



Agent-Led Tenant Experience in Data Center Services

Vinod Battapothu

Independent Researcher

vinod.battapothu.researcher@gmail.com

Abstract

A smart city is a complex ecosystem powered by a range of Information and Communication Technology (ICT) systems. Rapid urbanization is placing growing pressure on city officials to modernize services, and the adoption of smart technologies has become key to improving customer satisfaction, delivering better services, and optimizing resource utilization. Interconnected systems can significantly enhance the citizen experience — yet in practice, city services often rely on multiple, dispersed ICT systems that operate independently. This creates a "silos of innovation" problem, where new technologies improve one specific service without accounting for their cascading effects on others, producing outcomes that can feel more like opening a Pandora's box than a coordinated upgrade. Intelligent systems capable of autonomous operation offer a promising way to address this challenge. Agentic Artificial Intelligence (AI) — defined here as smart, autonomous agents capable of making decisions on behalf of users to accomplish defined goals, while coordinating with other agents to serve all users effectively — represents one such approach.

Agentic AI is also driving a new model of customer experience orchestration within enterprise data center service ecosystems, enabling holistic management of customer experience across all components, services, and service channels. Service orchestration refers to the automated, end-to-end coordination of services, where multiple cooperating services combine to deliver a unified outcome for the customer. By enabling ecosystems to continuously Learn and Adapt (L&A) as well as Satisfy and Delight (S&D) customers, Agentic AI supports digitally transformed enterprise service ecosystems that respond to evolving customer preferences and expectations — including real-time L&A and S&D across every customer service interaction.

Keywords :Agentic AI Customer Experience,Enterprise Service Orchestration,AI-Driven Data Center Operations,Autonomous CX Management,Intelligent Service Ecosystems,Enterprise Data Center Automation,AI-Powered Customer Journey Orchestration,Predictive Service Experience Optimization,Multi-Agent Enterprise Support Systems,Adaptive Infrastructure Experience Management.

1. Introduction

Digital business rebirth, AI-driven transformation, advanced data-rich capabilities and a people-centric approach are shaping the future of customer experience (CX). The importance is self-evident: a satisfied customer is not just a loyal customer, but also a good advertisement and salesperson. Enterprise data center service ecosystems are also subject to these radical shifts, aiming to enable system-


level optimization of service delivery and operations. The growing complexity of business, supported by advanced segment-specific data insights and analytics, requires these newly formed ecosystems go beyond segment-specific local optimization to support the entire business ecosystem.

Service provider organizations share large volumes of important segment data with data partners; it is essential to extract the maximum value from this segment/partner-



specific data through effect predictive and prescriptive insights/operational actions. Intelligent analytical recommendations or actions introduced to the delivery process at the appropriate moment benefit not just the service provider but the partner/supplier ecosystem as well. Achieving this type of intelligent CX delivery, where intelligent insights are woven seamlessly into the delivery process of every customer-facing function in real time, is driving the creation of advanced enterprise data-center service ecosystems.

Article Content Distribution by Section




Pie chart of approximate content distribution by major article section.

2. Conceptual Foundations

The theoretical foundations integrate perspectives on agentic AI and service orchestration from service science and management domains; operate with a concise set of concepts; make explicit the relevant sources and nature of agentic AI capabilities and service orchestration intents; and offer a heuristic that identifies Missing Maturity Areas for Agency via maturity modelling. The first topic discusses the constitutive properties of agentic AI: autonomous operational execution, situationally contextualized decision-making, goal-consistent functioning, and Sun Tzu's principle of "non-action". The second topic then elaborates on service orchestration as a class of activity underlying any service delivery in any service ecosystem.

Agentic AI implements the service-creating Gram architecture to establish and deliver service level agreements autonomously, within the confines of parameters defined by a human-in-the-loop. The Gram Eco-System Function implements service delivery protocols and creates the external service delivery containment achieving concurrent goals of fault-tolerant fulfilment of an outcome-as-a-service and de-risking the realisation of secondary to tertiary measured service targets. Agentic AI operates collaboratively across situated phases of function service delivery provisioning as an agent-based instantiation of service computing.



