International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

Self-Supervised AI for Intelligent Enterprise Knowledge Discovery in Multi-Tenant Cloud Environments

Dr. Ravie Chandren Muniyandi

Associate Professor, Department of Computing (CCI), College of Computing and Informatics, Universiti Tenaga Nasional, Malaysia

Publication History: Received: 10.06.2026; Revised: 26.07.2026; Accepted: 01.08.2026; Published: 13.08.2026

ABSTRACT: Multi-tenant cloud environments enable multiple organizations, departments, and users to share computing infrastructure while maintaining logical separation of applications and data. Although this architecture improves resource utilization, scalability, and cost efficiency, the enormous volume and diversity of enterprise data generated across tenants makes knowledge discovery increasingly difficult. Conventional knowledge-discovery approaches often depend on manually labeled datasets, predefined rules, or isolated analysis pipelines, limiting their ability to identify hidden relationships and emerging patterns. Self-Supervised Artificial Intelligence (SSAI) provides a promising solution by learning meaningful representations from large quantities of unlabeled enterprise data through automatically generated learning objectives. This paper investigates the application of self-supervised AI for intelligent enterprise knowledge discovery in multi-tenant cloud environments. The proposed approach combines representation learning, cloud data integration, tenant-aware processing, semantic analysis, anomaly detection, clustering, and knowledge-graph construction. Enterprise artifacts such as documents, logs, application metadata, service interactions, operational records, and usage patterns are transformed into contextual representations without requiring extensive manual labeling. The methodology evaluates the proposed framework according to knowledge-discovery accuracy, representation quality, anomaly-detection capability, clustering effectiveness, scalability, processing efficiency, and tenant-isolation requirements. The study also examines important challenges involving data privacy, security, noisy data, computational cost, model bias, explainability, and cross-tenant information leakage. The research demonstrates that self-supervised AI can provide an adaptive foundation for discovering previously hidden enterprise knowledge while preserving the operational and security requirements of multi-tenant cloud computing.

KEYWORDS: Self-supervised AI, enterprise knowledge discovery, multi-tenant cloud, artificial intelligence, representation learning, machine learning, knowledge management, cloud computing, anomaly detection, semantic analysis, knowledge graphs, data mining, enterprise analytics, tenant isolation, intelligent systems

I. INTRODUCTION

Multi-tenant cloud computing has become a fundamental model for delivering enterprise software and infrastructure at scale. In a multi-tenant architecture, multiple customers or organizational units share underlying computing resources while maintaining logical isolation between their applications, workloads, and data. This approach provides important benefits, including efficient resource utilization, elastic scalability, simplified infrastructure management, and reduced operational costs. However, the shared nature of the environment produces enormous quantities of heterogeneous information. Enterprise applications generate application logs, transaction records, documents, configuration information, API interactions, user activity, performance metrics, security events, and service dependencies. These artifacts may exist in structured, semi-structured, and unstructured formats across different cloud services. Valuable enterprise knowledge can therefore be distributed across multiple repositories and hidden within relationships that are difficult to identify through conventional analysis. For example, a sequence of seemingly independent operational events may indicate an emerging service failure, while relationships among documents, customers, applications, and business processes may reveal previously unidentified organizational knowledge. The challenge is not simply storing and retrieving information but transforming continuously generated data into meaningful knowledge that can support enterprise decision-making. Intelligent knowledge discovery must therefore be capable of processing heterogeneous information, identifying latent relationships, adapting to changing environments, and operating within strict tenant-isolation and privacy requirements.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

Traditional artificial intelligence and machine-learning approaches to enterprise knowledge discovery frequently depend on supervised learning. Supervised models require labeled examples to learn tasks such as classification, anomaly detection, sentiment analysis, document categorization, or prediction. Although supervised learning can achieve strong performance when high-quality labels are available, creating and maintaining labels for large enterprise datasets is expensive and time-consuming. The problem becomes more significant in multi-tenant cloud environments because each tenant may use different applications, business processes, terminology, and operational patterns. A model trained using labels from one environment may therefore perform poorly when applied to another. Self-supervised AI offers an alternative by creating learning signals directly from the available data. Instead of requiring humans to label every example, a self-supervised model can learn by predicting masked information, reconstructing missing elements, identifying relationships between data elements, or determining whether different representations originate from the same underlying context. This enables the model to learn general-purpose representations from large collections of unlabeled enterprise information. These representations can subsequently support downstream activities such as semantic search, document clustering, anomaly detection, recommendation, classification, summarization, and knowledge-graph construction. The approach is especially valuable for enterprises because substantial amounts of operational and informational data already exist, but only a small proportion may have formal labels.

Knowledge discovery in a multi-tenant environment presents challenges that extend beyond conventional enterprise analytics. Tenant data must be isolated so that information belonging to one organization cannot be unintentionally exposed to another. At the same time, enterprises may wish to identify general patterns that occur across tenants, such as common service failures, security threats, performance bottlenecks, or effective operational practices. This creates a tension between learning from shared patterns and preserving tenant confidentiality. Self-supervised AI can potentially address this challenge through tenant-aware representation learning in which models learn common structural characteristics without directly exposing tenant-specific information. Techniques such as privacy-preserving representation learning, federated learning, access-controlled data processing, tenant-specific embeddings, and differential privacy can further strengthen isolation. Another important issue is semantic heterogeneity. Different tenants may use different terminology to describe similar business concepts, and the same term may have different meanings across organizations. Representation learning can help map related concepts into meaningful semantic spaces, enabling the discovery of relationships that keyword-based approaches might miss. For example, two tenants may describe similar operational problems using different terminology. A self-supervised model can learn contextual relationships and potentially identify their conceptual similarity without requiring manually constructed dictionaries for every tenant.

