring. Requires careful validation and testing to make sure models behave in the way they are supposed to and also don't have any unanticipated biases or mistakes. Models are altered over time and are tracked using a version control mechanism, which aids the ability to reproduce and hold responsible for any changes. As data collection changes with time, the performance may degrade or the concept may change (concept drift), so continuous monitoring is necessary to detect this behavior (Al-Shabandar et al., 2019). Effective model governance practices can help ensure that the models are reliable and trustworthy, promoting the integrity and trustworthiness of the AI systems used by the organizations.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

5.3 Ethical and Regulatory Considerations

The integration of AI systems brings important ethical and regulatory considerations, such as bias, fairness, and transparency. Research prior to 2020 has focused on the significance of explainability, elaborating that decisions made in a model should be understood, and thus be trusted by stakeholders (Henfridsson and Bygstad, 2013). To establish responsibility for AI-based decision-making, particularly in critical fields, it's crucial to have accountability frameworks in place. Risk assessment processes can be used to ensure identification of potential harms and implementation of mitigation measures. These are key considerations in ensuring that AI systems are connected with society's values and societal expectations.

5.4 Operational Governance

Operational governance is about the management of AI systems in enterprise environments on a daily basis. This includes addressing system failures or anomalies through incident management, and setting up service-level agreements demonstrating the performance expectations. By adhering to continuous integration and deployment (CI/CD) principles, businesses can keep their AI systems up to date and continuously evolving, thus staying relevant to the evolving demands. Incorporating operational governance within the larger context of the enterprise allows them to maintain the reliability, efficiency and goal-setting abilities of their AI systems.

VI. CHALLENGES AND LIMITATIONS

6.1 Data Privacy and Security

For effectively handling sensitive data in the cloud, there is need to apply strong encryption and access control systems. With these increasingly distributed systems all critical to enterprise systems, the threat landscape grows, and robust identity management, encrypted data transport mechanisms and compliance-focused data handling methodologies must be established. When it comes to AI, especially when it deals with personal or proprietary data, confidentiality and data integrity are paramount.

6.2 Model Interpretability

These generative models are usually opaque and judgement on the outputs may be difficult and interpretations of the reasoning for the generated results difficult. The opacity can cast doubt on the reliability of AI systems, particularly in professional settings where concern with accountability and explainability play a vital function (Lauterbach and Bonime 2018).The opacity can undermine trust in AI systems, especially in enterprise settings where trust, accountability, and explainability are essential. There is a need for the development of interpretability techniques and model-agnostic explanation frameworks that can become useful in understanding the behavior of the model.

6.3 System Complexity

When implementing multiple components like data pipelines, model training systems, model deployment frameworks and agent orchestration mechanisms, it adds to the overall system complexity. This complexity can bring in technical debt and complexity to build systems that are difficult to maintain over time, debug and scale. These risks can be mitigated with effective architectural design and documentation and provide for long-term sustainability.

6.4 Resource Management

AI workloads require a lot of processing power, especially when it comes to training and deploying generative models. Resource allocation will then need to be done in an efficient way considering performance and costs. Cloud can offer elasticity but if not managed correctly, there could be a lot of wasted resources or resource constraint. One of the major ways to solve this problem includes optimizing model architectures, autoscaling and workload scheduling using techniques like these.

VII. CASE APPLICATIONS (PRE-2020 CONTEXT)

Initial deployments of AI into enterprise application systems offer insights into the actual viability of enterprise cloud integration with intelligent technologies.AI's initial use in enterprise systems offers valuable lessons in the practicality of embedding intelligent technologies as a part of cloudbased enterprise architecture. In the world of e-commerce, recommendation systems mark one shining example of the machine learning's application, with their models taking into account user preferences and actions and offering customized suggestions for products (Kant, 2015). Such systems struggle to make use of scalable data processing and real time, inference functionality. Fraud detection systems use machine learning to detect irregularities in transactions in the finance industry. Accuracy, availability of real-time data,

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

and effective model governance are crucial in these systems, with the potential for costly and damaging financial and reputational consequences if predictions are faulty. Another significant use of AI technology is in Chatbots and Virtual Assistants, which use NLP to engage with users and execute customer service tasks. While the initial ones were less sophisticated than modern systems, they do highlight the basis of agent-like behavior in enterprise AI applications. The examples together illustrate the potential for enterprise AI, reflecting both how possible it can be and what the issues might be regarding enterprise implementations.