Bar chart showing approximate section emphasis by word count.

2.1. Agentic AI and Service Orchestration

Both autonomy and the wielder of decision-making authority are core aspects of a customer experience delivery orchestrator. Decision-making authority refers to the physical, virtual, or hybrid entity that possesses the responsibility and accountability for the outcome of the decision that influences the delivery of customer experience. In addition, the logical mapping of goals and alignment of interests is foundational to effective customer experience delivery orchestration. Customer experience is executed by weaving together interdependent services offered by a complex ecosystem of delivery partners, third-party businesses, and shared economy suppliers. Achieving compelling customer experience requires a human touch, particularly in the final fulfillment, guidance, or education moments. Therefore, agents with source autonomy that do not appropriately engage a human touch in the correct

orchestration moment can degrade customer experience. These are layered orchestration roles that shape an experience and guide the delivery interaction in real-time based on an understanding of emotions and context.

Agentic AI is an instantiation of artificial intelligence capable of autonomously and continuously performing tasks of a specific operational area, guided by business goals and monitored by humans. It includes coherent planning and execution over decentralized implementations. The adoption of Agentic AI models in enterprises can facilitate the efficient and accurate orchestration of experience delivery across services and reduce operational costs by up to 30%.


Component	Function	Enterprise Benefit
Agentic AI Engine	Autonomous decision-making and execution	Faster service response
Service Orchestrator	Coordinates multi-service workflows	Improved operational efficiency
Enterprise Data Platform	Centralized data governance and analytics	Trusted data accessibility
Human-in-the-Loop	Supervises AI decisions	Ethical and reliable operations
CX Analytics Module	Monitors customer interactions	Enhanced customer satisfaction
Multi-Agent Collaboration Layer	Enables agent cooperation	Scalable intelligent operations

Table 1: Core Components of Agentic AI Service Orchestration


3. Architectural Framework

An integrated architectural design guides operationalization of agentic AI-enabled customer experience orchestration across enterprise data center service ecosystems. Comprising a data layer with governance mechanisms and an operational layer encapsulating multiagent architectures and customer journey coordination, the framework accounts for dependencies and delineates standards, components, and interfaces.

3.1. Data Layer and Enterprise Data Platform. Data from diverse sources underpin agentic AI-driven customer experience orchestration; alignment with enterprise data strategy, systems-of-record identification, and agency-enhancing governance processes ensure data quality, accessibility, and availability while mitigating risk to operations, compliance, and reputation. Established data sources feed into an enterprise data platform—combining data storage, processing capability, quality-assurance controls, lineage tracking, integration patterns, and regulatory compliance—to facilitate access and bolster trust. Configured to help executive, operations, and risk functions achieve targeted levels of performance throughout an enterprise data center service ecosystem’s maturity roadmap, the enterprise data platform blends functions associated with data management systems, data warehouses, data lakes, and data lakes houses.



Double Layer Flow Chart – Maturity Model & Roadmap Progression



Bar diagram of Enterprise Data Platform sublayer coverage

3.1. Data Layer and Enterprise Data Platform

The Data Layer consists of multiple sources reflecting the entire enterprise's activities on the Enterprise Data Platform, delivering valuable data to the domain-specific operational technology platforms. The platform is designed considering all facets of data: provenance; governance; lineage; quality; integration; storage; and access patterns. The Data Governance sublayer defines policies and processes for identification and management of data assets throughout their lifecycle, including those related to the ethical, trusted, and responsible use of data regulated by the Data Sovereignty sublayer. Checklists indicated in the Data Sovereignty sublayer ensure compliance of enterprise-level processing with international, national, and local laws, the resources assigned to monitor potential breaches, and the ones authorized to assure compliance.

The Data Quality sublayer specifies criteria and responsibilities for validation of data availability, consistency, accuracy, timeliness, and uniqueness. The Data Integration sublayer defines processes for ingestion of data in accordance with enterprise policies, best-fit classification into one of the three data patterns, and processing by ETL or ELT. The Data Storage sublayer specifies storage capacities and

technologies for operational and asset historical data along with the respective data retention periods. Finally, the Data Access sublayer provides an inventory of enterprise-level services, logical or physically replicated, which allow operation and processing modules to exploit the information in a fast and efficient way.