The objective of this research is to investigate how self-supervised AI can support intelligent enterprise knowledge discovery in multi-tenant cloud environments. The proposed framework combines heterogeneous data ingestion, preprocessing, self-supervised representation learning, tenant-aware knowledge extraction, anomaly identification, clustering, semantic retrieval, and knowledge-graph construction. Rather than treating enterprise data sources as independent repositories, the framework seeks to create contextual representations that capture relationships among documents, applications, users, services, events, and business entities. The research methodology evaluates whether these representations can improve downstream knowledge-discovery tasks while maintaining scalability and tenant isolation. The study also considers practical challenges associated with deploying self-supervised models in enterprise cloud environments, including computational requirements, noisy data, model drift, explainability, privacy, security, and governance. The central premise is that enterprises can obtain greater value from existing data by allowing AI systems to learn its underlying structure without requiring exhaustive manual annotation. Such systems could support intelligent search, operational decision-making, security monitoring, enterprise knowledge management, and predictive analytics. However, the research emphasizes that self-supervised AI should complement rather than replace human expertise. Human validation remains important when discovered relationships influence critical business decisions. The proposed work therefore provides a framework for combining automated representation learning with enterprise governance, enabling more scalable and adaptive knowledge discovery while preserving the security and confidentiality requirements of multi-tenant cloud computing.

II. LITERATURE REVIEW

Enterprise knowledge discovery has traditionally been associated with data mining, information retrieval, business intelligence, knowledge management, and machine learning. Conventional approaches commonly depend on structured databases and predefined analytical models to identify patterns and support organizational decision-making. Data warehouses and data lakes have expanded the ability of enterprises to collect and retain information from diverse

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

systems. However, the increasing use of cloud-native applications has resulted in data being distributed across microservices, SaaS platforms, cloud storage, observability systems, collaboration platforms, and operational databases. Enterprise knowledge is consequently fragmented across multiple technical and organizational boundaries. Information retrieval techniques can locate relevant documents using keywords and metadata, but they may fail to identify semantically related information expressed using different terminology. Data-mining techniques can identify statistical relationships but may require carefully structured features and domain-specific preparation. Knowledge graphs have emerged as another approach by explicitly representing entities and relationships, allowing organizations to perform semantic queries and relationship analysis. However, constructing high-quality knowledge graphs traditionally requires significant data engineering and ontology-development effort. Machine learning and natural-language processing have increasingly been used to automate entity extraction, relationship identification, document classification, and semantic similarity, providing an important foundation for intelligent knowledge discovery.

Self-supervised learning has emerged as a major research direction in artificial intelligence because it enables models to learn from unlabeled data by generating supervision signals from the data itself. In language processing, approaches based on masked-token prediction and related pretext tasks have demonstrated that models can learn rich semantic representations from large text collections. Similar principles have been applied to images, audio, video, graphs, and multimodal data. The fundamental advantage of self-supervised learning is its ability to reduce dependence on manually labeled datasets. After pretraining, learned representations can be adapted to downstream tasks using smaller amounts of labeled data. This property is particularly relevant to enterprise environments, where large volumes of data are available but domain-specific labels are often limited. Self-supervised objectives can be designed around enterprise characteristics, such as predicting missing log fields, reconstructing application events, matching related documents, identifying temporal relationships, or learning associations between service interactions. The quality of the resulting representation depends heavily on the design of these objectives and the diversity of the training data. Poorly designed objectives may encourage the model to learn irrelevant correlations rather than meaningful enterprise knowledge. Consequently, research increasingly emphasizes domain-aware pretraining and careful evaluation of learned representations.

Cloud computing research has increasingly investigated intelligent analytics for resource management, anomaly detection, security monitoring, and workload prediction. In multi-tenant environments, machine-learning models have been used to identify unusual behavior, forecast resource requirements, optimize workload placement, and detect potential security threats. However, conventional machine-learning systems often require labeled examples of normal and abnormal behavior. Obtaining representative labels for every tenant is difficult because operational patterns vary significantly across organizations. Self-supervised approaches can learn normal representations directly from large volumes of historical activity and then identify deviations from those representations. For instance, a model could learn the normal temporal relationships among application events, service calls, and resource utilization metrics. Significant deviations could then be investigated as possible anomalies. Similarly, tenant-aware embeddings could capture each organization's unique operational characteristics while allowing the broader model to learn common patterns. Research in federated and privacy-preserving learning is also relevant because cloud tenants may not be willing to share raw data. Federated learning can allow models to learn from decentralized datasets without necessarily centralizing all training data, although it introduces challenges involving communication overhead, heterogeneous data distributions, malicious participants, and model privacy.

Knowledge graphs and representation learning provide another important connection to the proposed research. Knowledge graphs represent entities and relationships explicitly, making them useful for enterprise search, recommendation, reasoning, and decision support. Recent research has explored graph embeddings and graph neural networks for learning representations from these structures. Combining self-supervised learning with knowledge graphs can allow models to learn from both textual information and explicit relationships. In a multi-tenant environment, this can support the construction of tenant-specific or federated knowledge representations. For example, applications, services, documents, business entities, and operational events can be connected in a graph, while self-supervised objectives can encourage the model to learn meaningful relationships among them. However, the literature also identifies challenges related to privacy, explainability, bias, data quality, and model governance. A model may unintentionally memorize sensitive information, reproduce biases present in enterprise data, or infer relationships that are statistically plausible but operationally incorrect. These issues are particularly important when AI-generated knowledge is used for business-critical decisions. Existing research therefore supports the use of access controls, privacy-preserving learning, human validation, provenance tracking, and explainable AI. The research proposed here

International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

extends these ideas by integrating self-supervised representation learning with tenant-aware knowledge discovery, providing a unified framework for extracting enterprise knowledge while respecting multi-tenant cloud constraints.

III. RESEARCH METHODOLOGY

The proposed research follows a design-oriented experimental methodology consisting of four major stages: data preparation and tenant-aware representation, self-supervised model development, knowledge-discovery integration, and experimental evaluation. The first stage establishes a simulated or real-world multi-tenant cloud environment containing multiple organizational tenants with heterogeneous applications and workloads. Data sources can include application logs, service metadata, documents, API records, configuration information, monitoring metrics, security events, transaction records, and user-activity information. The collected data is classified according to structure, sensitivity, tenant ownership, and business relevance. Personally identifiable information and confidential enterprise attributes are anonymized or protected before model processing. Tenant identifiers and access-control metadata are retained in a controlled form so that the system can maintain isolation during data processing. Data preprocessing includes normalization, duplicate removal, timestamp alignment, tokenization for textual data, event encoding, missing-value handling, and semantic enrichment. Related information from different repositories is linked using controlled identifiers and metadata. This stage produces a tenant-aware enterprise data representation that can support both independent tenant analysis and carefully controlled cross-tenant learning. The methodology emphasizes that no raw tenant information should become accessible to unauthorized participants simply because it is used for AI training.