VIII. FUTURE DIRECTIONS (AS OF 2020 PERSPECTIVE)

Before 2020, there were a few road maps that were thought to be followed by the growth of AI-enabled enterprise cloud architectures. An important track has been the development of generative models to make them more accurate, scalable, and accommodating to different data modalities. These developments were anticipated to broaden the scope of generative AI applications across various sectors. The other significant way went along strengthening agentic systems' autonomy. The researchers investigated more advanced decision making processes which allowed agents to be more autonomous, yet align with organizational goals. This involved improvement of reinforcement learning and multi-agent coordination techniques. Another area identified as a potential growth opportunity for AI systems covered integration with "edge computing" environments (Renda, 2019). This shift from the cloud towards the edge would enable organizations to lower latency and enhance the responsiveness of their applications, especially in cases where instant feedback is crucial. The trend pointed to the necessity of a new paradigm of hybrid architectures and arrangements of centralized cloud resources and decentralized edge capabilities. The expected changes reinforced the ever-changing nature of AI technologies and their relevance to enterprise systems.

IX. CONCLUSION

Enabling generative AI and agentic systems into enterprise cloud architecture needs to be done in a comprehensive way involving frameworks and governance. The paper presents an organized and systematic approach using the knowledge gained from prior studies on the development of scalable, secure, and efficient AI-based systems prior to 2020. With the increasing adoption of AI technologies, comprehensive governance and architectural frameworks will become properties of enduring significance in efforts to ensure sustainable implementation.

REFERENCES

1. Al-Shabandar, R., Lightbody, G., Browne, F., Liu, J., Wang, H. and Zheng, H., 2019, October. The application of artificial intelligence in financial compliance management. In Proceedings of the 2019 international conference on artificial intelligence and advanced manufacturing (pp. 1-6). https://dl.acm.org/doi/pdf/10.1145/3358331.3358339
2. Arora, A., 2017. Evaluating Ethical Challenges in Generative AI Development and Responsible Usage Guidelines. INTERNATIONAL JOURNAL OF RESEARCH IN ELECTRONICS AND COMPUTER ENGINEERING. https://www.academia.edu/download/122909992/Title_16_AA_Oct_17_.pdf
3. Balaganski, A., 2015. API Security Management. KuppingerCole Report, 70958, pp.20-27. https://cpl.thalesgroup.com/sites/default/files/content/analyst_research/Leadership-Compass-API-Security-and-Management-25.pdf
4. Benlian, A., Kettinger, W.J., Sunyaev, A., Winkler, T.J. and Guest Editors, 2018. The transformative value of cloud computing: a decoupling, platformization, and recombination theoretical framework. Journal of management information systems, 35(3), pp.719-739. https://www.jmis-web.org/articles/1394
5. Ferreira, L., Putnik, G., Cunha, M.M.C., Putnik, Z., Castro, H., Alves, C., Shah, V. and Varela, L., 2017. A cloud-based architecture with embedded pragmatics renderer for ubiquitous and cloud manufacturing. International Journal of Computer Integrated Manufacturing, 30(4-5), pp.483-500. https://ciencipca.ipca.pt/bitstreams/34e8e463-d801-4769-a371-7c4a2065bb40/download
6. Gelev, S., 2017. Requirements for next generation business transformation and their implementation in 5G architecture. International Journal of Computer Applications. https://www.academia.edu/download/91823956/IJCA_20CRC_20_20-_20Requirements_20for_20next_20generation_20business_20transformation_20and_20their_20implementation_20in_205G_20archit.pdf
7. Henfridsson, O. and Bygstad, B., 2013. The Generative Mechanisms of Digital Infrastructure Evolution1. MIS quarterly, 37(3), pp.907-931. https://www.researchgate.net/profile/Ola-

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 4, Issue 3, May- June 2021||