Capability	Description	Operational Impact
Predictive Analytics	Forecasts incidents and service needs	Reduced downtime
Autonomous Automation	Self-executing workflows	Lower operational costs
Adaptive Learning	Continuous learning from interactions	Personalized CX
Real-Time Monitoring	Tracks infrastructure and service metrics	Faster issue resolution
Intelligent Routing	Dynamically assigns service requests	Improved SLA compliance
Decision Intelligence	AI-assisted operational decisions	Better governance

Table 2: Agentic AI Capabilities in Enterprise Data Centers

4. Methodologies for Deployment

Deployment of the Architectural Framework requires systematic methods for operationalizing the design in practice. Such methods should be backed by robust, repeatable processes. Maturity models and roadmapping are two commonly used approaches that guide organizations through the roll-out of complex capabilities by mapping levels of capability maturity, providing criteria for evaluating the current state, defining key performance indicator targets,



establishing timelines for achieving the roadmap, and defining overall governance for the progression.

Maturity models help organizations break down a complex capability area, such as the deployment of agentic AI in service orchestration, into several easier-to-manage segments. Each segment is supported by criteria used to assess maturity against that capability segment. To facilitate this, the escalation of maturity is typically mapped across five discrete levels: initial; repeatable; defined; managed; and optimized. The resulting maturity model for agentic AI required for service orchestration includes 14 segments, consisting of people, process, governance, and technology, across which assessment criteria and levels of maturity have been developed. Roadmaps then define the required changes, together with the key activities, timing, costs, and responsibilities associated with achieving the desired KPI targets associated with the segmented capability.

Collaboration Type	Human Involvement	AI Decision Authority	Use Case
Consulted	High	Medium	Service recommendations
Informed	Medium	High	Automated notifications
In Charge	High	Low	Critical incident management
Not Involved	Low	Full AI autonomy	Routine operational automation

Table 3: Human–AI Collaboration Models

Mathematical Formulas:

1. Customer Experience Satisfaction Function

Measures overall customer satisfaction across enterprise services.

$$CX_{score} = \sum_{i=1}^n w_i \times S_i$$

Where:

- CX_{score} = Overall customer experience score
- w_i = Weight of service factor
- S_i = Satisfaction level of factor i

2. Agentic AI Decision Optimization

Represents autonomous decision efficiency.

$$D_{opt} = \frac{A_c}{T_r}$$

Where:

- D_{opt} = Decision optimization score
- A_c = Accurate autonomous actions
- T_r = Total requests processed

3. Service Orchestration Efficiency

Evaluates orchestration performance across services.

$$SOE = \frac{S_d}{S_t}$$

Where:



- SOE = Service orchestration efficiency
- S_d = Successfully delivered services
- S_t = Total requested services

4. AI-Driven Predictive Accuracy

Used for predictive customer experience analytics.

$$PA = \frac{TP + TN}{TP + TN + FP + FN}$$

Where:

- TP = True positives
- TN = True negatives
- FP = False positives
- FN = False negatives

5. Enterprise Data Quality Index

Measures governance and quality of enterprise data.

$$DQI = \frac{C + A + T + U}{4}$$

Where:

- C = Consistency
- A = Accuracy
- T = Timeliness
- U = Uniqueness

6. Human-AI Collaboration Index

Represents collaboration effectiveness between humans and AI agents.

$$HAI = \alpha H + \beta AI$$

Where:

- H = Human contribution score
- AI = AI agent contribution score
- α, β = Collaboration weights

7. System Reliability Equation

Used for enterprise data center operational reliability.

$$R(t) = e^{-\lambda t}$$

Where:

- $R(t)$ = Reliability at time t
- λ = Failure rate
- t = Operational time

8. Service Latency Optimization

Measures response latency improvement.

$$L_o = \frac{L_b - L_a}{L_b} \times 100$$

Where:

- L_o = Latency optimization percentage



- L_b = Latency before AI orchestration
- L_a = Latency after AI orchestration

$$RU = \frac{R_u}{R_a} \times 100$$

9. Risk Mitigation Score

Represents governance and compliance effectiveness.

$$RM = \frac{R_p}{R_t}$$

Where:

- RM = Risk mitigation ratio
- R_p = Prevented risks
- R_t = Total identified risks

10. Multi-Agent Coordination Function

Models interaction among autonomous AI agents.

$$MAC = \sum_{i=1}^n \sum_{j=1}^m A_i \cdot C_{ij}$$

Where:

- A_i = Agent activity
- C_{ij} = Coordination coefficient between agents

11. Resource Utilization Optimization

Represents efficient utilization in data center ecosystems.