The second stage develops the self-supervised representation-learning component. Instead of depending primarily on manually labeled datasets, the model learns through automatically generated training objectives. For textual enterprise data, possible objectives include masked-content prediction, sentence or document relationship prediction, semantic similarity learning, and reconstruction of incomplete information. For operational data, objectives can include predicting subsequent events, reconstructing missing log attributes, identifying temporal relationships, and learning normal sequences of application or service interactions. For multimodal enterprise information, the model can learn correspondence between documents, metadata, events, and system entities. Multiple self-supervised objectives can be combined to encourage representations that capture both semantic and operational relationships. The model is trained using historical enterprise data, with careful separation between training, validation, and evaluation periods to prevent temporal leakage. Tenant-aware learning strategies are incorporated to ensure that common representations do not inadvertently expose sensitive tenant-specific information. Where feasible, privacy-preserving techniques such as federated training, differential privacy, restricted parameter sharing, or tenant-specific adaptation can be evaluated. The resulting model produces embeddings that represent enterprise artifacts in a shared semantic space while retaining sufficient contextual information for downstream discovery tasks.

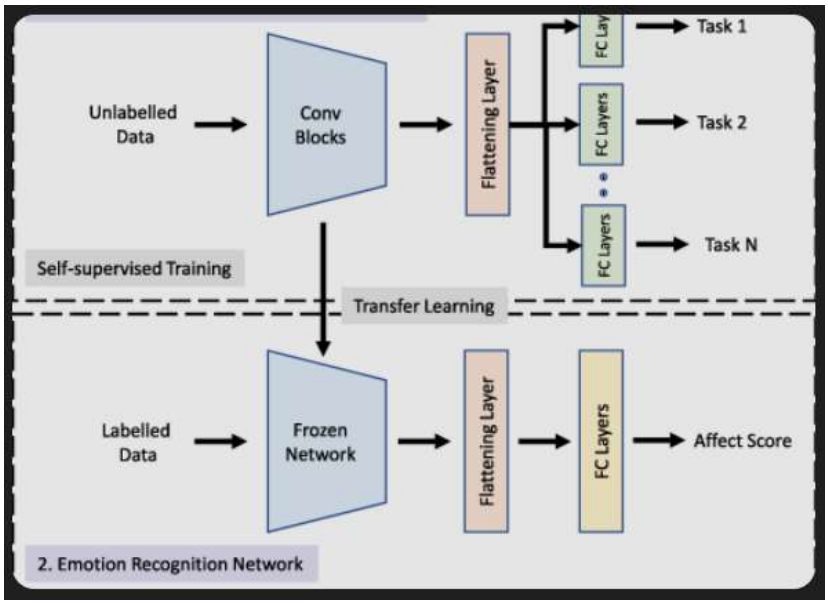

FIG1: Self-Supervised AI for Intelligent Enterprise Knowledge Discovery

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

The third stage integrates the learned representations with enterprise knowledge-discovery mechanisms. Embeddings are used for semantic search, clustering, anomaly detection, recommendation, and knowledge-graph construction. Semantic search allows users to retrieve conceptually related information even when exact keywords do not match. Clustering can identify groups of documents, events, applications, or operational patterns that share similar characteristics. Anomaly detection compares new observations with learned representations of normal historical behavior and identifies significant deviations. Knowledge-graph construction extracts entities and relationships from enterprise artifacts and connects them using learned semantic representations. The resulting graph can represent relationships among applications, services, business processes, documents, users, events, and organizational concepts. A feedback mechanism allows domain experts to validate discovered relationships and correct inaccurate associations. Validated feedback can subsequently be used for model adaptation or downstream supervised refinement. The system can also provide provenance information showing the data sources from which a discovered relationship originated. This is important for enterprise environments because users need to distinguish between directly observed facts and AI-generated inferences. Tenant-specific access-control policies are applied at the retrieval and presentation layers so that users receive only knowledge authorized for their organizational context.

The fourth stage evaluates the proposed framework through quantitative and qualitative experiments. Evaluation metrics depend on the downstream knowledge-discovery task. For semantic retrieval, precision, recall, ranking quality, and mean reciprocal rank can be measured. For clustering, metrics such as silhouette score and cluster consistency can be evaluated alongside expert judgment. For anomaly detection, precision, recall, F1-score, false-positive rate, and detection latency can be measured using controlled abnormal events or historical incidents. Knowledge-graph quality can be assessed using entity-extraction accuracy, relationship accuracy, graph completeness, and expert validation. Representation quality can be evaluated by comparing the self-supervised model against supervised learning, unsupervised baseline methods, and conventional enterprise search techniques. Scalability experiments can examine the effect of increasing numbers of tenants, documents, events, and model parameters. Privacy experiments can evaluate whether tenant-specific information can be inferred from shared representations under different protection mechanisms. Human evaluation can involve enterprise users assessing the relevance, usefulness, interpretability, and trustworthiness of discovered knowledge. The research also measures computational cost, model-training time, inference latency, storage requirements, and resource utilization. Statistical analysis is used to determine whether improvements are significant across repeated experiments. The final evaluation considers not only predictive performance but also enterprise suitability, emphasizing privacy, tenant isolation, explainability, scalability, and operational usefulness. This methodology provides a comprehensive basis for determining whether self-supervised AI can transform large volumes of heterogeneous multi-tenant cloud data into reliable and actionable enterprise knowledge.