DOI:10.15662/IJRPETM.2021.0503003

Henfridsson/publication/285924538_The_Generative_Mechanisms_of_Digital_Infrastructure_Evolution/links/6730e6d92326b47637d6f2ed/The-Generative-Mechanisms-of-Digital-Infrastructure-Evolution.pdf

8. Iqbal, J. and Saleh, A., 2020. ARTIFICIAL INTELLIGENCE–DRIVEN DECISION SUPPORT SYSTEMS FOR SMART ENTERPRISES. International Research Journal of Advanced Science, 1(1), pp.1-11. https://irjas.com/index.php/sciencejournal/article/download/49/43
9. Kant, K., 2015. AI and ML in Predictive Consumer Analytics: A Conceptual Model for Personalized Marketing. https://www.researchgate.net/profile/Kamal-Kant-5/publication/398814585_AI_and_ML_in_Predictive_Consumer_Analytics_A_Conceptual_Model_for_Personalized_Marketing/links/6943d0c027359023a00db933/AI-and-ML-in-Predictive-Consumer-Analytics-A-Conceptual-Model-for-Personalized-Marketing.pdf
10. Kaur, H., 2020. Building Smart Applications: A Guide to Integrating AI and Machine Learning into Salesforce. https://www.researchgate.net/profile/John-Mathew-26/publication/396262619_Building_Smart_Applications_A_Guide_to_Integrating_AI_and_Machine_Learning_into_Salesforce/links/68e4c9ac02d6215259b9c52d/Building-Smart-Applications-A-Guide-to-Integrating-AI-and-Machine-Learning-into-Salesforce.pdf
11. Kop, M., 2019. AI & intellectual property: Towards an articulated public domain. Tex. Intell. Prop. LJ, 28, p.297. https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=3409715
12. Lauterbach, A. and Bonime-Blanc, A., 2018. The artificial intelligence imperative: a practical roadmap for business. Bloomsbury Publishing USA. https://books.google.com/books?hl=en&lr=&id=_o3DEAAAQBAJ&oi=fnd&pg=PP1&dq=INTEGRATING+GENERATIVE+AI+AND+AGENTIC+SYSTEMS+INTO+ENTERPRISE+CLOUD+ARCHITECTURE:+FRAMEWORKS+AND+GOVERNANCE+MODELS&ots=2wkguTItnP&sig=RNrbpr4oCLi48IE1IHK9Z2rnaSM
13. Lauterbach, A., 2019. Artificial intelligence and policy: quo vadis?. Digital Policy, Regulation and Governance, 21(3), pp.238-263. https://www.emerald.com/dprg/article-pdf/21/3/238/580746/dprg-09-2018-0054.pdf
14. Nippatla, R.P., 2018. AI-Driven Cloud BI: Enhancing Predictive Analytics for Financial Insights. International Journal of Marketing Management. https://www.academia.edu/download/123203417/AI_Driven_Cloud_BI_Enhancing_Predictive_Analytics_for_Financial_Insights.pdf
15. Pantazis, E. and Gerber, D., 2018. A framework for generating and evaluating façade designs using a multi-agent system approach. International Journal of Architectural Computing, 16(4), pp.248-270. https://www.academia.edu/download/57885578/19_181128_IJAC_journal_01_PantazisGerber_published_version.pdf
16. Renda, A., 2019. Artificial Intelligence. Ethics, governance and policy challenges. CEPS Centre for European Policy Studies. https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=3420810
17. Thomas, P., 2020. Multi-Agent Generative AI Frameworks for Adaptive Human–Machine Collaboration. https://www.researchgate.net/profile/Phillip-Thomas-14/publication/403113198_Multi-Agent_Generative_AI_Frameworks_for_Adaptive_Human-_Machine_Collaboration/links/69c362c260c0371a60eeaa3b/Multi-Agent-Generative-AI-Frameworks-for-Adaptive-Human-Machine-Collaboration.pdf
18. Wareham, J., Fox, P.B. and Cano Giner, J.L., 2014. Technology ecosystem governance. Organization science, 25(4), pp.1195-1215. https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=2201688