Where:

- RU = Resource utilization percentage
- R_u = Used resources
- R_a = Available resources

12. AI Trust and Explainability Score

Models customer trust in AI-driven orchestration.

$$TS = \frac{E + T + C}{3}$$

Where:

- E = Explainability score
- T = Transparency score
- C = Compliance score

13. Workflow Automation Gain

Measures operational improvement through automation.

$$WAG = \frac{P_a - P_m}{P_m} \times 100$$

Where:

- WAG = Workflow automation gain
- P_a = Automated process performance

- P_m = Manual process performance

14. Adaptive Learning Function for Agentic AI

Represents continuous AI learning capability.

$$L(t+1) = L(t) + \eta \times \Delta E$$

Where:

- $L(t)$ = Current learning state
- η = Learning rate
- ΔE = Error correction

15. Customer Retention Prediction

Used for AI-enabled customer retention analytics.

$$CR = \frac{C_r}{C_t} \times 100$$

Where:


- CR = Customer retention rate
- C_r = Retained customers
- C_t = Total customers

4.1. Maturity Models and Road mapping

Infrastructure-as-code, orchestration, artificial intelligence, and digital twins significantly simplify and quicken enterprise data platform development, but a pragmatic, informatized, disciplined approach is vital. Maturity roadmaps allow enterprises to define their desired enterprise data platform maturity level and operational performance trajectory,

anticipate the necessary evolution of enabling capabilities in business functions, processes, organization, data sources, and technology so that all parts of the enterprise progress on plan, and deliver value quickly while managing risk.

So far, no business function has operated without applying Stages 1 and 2 of the Digital Operating Model. High priority is given to the development of production automation capabilities and the use of these capabilities by enterprise-wide process execution, but the time required is substantial. System downtime can continue to be significantly further reduced, but not eliminated, and customer satisfaction does not yet meet target. The Enterprise Data Platform has provided the Enterprise Data Ecosystem with fundamental data function capabilities, such as volume, governance, and data quality recommended by the Maturity Model.



Graph of the five-level maturity progression: initial to optimized.




Double Layer Flow Chart – Human-AI Collaborative Service Ecosystem

5. Operational Impacts and Case Considerations

Enhancing operations and customer experience through agentic AI presents inherent challenges. Effective development, integration, and continuous innovation require engagement, support, and governance from a broad spectrum of business and IT stakeholders. Consequently, the operational impact of an agentic foundation on service delivery encompasses various facets. The AI model must combine multiple sources of expertise and skill, both historical and current, to execute tasks that involve diverse expertise. Determining how the operational model changes—and what information requirements are necessitated beyond those of existing operations—is vital for management and production support personnel.

A common enterprise data platform provides the underlying foundation, unifying all data, services, and applications. Maturity models define the operational roadmap and assist in achieving business objectives. Consequently, agentic-enabled services can be broadly categorized as: (a) enabling the enterprise data platform and assisting data governance, management, and engineering operations; (b) enhancing service delivery and customer experience through human-AI collaboration; (c) supporting software engineering processes; (d) enhancing security operations; and (e) advancing service operation management processes. The associated changes in operation, performance, risk, and value realization remain consistent across these categories, but the underlying enabling requirements differ significantly, as described subsequently.

Agentic-Enabled Service Categories



Pie chart of agentic-enabled service categories; equal conceptual weighting because no numeric weights are stated.