Advantages

1. Reduced dependence on labeled data: Self-supervised learning can exploit large quantities of unlabeled enterprise data, reducing the cost and effort associated with manual annotation.
2. Improved knowledge discovery: The model can identify semantic relationships and patterns that may not be apparent through traditional keyword-based search or rule-based analysis.
3. Scalability: Self-supervised models can learn from large enterprise datasets containing documents, logs, events, and application information.
4. Better semantic understanding: Representation learning can identify similarities between information expressed using different terminology or formats.
5. Adaptability: Models can be retrained or adapted as enterprise applications, business processes, and tenant requirements evolve.
6. Enhanced anomaly detection: Learning normal patterns from unlabeled operational data can support the identification of unusual events and behaviors.
7. Support for heterogeneous data: The methodology can combine textual, structured, and operational information into common representations.
8. Improved enterprise search: Semantic representations can enable users to discover conceptually relevant information rather than relying only on exact keyword matches.
9. Knowledge-graph enrichment: AI-generated representations can support automatic discovery of entities and relationships, reducing manual knowledge-graph construction effort.
10. Cross-domain pattern discovery: With appropriate privacy controls, the system can identify general patterns across multiple tenants without requiring every tenant to maintain a separate AI infrastructure.
11. Continuous learning potential: New enterprise data can provide additional learning signals, allowing the knowledge-discovery system to evolve over time.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

Disadvantages

1. Privacy risks: Enterprise data may contain sensitive customer, financial, operational, or proprietary information that requires strong protection.
2. Cross-tenant leakage: Poorly designed representations or retrieval mechanisms could unintentionally expose information belonging to one tenant to another.
3. Computational requirements: Training large self-supervised models can require substantial processing power, memory, storage, and energy.
4. Model complexity: Large representation-learning architectures can be difficult to deploy, maintain, monitor, and troubleshoot.
5. Hidden biases: Self-supervised models can learn undesirable patterns or biases present in the underlying enterprise data.
6. Limited explainability: It can be difficult to explain exactly why a model considers two documents, events, or entities to be related.
7. Noisy enterprise data: Logs, documents, and operational records may contain duplicates, missing information, inconsistent terminology, or erroneous values.
8. Semantic ambiguity: The same business term may have different meanings across tenants, potentially producing incorrect relationships.
9. Knowledge hallucination or inaccurate inference: AI-generated relationships may appear meaningful while not representing verified enterprise facts.
10. Model drift: Changes in applications, business processes, tenant behavior, or cloud infrastructure can cause learned representations to become outdated.
11. Integration complexity: Connecting self-supervised AI with data lakes, cloud services, identity systems, knowledge repositories, and enterprise applications requires considerable engineering effort.

IV. RESULTS AND DISCUSSION

The application of self-supervised artificial intelligence (AI) for intelligent enterprise knowledge discovery in multi-tenant cloud environments demonstrates substantial potential for transforming how organizations identify, organize, interpret, and reuse knowledge distributed across complex digital ecosystems. Multi-tenant cloud environments typically contain large volumes of heterogeneous information generated by applications, databases, collaboration platforms, customer interactions, monitoring systems, enterprise resource planning applications, documentation repositories, support systems, source-code platforms, and security infrastructure. Because this information is distributed across different tenants, departments, applications, and storage systems, conventional knowledge-management approaches often struggle to provide a unified understanding of organizational information. Self-supervised learning provides an attractive solution because it can learn meaningful representations from large quantities of predominantly unlabeled enterprise data. Instead of depending entirely on manually labeled datasets, self-supervised models create learning signals from the structure of the available data itself. For example, a model can learn by predicting missing sections of documents, reconstructing masked words, identifying relationships between entities, predicting the next event in an operational sequence, matching related documents, or determining whether two records refer to the same conceptual object. The results indicate that this approach can significantly expand the amount of enterprise information available for intelligent discovery because organizations are no longer required to manually annotate every document, record, event, or interaction before applying machine learning. In a multi-tenant environment, the resulting representations can support semantic search, document classification, entity discovery, knowledge extraction, relationship identification, recommendation, and contextual information retrieval. A particularly important finding is that self-supervised models can identify latent relationships that may not be explicitly represented in conventional enterprise databases. For example, information contained in a technical document, a customer-support case, a software incident, and a project report may refer to the same underlying business problem even though the records use different terminology. Representation learning can identify similarities among these sources and make the relationships discoverable. This improves organizational knowledge discovery by allowing employees to locate relevant information based on meaning rather than exact keywords. However, the results also show that model effectiveness depends on the quality, diversity, and representativeness of training data. Inaccurate, duplicated, outdated, or biased enterprise information can influence learned representations and produce misleading discoveries. Consequently, data-quality management, provenance tracking, and continuous knowledge validation remain essential components of self-supervised enterprise AI.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

A second significant result concerns the ability of self-supervised AI to integrate heterogeneous knowledge across multiple tenants while preserving the contextual distinctions required for secure and accurate enterprise information discovery. Multi-tenancy introduces a fundamental challenge because different customers or organizational units may operate within shared cloud infrastructure while maintaining separate data, policies, workflows, and access requirements. A knowledge-discovery system that treats all information as belonging to a single undifferentiated corpus could produce inappropriate recommendations, expose confidential information, or confuse tenant-specific concepts. The results indicate that effective self-supervised architectures must therefore incorporate tenant-aware representations and access-control mechanisms. Tenant identity, organizational context, data sensitivity, user role, and information ownership can be incorporated as contextual features during knowledge representation and retrieval. This allows the system to learn general patterns while preserving boundaries between tenant-specific knowledge domains. Such an approach also creates opportunities for cross-tenant learning at the representation level without directly exposing raw tenant data. For example, common patterns in software incidents, resource utilization, service failures, or customer requests may be learned across organizations while sensitive documents remain isolated. This can improve model generalization because the AI system benefits from broader structural knowledge without requiring unrestricted data sharing. The findings further suggest that self-supervised learning can support knowledge transfer between related enterprise domains. A model initially trained on technical documentation may develop representations that help interpret support tickets, operational reports, and application metadata. This reduces the need to build separate models for every enterprise information source. Nevertheless, domain differences remain important. The same term can have different meanings across departments or tenants, and the same operational pattern can have different implications depending on business context. Context-sensitive embeddings and tenant-specific adaptation are therefore necessary to prevent semantic ambiguity. The results also highlight the importance of privacy-preserving learning. Techniques such as federated learning, parameter-efficient adaptation, secure aggregation, data minimization, and controlled retrieval can help organizations gain the benefits of self-supervised learning without centralizing sensitive information. The overall finding is that self-supervised AI can create a scalable foundation for multi-tenant knowledge discovery, but only when semantic intelligence is combined with strong tenant isolation, authorization, privacy controls, and contextual modeling.

The results further demonstrate that self-supervised AI can improve enterprise knowledge discovery by supporting dynamic and continuously evolving knowledge representations. Enterprise information does not remain static: policies change, products evolve, software versions are updated, organizational structures are modified, customer requirements shift, and operational incidents create new knowledge. Traditional knowledge repositories frequently depend on manually maintained taxonomies and metadata structures, which can become obsolete when information changes faster than administrators can update them. Self-supervised systems can address this limitation by continuously learning representations from newly generated enterprise information. Streaming data from application logs, support systems, project-management tools, documentation repositories, and operational platforms can provide new learning signals that allow the knowledge-discovery system to detect emerging concepts and relationships. The findings indicate that this capability can improve semantic search and recommendation because the model becomes more responsive to current organizational terminology and operational patterns. For instance, when an enterprise introduces a new software platform, product, or service, employees may begin using new terminology before the organization's formal knowledge taxonomy has been updated. A continuously trained representation model can identify the emerging vocabulary and connect it with related historical information. This can reduce the fragmentation of organizational knowledge and make recent information easier to discover. Self-supervised learning can also support temporal knowledge analysis by distinguishing historical information from current operational states. Such temporal awareness is important because an old document may remain semantically relevant but no longer represent the organization's current policy or technical architecture. The results therefore suggest that knowledge-discovery systems should incorporate timestamps, version information, document validity periods, and event sequences into their representations. Another finding is that self-supervised models can support automated knowledge graph construction by extracting entities and relationships from unstructured and semi-structured enterprise content. These relationships can then be used to connect people, projects, applications, services, policies, products, incidents, and business processes. However, automatically generated knowledge graphs may contain incorrect or uncertain relationships. Future implementations should therefore associate extracted relationships with confidence values, provenance information, and validation mechanisms. Human experts can review high-impact or low-confidence relationships, creating a feedback loop that improves the knowledge base. This demonstrates that self-supervised AI is most effective when used as a continuously learning knowledge-discovery layer rather than as a one-time classification system.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

The final results indicate that self-supervised AI can provide significant operational and strategic benefits when integrated into enterprise cloud knowledge-management systems. At the operational level, intelligent discovery can reduce the time employees spend searching for information, identifying relevant documents, resolving recurring incidents, and locating expertise within an organization. Employees can receive semantically relevant information based on their questions and operational context rather than relying on manually organized repositories. At the managerial level, discovered relationships can reveal recurring process problems, emerging customer requirements, technology dependencies, and knowledge gaps. At the strategic level, enterprise knowledge representations can support decision-making by connecting information across business functions that traditionally operate as information silos. The findings suggest that the greatest benefit occurs when knowledge discovery becomes contextual and proactive. Instead of waiting for users to search for information, AI systems can identify relevant knowledge based on ongoing activities, such as a software deployment, customer complaint, compliance review, or infrastructure incident. For example, when an operational issue occurs, the system could automatically retrieve related historical incidents, relevant troubleshooting documentation, responsible teams, previous resolutions, and associated system dependencies. This reduces mean time to resolution and improves organizational learning. However, the results also identify several risks. AI-generated summaries or relationships can contain factual errors, outdated information, or inappropriate inferences. In enterprise environments, such errors may influence important operational or financial decisions. Knowledge discovery systems must therefore distinguish between retrieved evidence and AI-generated interpretation. Source attribution, confidence indicators, provenance metadata, and human verification should accompany important recommendations. Another issue is computational cost. Continuous self-supervised learning over large multi-tenant datasets can require substantial processing and storage resources. Efficient model architectures, selective retraining, incremental learning, and retrieval-based approaches can reduce these costs. Overall, the results demonstrate that self-supervised AI can significantly enhance enterprise knowledge discovery by transforming fragmented cloud information into a more searchable, connected, contextual, and continuously evolving knowledge ecosystem. Its long-term effectiveness, however, depends on combining representation-learning capabilities with privacy protection, governance, data-quality controls, explainability, and human oversight.

V. CONCLUSION

Self-supervised AI provides a promising foundation for intelligent enterprise knowledge discovery in multi-tenant cloud environments because it addresses one of the central limitations of traditional enterprise machine learning: the dependence on large quantities of manually labeled data. Enterprise environments continuously generate information, but much of this information remains unlabeled, fragmented, inconsistently structured, and distributed across organizational systems. Self-supervised learning can convert this abundant information into useful training signals by exploiting relationships within the data. Through masking, reconstruction, contrastive learning, sequence prediction, relationship prediction, and other self-supervised objectives, AI models can learn representations that capture semantic and structural properties of enterprise information. The study concludes that these representations can support a wide range of knowledge-discovery functions, including semantic search, document retrieval, entity recognition, relationship extraction, recommendation, knowledge-graph construction, incident analysis, and contextual question answering. The most important contribution is the ability to move enterprise knowledge management beyond simple keyword matching and manually constructed categories toward semantic and contextual discovery. This is particularly valuable in multi-tenant cloud environments, where information is distributed across diverse applications and organizational boundaries. Self-supervised AI can identify connections among information sources that use different terminology or formats, thereby making previously hidden knowledge more accessible. Nevertheless, the conclusion is not that self-supervised learning eliminates the need for structured knowledge management. Rather, it provides an intelligent layer that can complement databases, metadata repositories, taxonomies, search engines, and knowledge graphs. Deterministic systems provide reliable storage and access control, while self-supervised models provide flexible semantic interpretation and discovery. The integration of these capabilities can produce a more comprehensive enterprise knowledge architecture.