5.1. Human-AI Collaboration in Service Delivery

The operationalization of agentic AI capabilities alters the customer experience delivery processes of an enterprise data center service ecosystem. Risks must be mitigated to ensure that the envisaged operational enhancements are realized. In the customer experience delivery process of an enterprise data center service ecosystem, factors such as CxOs' trust in AI agents' competence, roles in service delivery collaboration and decision-making authority are considered vital for successful adoption. The delivery process might require decision-making changes grounded in CxOs' confidence in AI to deliver better results than humans. CxOs' concern about job loss as a result of introducing agentic AI could also be addressed using a change management approach.

Collaboration between CxOs and AI agents for service delivery can be classified into four categories based on CxOs' task involvement with the help of two perspectives: decision authority and decision involvement (the degree to which an individual takes part in the process leading to the choice). The four categories—consulted, informed, in charge, and not involved—help define the preconditions of successful collaboration with agents in each role. The nature of human–AI collaboration is essential in determining success or failure; decisions made by humans seldom remain unchallenged when the outcome is less than expected. Therefore, the decision authority for these role types must be clearly established. Training for the roles defined in the delivery process is essential. Change management is necessary to create an environment that inspires CxOs to collaborate closely with agents and enables them to accept the changed decision authority gracefully.

Governance Layer	Primary Responsibility	Compliance Focus
Data Sovereignty	Regional data control	GDPR, CCPA, PDPA
Data Quality	Validation and consistency	Accurate analytics
Data Integration	ETL/ELT orchestration	Unified enterprise data
Data Storage	Historical and operational storage	Retention policies
Data Access	Secure information access	Authorization control
Audit & Assurance	Monitoring AI decisions	Regulatory transparency


Table 4: Data Governance and Compliance Layers

6. Ethical, Legal, and Compliance Implications

Governance and control of operational models are primary components of successful AI deployment and will determine how responsibility and risk are managed across a diverse cloud ecosystem. Enabling a higher level of autonomy creates new trust models and changes the decision-making landscape in service delivery. While accountability remains a cornerstone, it now comprises an ecosystem approach where aligned interests support the delivery of value across multiple business decisions, and not just within one organization. Trust remains a key consideration and is typically achieved by enabling an appropriate level of explainability of AI-driven decisions. The application of AI at scale introduces risks associated with the behavior of complex systems that are constantly changing. Appropriate levels of monitoring, evaluation, and risk modeling are essential to ensure and

minimize the impact of these unknowns. An enterprise operating with the utmost transparency while at the same time protecting confidential information provides the highest level of trust in its AI models.

A customer-centric deployment of agentic AI requires respect for privacy, data sovereignty, and compliance. The notification and consent of individuals involved in complex data- and AI-driven operations are essential elements for the alignment of AI services. Current and future laws governing the handling of data, from its acquisition through its use and consumption, must be taken into consideration. To mitigate the risks arising from the mounting regulatory burden, organizations need to conduct a thorough assessment of the laws impacting their AI services and the specific requirements of the law for each role. Solid data lifecycle controls reduce the risks of noncompliance or misuse of data. Data should be classified according to its attributes, location and geographical requirements, anonymized or pseudonymized to the extent possible, and utilized for its primary purpose. All AI services should incorporate auditing and assurance services that may be leveraged for regulatory compliance.



Bar diagram of operational impact categories covered in the article.

6.1. Privacy, Data Sovereignty, and Compliance

Governance, accountability, compliance, and risk management are essential in sustainable operations. Insights from related domains help balance AI deployment benefits



with privacy needs and legal requirements. Given the complexity of deploying AI techniques in data-centric organizations, governance at all levels is crucial to ensuring transparency, consistency, and control of the implemented AI approach. The following major categories play a pivotal role in implementing transparency.

Enterprise data platform and service management leaders must define procedures that address privacy and data protection concerns for customers and citizens—such as data sovereignty and handling personal identifiable information (PII)—(in-transit) consent, data locality, and processing. A privacy policy that covers these aspects, is shared with users of the enterprise data platform and services, and is consistently implemented will limit the legal exposure during an internal or external audit. Privacy policies should also incorporate sufficiently detailed descriptions of training and test data needed for the deployment and the implication of such data when establishing controls for AI-generated/assisted data.