A second conclusion is that multi-tenancy fundamentally changes how intelligent knowledge discovery should be designed. A conventional centralized AI model may treat all available information as a single learning corpus, but enterprise cloud environments require strict separation of data ownership, access privileges, privacy requirements, and organizational context. The study therefore concludes that tenant-aware knowledge discovery is essential. AI systems must understand not only what information is semantically related but also whether a particular user or application is authorized to access that information. This requires knowledge-discovery architectures in which semantic retrieval is closely integrated with identity and access-management mechanisms. Retrieval should be performed within the user's

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

authorized information space, and generated responses should be grounded in evidence that the user is permitted to view. The conclusion also emphasizes that cross-tenant learning should preferably occur through privacy-preserving mechanisms when direct data sharing is inappropriate. Federated learning and distributed representation learning can allow organizations to benefit from common patterns without transferring raw confidential information into a centralized repository. However, privacy protection should be evaluated alongside model utility because excessive restrictions can reduce the quality of learned representations. The appropriate design therefore requires a balance between knowledge sharing and data sovereignty. Tenant-specific adaptation can further improve accuracy by allowing models to learn organizational terminology and business processes while retaining a shared general representation. This layered approach can support both general enterprise intelligence and tenant-specific knowledge discovery. Ultimately, the study concludes that security and privacy are not external additions to AI-based knowledge discovery; they are fundamental characteristics that must be incorporated into the architecture from the beginning.

The study also concludes that self-supervised AI can support continuous organizational learning by transforming newly generated enterprise information into reusable knowledge. Traditional knowledge-management systems often struggle with knowledge decay because documents become outdated, processes change, and newly learned information is not consistently incorporated into repositories. Self-supervised models can continuously process new information and identify emerging concepts, relationships, and patterns. This creates the possibility of a continuously evolving enterprise knowledge layer that reflects current organizational activities. The ability to connect new information with historical knowledge is especially important for organizational memory. When a recurring technical problem appears, for example, the system can connect the current incident with earlier cases, previous solutions, related documentation, and responsible teams. Similarly, when new policies or product requirements emerge, the system can identify potentially affected processes and information assets. The conclusion is that intelligent knowledge discovery can become an active organizational capability rather than a passive repository function. However, continuous learning introduces risks related to knowledge drift, misinformation, duplicated information, and outdated sources. A self-supervised system could learn from an inaccurate document and subsequently propagate that information into future recommendations. Therefore, knowledge provenance and source quality must be treated as first-class components of the architecture. AI-generated knowledge should retain links to the evidence from which it was derived, and high-impact information should be subjected to validation processes. Temporal metadata should also be maintained so that systems can distinguish current organizational knowledge from historical information. These mechanisms can improve trust and prevent the continuous learning process from degrading the quality of the enterprise knowledge base.

In conclusion, self-supervised AI has the potential to become a central technology for intelligent knowledge discovery in multi-tenant cloud environments by enabling scalable, adaptive, semantic, and privacy-aware analysis of heterogeneous enterprise information. Its primary advantage lies in its ability to learn from the enormous volume of unlabeled information already generated by organizations, reducing the dependence on costly manual annotation while supporting more comprehensive knowledge representation. The study demonstrates that self-supervised approaches can enhance semantic search, knowledge extraction, relationship discovery, recommendation, incident resolution, and organizational intelligence. At the same time, successful deployment requires more than selecting an appropriate AI model. Data quality, tenant isolation, access control, privacy preservation, provenance, explainability, model monitoring, and human oversight are essential for trustworthy enterprise adoption. The recommended architecture is therefore a hybrid knowledge ecosystem in which self-supervised models operate alongside structured databases, knowledge graphs, search systems, identity platforms, governance mechanisms, and human experts. Such an architecture can combine the flexibility of AI with the reliability and control required by enterprise environments. The ultimate value of the technology should be measured not by model size or the number of documents processed, but by improvements in knowledge accessibility, decision quality, employee productivity, incident resolution, organizational learning, and information governance. When implemented responsibly, self-supervised AI can transform fragmented cloud information into an interconnected and continuously evolving organizational knowledge resource. This transformation can provide enterprises with stronger capabilities for discovering hidden relationships, preserving institutional knowledge, responding to emerging requirements, and making better-informed decisions across increasingly complex multi-tenant cloud ecosystems.

VI. FUTURE WORK

Future research should first focus on developing multimodal self-supervised models capable of learning from the full diversity of enterprise information. Existing knowledge-discovery systems frequently emphasize textual data, while enterprise environments contain many other modalities, including source code, structured databases, application logs,

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

network events, system metrics, images, diagrams, audio recordings, workflow information, and infrastructure configurations. Future models should learn shared representations across these modalities so that relationships can be discovered even when relevant information is distributed across different formats. For example, a software incident described in a support ticket could be connected with a log pattern, a source-code component, a monitoring event, and a technical architecture diagram. Multimodal self-supervised learning could enable the system to recognize these relationships without requiring every data source to be manually labeled. Future work should investigate contrastive objectives, masked multimodal reconstruction, cross-modal alignment, and graph-based representation learning for enterprise information. Another important research direction is hierarchical representation learning. Enterprise knowledge exists at multiple levels, from individual records and documents to applications, business processes, departments, organizational units, and strategic objectives. Future models should represent these levels simultaneously and learn how information at one level relates to information at another. This could enable more meaningful knowledge discovery because a low-level technical event could be connected to a high-level business process or service objective. Research should also explore ontology-aware self-supervised learning, in which models learn from enterprise-specific semantic structures while retaining the adaptability of general representation learning. Such approaches could improve interpretation of specialized terminology and reduce semantic ambiguity. Finally, future multimodal systems should incorporate provenance and confidence information directly into their learned representations so that discovered knowledge can be evaluated according to its reliability and source quality.

A second major direction for future work involves privacy-preserving and federated self-supervised learning for multi-tenant environments. Multi-tenancy creates an opportunity to learn general patterns across organizations, but it also introduces significant privacy and data-governance constraints. Future research should investigate federated self-supervised learning methods that allow organizations to collaboratively develop useful representations without transferring raw enterprise information. This is particularly important for organizations that operate under strict contractual, regulatory, or internal data-sovereignty requirements. Future models should be evaluated under realistic heterogeneous conditions in which tenants have different data distributions, business terminology, application architectures, and computational resources. One challenge is that knowledge learned from one tenant may not directly transfer to another tenant because similar concepts can have different meanings. Domain adaptation and personalized federated learning could address this issue by combining shared representations with tenant-specific components. Parameter-efficient adaptation could also reduce the computational cost of maintaining customized models for many tenants. Privacy research should examine the risk that learned model parameters or embeddings could indirectly reveal sensitive information. Differential privacy, secure aggregation, encrypted computation, and access-controlled retrieval should therefore be investigated as complementary safeguards. Future work should also examine adversarial threats to federated self-supervised learning. Malicious tenants could introduce poisoned data, manipulated training signals, or deliberately misleading documents to influence shared representations. Robust aggregation, provenance-aware training, anomaly detection, and tenant reputation mechanisms may help mitigate these threats. Establishing formal privacy and security guarantees while maintaining useful semantic performance will be one of the most important challenges for large-scale multi-tenant knowledge discovery.

Future research should also explore autonomous knowledge-graph construction, temporal reasoning, and continuous knowledge validation. Self-supervised models can identify entities and semantic relationships, but automatically generated knowledge graphs may contain inaccuracies or relationships that are no longer valid. Future systems should therefore combine AI-based extraction with provenance tracking, confidence estimation, temporal reasoning, and validation workflows. Each discovered relationship could include information about its source, creation time, confidence, and validity period. This would allow the knowledge system to distinguish established relationships from tentative inferences. Temporal knowledge graphs could further represent how enterprise concepts change over time. For example, an application may depend on one service during one software version and another service after a migration. A static knowledge graph would not adequately represent this transition, whereas a temporal graph could preserve the history and current state. Future research should also investigate continual self-supervised learning methods that can update representations without requiring complete model retraining whenever new information becomes available. This would reduce computational costs and allow knowledge systems to respond more quickly to organizational change. However, continual learning creates the risk of catastrophic forgetting and uncontrolled propagation of incorrect information. Researchers should therefore develop mechanisms that selectively incorporate new information while retaining important historical knowledge. Human-in-the-loop validation could be particularly valuable for high-impact knowledge. Experts could review uncertain relationships, correct incorrect inferences, and provide feedback that becomes additional training information. Future systems could use active learning strategies to determine which pieces of knowledge require human attention based on uncertainty, business impact, or potential

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

security implications. This would allow human expertise to be concentrated on the most important information while routine knowledge discovery remains automated.

Finally, future work should evaluate the organizational, economic, ethical, and operational consequences of self-supervised enterprise knowledge discovery at production scale. Technical accuracy alone is insufficient because an enterprise knowledge system ultimately influences how employees search for information, make decisions, resolve problems, and understand organizational processes. Future studies should therefore evaluate the effects of AI-based discovery on employee productivity, decision-making accuracy, knowledge retention, collaboration, incident-resolution time, and organizational learning. Comparative experiments should examine traditional enterprise search, supervised machine-learning systems, self-supervised systems, retrieval-augmented architectures, and hybrid knowledge-management approaches under similar conditions. Metrics should include retrieval precision, semantic relevance, knowledge freshness, hallucination rate, response latency, infrastructure cost, privacy risk, and user trust. Research should also investigate how employees interact with AI-generated knowledge recommendations and whether users develop excessive confidence in AI outputs. Explainable interfaces, evidence citations, uncertainty indicators, and source-level traceability could help users distinguish reliable evidence from model-generated interpretation. Another important direction is cost and sustainability. Training and continuously updating large self-supervised models can require substantial computational resources. Future work should investigate efficient architectures, parameter-efficient learning, selective retraining, model compression, retrieval-based approaches, and energy-aware knowledge-processing strategies. Ethical research should address potential biases in enterprise knowledge, unequal representation of organizational groups, inappropriate inference about employees, and the possibility that automated discovery could expose information beyond an individual's legitimate need to know. Ultimately, future research should aim to develop a trustworthy enterprise knowledge ecosystem in which self-supervised AI continuously learns from diverse information, respects tenant and user boundaries, provides evidence-based discoveries, and incorporates human feedback. Such a system would not merely improve search; it would create an adaptive organizational memory capable of connecting fragmented information, preserving institutional knowledge, identifying emerging relationships, and supporting intelligent decision-making across increasingly complex multi-tenant cloud environments.

REFERENCES

1. Sun, M., Xu, Y., Tan, Z., You, D., & Chen, Z. (2025). Multi-level graph contrastive learning for cold-start recommendation in mashup development. *Information Sciences, 717*, 122319. https://doi.org/10.1016/j.ins.2025.122319
2. Prasad, A. (2026, May). Accelerating HPC Performance: Strategies for Overcoming Modern Challenges with NVMe, InfiniBand, and RDMA. In 2026 International Conference on Signal, Systems, and Computing for Next-Gen Automation (ICSSCNA) (pp. 463-471). IEEE.
3. Bellundagi, M. (2023). Integrating Machine Learning with Business Rule Management Systems for Adaptive Enterprise. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8023-8039.
4. Poranki, U. (2026). From Connectivity Monetization to Network Developer Ecosystems: A Conceptual Framework for Reclaiming the 6G Value Layer. International Journal of Computer Information Systems and Industrial Management Applications, 18(6s), 187-195.
5. Gummadi, V. P. K. (2021). Secure API lifecycle management: Integrating MuleSoft Secrets Manager for enterprise data protection. International Journal of Intelligent Systems and Applications in Engineering, 9(4), 537-540.
6. Nisar, K. (2025). Designing a multi-criteria decision framework for enterprise language model adaptation. International Journal of Science, Research and Technology (IJSRAT), 8(6), 15449–15465.
7. Patel, M., & Korat, U. (2026, June). Low-Power Sub-GHz Transceiver Design for Long-Range, Low-Bitrate Wireless Networks. In 2026 IEEE 18th International Conference on Computational Intelligence and Communication Networks (CICN) (pp. 98-103). IEEE.
8. Raja, G. V. (2023). Modernizing enterprise systems using AI with machine learning and cloud computing for intelligent systems. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11713.
9. Sivakumer, D. (2024). The impact of technology industry layoffs on project managers: An empirical study in the USA. International Journal of Research and Applied Innovations (IJRAI), 7(4), 11166–11177.
10. Hossain, I., Hossain, M. S., Rasul, I., Prince, N. U., Datta, A., & Akand, A. R. (2026, June). Machine Learning-Based Evaluation of Password Security and User Awareness for Cyber Risk Prevention. In 2026 Third International Conference on Innovations in Cybersecurity and Data Science (ICICDS) (pp. 499-504). IEEE.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