With governments worldwide yielding to the pressure of privacy advocates over the past decade, companies have started putting their infrastructure and process definitions in place to comply with laws such as the California Consumer Privacy Act (CCPA), the General Data Protection Regulation (GDPR), and Singapore's Personal Data Protection Act (PDPA). Privacy-related requirements at federal and state levels must be compiled based on the company's geographical footprint, and action plans must be formulated to meet the minimum at the right time. AI techniques that require customers' PII data for training or testing demand a more stringent approach, since the laws cover the handling of sensitive personal data as well as the consent requirement before sharing or processing such data.

Maturity Level	Characteristics	Organizational Readiness
Initial	Manual operations	Low

Maturity Level	Characteristics	Organizational Readiness
Repeatable	Basic automation introduced	Moderate
Defined	Standardized AI workflows	Medium
Managed	KPI-driven optimization	High
Optimized	Autonomous intelligent ecosystem	Very High

Table 5: Enterprise AI Maturity Model

7. Conclusion

Operationalization of the presented architectural framework creates a robust system that orchestrates the delivery of services provided by an enterprise data center service ecosystem and introduces a new layer of service delivery players. Methods for implementation are as rigorous and repeatable as an artistic endeavor can be, and any enterprise data center service ecosystem pursuing operational excellence, innovation, investment reduction, and resilience should start with them. The methods emphasize the value of progression through maturity levels, supported by a roadmap of actions necessary for moving toward the desired future state while keeping the focus on business rather than IT performance. Execution should be improved through external supplier-assisted delivery models that use highly specialized vendors for selected service delivery components without compromising end-to-end ...

The orchestration requirement across the service ecosystem indicates that human–AI collaboration is paramount for operational effectiveness and efficiency. As the use of advanced AI technologies becomes ubiquitous, a closer collaboration between robots and humans also becomes



paramount for enterprise data center service ecosystem operations. A matrix of potential collaboration models recognizes the actors involved, their focus area, decision rights regarding the specific service or activity domain, training and upskilling directions, and change management support required for successful adoption. Many ensure optimum decision-making at every level by defining where human judgment is required most and when machines can act effectively without human emotional quotient and/or IQ.

References

1. Baier, L., Kühl, N., Schüritz, R., & Satzger, G. (2021). Will the customers be happy? Identifying unsatisfied customers from service encounter data. *Journal of Service Management*, 32(2), 265–288.
2. Ashfaq, M., Yun, J., Yu, S., & Loureiro, S. M. C. (2020). I, chatbot: Modeling the determinants of users' satisfaction and continuance intention of AI-powered service agents. *Telematics and Informatics*, 54, 101473.
3. Sheehan, B., Jin, H. S., & Gottlieb, U. (2020). Customer service chatbots: Anthropomorphism and adoption. *Journal of Business Research*, 115, 14–24.
4. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*, 1(1), 1–13.
5. Li, L., Lee, K. Y., Emokpae, E., & Yang, S.-B. (2021). What makes you continuously use chatbot services? Evidence from Chinese online travel agencies. *Electronic Markets*, 31(3), 575–599.
6. Følstad, A., & Taylor, C. (2021). Investigating the user experience of customer service chatbot interaction: A framework for qualitative analysis of chatbot dialogues. *Quality and User Experience*, 6, Article 6.
7. Hsiao, K.-L., & Chen, C.-C. (2022). What drives continuance intention to use a food-ordering chatbot? An examination of trust and satisfaction. *Library Hi Tech*, 40(4), 929–946.
8. Chen, Q., Gong, Y., Lu, Y., & Tang, J. (2022). Classifying and measuring the service quality of AI chatbot in frontline service. *Journal of Business Research*, 145, 552–568.
9. Kharub, I., Lwin, M., Khan, A., & Mubin, O. (2021). Perceived service quality in HRI: Applying the SERVBOT framework. *Frontiers in Robotics and AI*, 8, 746674.
10. Wang, X., Lu, J., Feng, Y., & Liu, L. (2021). Antecedents and mediating role of IT service climate in IT service quality: A mixed methods study. *International Journal of Information Management*, 57, 102290.
11. Al-Hawari, F., & Barham, H. (2021). A machine learning based help desk system for IT service management. *Journal of King Saud University—Computer and Information Sciences*, 33(6), 702–718.
12. Mangalampalli, B. M. (2022). Automated Invoice Validation Systems Using Advanced SQL Analytics in Healthcare Insurance. *Front Health Inform*, 11.
13. Notaro, P., Cardoso, J., & Gerndt, M. (2021). A survey of AIOps methods for failure management. *ACM Transactions on Intelligent Systems and Technology*, 12(6), Article 81.
14. Dong, W. (2022). AIOps architecture in data center site infrastructure monitoring. *Computational Intelligence and Neuroscience*, 2022, 1988990.
15. Li, X.-Y., Liu, Y., Lin, Y.-H., Xiao, L.-H., Zio, E., & Kang, R. (2021). A generalized Petri net-based modeling framework for service reliability evaluation and management of cloud data centers. *Reliability Engineering & System Safety*, 207, 107381.
16. Moghaddam, S. M. M., O'Sullivan, M., Walker, C. G., Fotuhi Piraghaj, S., & others. (2020). Embedding individualized machine learning prediction models for energy efficient VM consolidation within cloud data centers. *Future Generation Computer Systems*, 106, 221–233.
17. Mangala, N. (2021). Optimizing Large-Scale ETL Pipelines Using Medallion Architecture on Azure Data Lake. *Journal of Artificial Intelligence and Big Data*, 1(1), 1–20.
18. Liang, B., Dong, X., Wang, Y., & Zhang, X. (2020). Memory-aware resource management algorithm for low-



- energy cloud data centers. *Future Generation Computer Systems*, 113, 329–342.
19. Mirhoseininejad, S. M., Badawy, G., & Down, D. G. (2021). Holistic thermal-aware workload management and infrastructure control for heterogeneous data centers using machine learning. *Future Generation Computer Systems*, 118, 208–218.
 20. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. *World Journal of Clinical Medicine Research*, 1(1), 1-18.
 21. Wang, B., Liu, F., & Lin, W. (2021). Energy-efficient VM scheduling based on deep reinforcement learning. *Future Generation Computer Systems*, 125, 616–628.
 22. Xiao, X., Sun, J., & Yang, J. (2021). Operation and maintenance (O&M) for data center: An intelligent anomaly detection approach. *Computer Communications*, 178, 141–152.
 23. Shoukourian, H., & Kranzlmüller, D. (2020). Forecasting power-efficiency related key performance indicators for modern data centers using LSTMs. *Future Generation Computer Systems*, 112, 362–382.
 24. Saxena, D., & Singh, A. K. (2021). A proactive autoscaling and energy-efficient VM allocation framework using online multi-resource neural network for cloud data center. *Neurocomputing*, 426, 248–264.
 25. Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. *International Journal of Scientific Research and Modern Technology*, 1(12), 227-237.
 26. Caviglione, L., Gaggero, M., Paolucci, M., & Ronco, R. (2021). Deep reinforcement learning for multi-objective placement of virtual machines in cloud datacenters. *Soft Computing*, 25, 12569–12588.
 27. Asghari, A., Sohrabi, M. K., & Yaghmaee, F. (2020). A cloud resource management framework for multiple online scientific workflows using cooperative reinforcement learning agents. *Computer Networks*, 179, 107340.
 28. Hu, X., & Sun, Y. (2020). A deep reinforcement learning-based power resource management for fuel cell powered data centers. *Electronics*, 9(12), 2054.
 29. Viola, L., Ronchieri, E., & Cavallaro, C. (2022). Combining log files and monitoring data to detect anomaly patterns in a data center. *Computers*, 11(8), 117.
 30. Li, J., Zhang, R., & Zheng, Y. (2022). QoS-aware and multi-objective virtual machine dynamic scheduling for big data centers in clouds. *Soft Computing*, 26, 10239–10252.
 31. Mangala, N. (2021). CI/CD Pipeline Automation for Enterprise Data Artifacts Using Azure DevOps. *Universal Journal of Business and Management*, 1(1), 1-18.