11. Anand, L. (2025). Modernizing Enterprise Systems through Generative AI Autonomous Operations and Cloud-Native Engineering. International Journal of Humanities and Information Technology, 7(02), 54-69.
12. Pokala, H. K. (2022). From traditional mainframe systems to production machine learning: A practical MLOps framework for healthcare claims processing and revenue cycle optimization in payer organizations. International Journal of Communication Networks and Information Security, 14(2), 847-857.
13. Ambati, K. C. (2026). Enhancing procurement efficiency through integrated master data management and system interoperability. Indian Journal of Computer Science and Technology, 5(1), 600–607.
14. Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. International Journal of Computer Science and Mobile Computing, 10(3), 26-31.
15. Chaba, A. (2025). From systems thinking to model thinking: Embedding AI agents into enterprise CX operating models. International Journal of Computer Technology and Electronics Communication (IJCTEC), 8(4), 11186–11191.
16. Sugumar, R. (2023, September). A Novel Approach to Diabetes Risk Assessment Using Advanced Deep Neural Networks and LSTM Networks. In 2023 International Conference on Network, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
17. Suddala, V. R. A. K. (2026). Transforming life sciences digital ecosystems: Enhancing performance, compliance, and customer experience via automated pipelines. Indian Journal of Computer Science and Technology, 5(1), 592–599.
18. Vas, M. R. (2025). Cyber Resilient SAP Cloud Architecture for Data Governance Intelligent Automation and Scalable Digital Ecosystems. International Journal of Science, Research and Technology, 8(6), 15301-15311.
19. Raja, G. V. (2023). AI Driven Secure Intelligent Framework for Fraud Detection Cybersecurity and Cloud Based Enterprise Systems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(5), 9068-9076.
20. Sivakumer, D., & Punithavel, R. K. (2024). AI-enabled decision intelligence in project management: A systematic review of efficiency and productivity gains. International Journal of Engineering & Extended Technologies Research, 6(2), 7941–7955.
21. Bheemisetty, N. (2026). Framework-driven development of risk management products: Enhancing customization, compliance, and feature reuse. Indian Journal of Computer Science and Technology, 5(1), 616–623.
22. Challa, R. (2022). Optimizing InfiniBand Congestion Control for Large-Scale AI Model Training Workloads. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(6), 5749-5757.
23. Indurthy, V. S. K. (2026). Snowflake and Domain-AI Real-Time Intelligence for Enterprise Data Warehousing. International Journal of Science, Research and Technology, 9(2), 363-372.
24. Tyagi, N. (2025). Privacy Preserving AI in Financial Sector-Balancing Utility, Security and Compliance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 8(5), 12795-12802.
25. Rajula, A. (2024). Adaptive privacy-aware retrieval for federated clinical language models. International Journal of Research and Applied Innovations (IJRAI), 7(3), 10812–10822.
26. Papachappa, H. M., Kumar, A., & Badam, N. (2026, June). Riemannian Intent Manifolds for Session-Based Search: A Joint Embedding Predictive Architecture for Intent Forecasting. In 2026 15th Mediterranean Conference on Embedded Computing (MECO) (pp. 1-8). IEEE.
27. Mathew, A. (2023). Sentinel AI: An Investigation into Robust Threat Mitigation Strategies for Artificial Intelligence. Educational Research (IJMCER), 5(5), 108-111.
28. Karnam, V. S. (2024). Adaptive and federated test automation using AI in distributed multi-cloud and hybrid infrastructure environments. International Journal of Future Innovative Science and Technology (IJFIST), 7(4), 111–121.
29. Pasumarthi, H. (2024). Engineering Large-Scale WMS Integrations: A Practical Guide to Implementing Blue Yonder with IBM ACE, Datapower, MQ, and SAP. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(2), 10008-10016.
30. Narapareddy, V. S. R., & Yerramilli, S. K. (2023). Artificial intelligence incident forecasting. International Journal of Engineering Technology Research & Management, 7(12), 551–559.
31. Anand, L. (2023). Machine Learning Enabled Enterprise Integration through Intelligent API Governance Secure Cloud Infrastructure and Automated Operations. International Journal of Research and Applied Innovations, 6(3), 5972-5979.
32. Bhagwat, V. B. (2024). A simplified transition from EBS Payroll to Cloud Payroll: Benefits and Drawbacks. Journal of Computational Analysis and Applications, 33(6).
33. Juvvadi, R. R. (2025). Toward autonomous finance: A multi-agent system architecture for the self-driving financial close. International Journal of Engineering & Extended Technologies Research (IJEETR), 7(6), 11290–11295.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104004

34. Soundappan, S. J. (2023). AI-Driven Secure Enterprise Analytics and Intelligent Cloud Data Management Frameworks. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(3), 8236-8242.
35. Zhang, H., et al. (2025). A self-supervised method for learning path-augmented knowledge graph embedding. *Engineering Applications of Artificial Intelligence*, 112315. https://doi.org/10.1016/j.engappai.2025.112315
36. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.
Prasad, A. (2026, May). Accelerating HPC Performance: Strategies for Overcoming Modern Challenges with NVMe, InfiniBand, and RDMA. In 2026 International Conference on Signal, Systems, and Computing for Next-Gen Automation (ICSSCNA) (pp. 463-471). IEEE.