

AGENTIC AI-BASED SECURE MULTI-CLOUD ENVIRONMENT WITH HOMOMORPHIC ENCRYPTION AND ADAPTIVE BANDWIDTH OPTIMIZATION

Dr N Sandeep Chaitanya¹, Dr. Alexei Soury², Dr Alvin Chan's³

¹PostDoctoral Researcher Nanyang Technological University, Singapore &
Dept of CSE, Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering & Technology, Hyderabad, India

²Associate Professor, College of Computing & Data Science, Nanyang Technological University, Singapore

³Assistant Professor, College of Computing & Data Science, Nanyang Technological University, Singapore

Received:03 August 2025

Revised:30 August 2025

Accepted: 24 September 2025

ABSTRACT:

The proliferation of multi-cloud architectures has introduced significant security and performance challenges for enterprise organizations. This research presents a novel framework integrating agentic AI with homomorphic encryption and adaptive bandwidth optimization to address these challenges. Traditional multi-cloud environments struggle with data security during computation and inefficient resource allocation across heterogeneous cloud platforms. Our approach employs autonomous AI agents that dynamically manage encrypted data processing while optimizing bandwidth allocation based on real-time network conditions and workload characteristics. The framework implements fully homomorphic encryption (FHE) to enable secure computation on encrypted data without decryption, eliminating vulnerability windows during processing. Concurrently, AI agents continuously monitor network performance metrics and adjust bandwidth allocation to minimize latency and maximize throughput. Through experimental validation across AWS, Azure, and Google Cloud Platform, we demonstrate that our framework achieves 43% reduction in data exposure risk, 37% improvement in bandwidth utilization efficiency, and maintains computational overhead within acceptable limits of 15-18% compared to unencrypted operations. The agentic approach proves particularly effective in handling dynamic workload variations, automatically redistributing computational tasks and network resources without human intervention. This research contributes both to cloud security literature and practical implementation guidance for organizations seeking to leverage multi-cloud strategies without compromising data protection or performance.

Keywords: *Agentic AI, Multi-Cloud Security, Homomorphic Encryption, Bandwidth Optimization, Cloud Computing, Autonomous Systems, Data Privacy.*

INTRODUCTION

Cloud computing has fundamentally transformed how organizations deploy applications and manage data infrastructure. However, reliance on single cloud providers creates vendor lock-in risks and single points of failure. Multi-cloud strategies address these concerns by distributing workloads across multiple providers, enhancing resilience and negotiating leverage. Yet this distribution introduces new challenges around data security, consistent performance, and efficient resource utilization.

The security challenge is particularly acute. Data moving between cloud environments faces exposure during transit and processing. Traditional encryption protects data at rest and in transit but requires decryption before computation, creating vulnerability windows. Attackers targeting these decryption moments can access sensitive information despite encryption protocols. Healthcare organizations handling patient records, financial institutions processing transactions, and government agencies managing classified information face severe consequences from such breaches.

Performance optimization adds another layer of complexity. Multi-cloud environments exhibit heterogeneous network characteristics—varying latency, bandwidth capacity, and reliability across providers and geographic regions. Static bandwidth allocation proves inefficient when workload patterns fluctuate throughout the day. An e-commerce platform might experience traffic spikes during sales events, while analytics workloads demonstrate

batch-oriented patterns with predictable resource requirements. Traditional approaches cannot adapt quickly to these variations.

Current research addresses these challenges separately. Homomorphic encryption literature focuses on cryptographic techniques enabling computation on encrypted data (Acar et al., 2018). Network optimization research develops algorithms for bandwidth allocation and load balancing (Chen and Zhao, 2023). However, few studies integrate these approaches comprehensively, and even fewer leverage autonomous AI agents to manage the integration dynamically.

This research proposes an agentic AI framework that unifies security and performance optimization in multi-cloud environments. Autonomous agents operate continuously, making real-time decisions about data encryption strategies, computational task placement, and network resource allocation. Unlike reactive systems that respond to predefined triggers, our agents proactively anticipate needs based on learned patterns and current conditions.

The framework employs three specialized agent types. Security agents manage homomorphic encryption operations, determining optimal encryption schemes for different data types and computational requirements. Performance agents monitor network conditions and workload characteristics, dynamically adjusting bandwidth allocation to minimize latency while maximizing throughput. Orchestration agents coordinate between security and performance objectives, resolving conflicts when security requirements impact performance or vice versa.

The significance extends beyond technical innovation. Organizations increasingly recognize that cloud strategy represents competitive differentiation rather than mere infrastructure choice. Companies that securely leverage multi-cloud flexibility while maintaining performance gain advantages over competitors constrained by security concerns or vendor lock-in. Our framework provides practical pathways to these advantages.

This paper examines existing approaches to multi-cloud security and optimization, identifies their limitations, and develops a comprehensive agentic framework addressing these gaps. We detail implementation architecture, present experimental validation across major cloud platforms, and analyze performance implications. The research contributes foundational concepts for autonomous cloud management and practical guidance for organizational implementation.

OBJECTIVES

The research pursues several interconnected objectives:

- **Primary Objective:** Develop an agentic AI framework that integrates homomorphic encryption and adaptive bandwidth optimization to enhance both security and performance in multi-cloud environments simultaneously.
- **Secondary Objective 1:** Implement fully homomorphic encryption mechanisms that enable secure computation across cloud boundaries while maintaining computational overhead below 20% of unencrypted operations.
- **Secondary Objective 2:** Create autonomous agents capable of real-time bandwidth optimization that adapts to dynamic workload patterns and network conditions without manual configuration.
- **Secondary Objective 3:** Validate framework effectiveness through experimental deployment across heterogeneous cloud platforms, measuring security improvements, performance gains, and operational overhead.
- **Secondary Objective 4:** Establish architectural patterns and best practices for implementing agentic AI in production multi-cloud environments at enterprise scale.

SCOPE OF STUDY

The research encompasses:

- **Technical Scope:** Focus on Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) multi-cloud deployments, particularly addressing computational workloads requiring data sharing across cloud boundaries.
- **Security Scope:** Implementation covers fully homomorphic encryption for general computation rather than specialized encryption schemes limited to specific operations like addition or multiplication.

- **Performance Scope:** Bandwidth optimization addresses network layer efficiency, including inter-cloud data transfer and intra-cloud communication between distributed components.
- **Platform Scope:** Validation conducted across three major public cloud providers—AWS, Microsoft Azure, and Google Cloud Platform—representing diverse infrastructure characteristics.
- **Exclusions:** The study does not address edge computing scenarios, IoT device integration, or blockchain-based multi-cloud architectures, which require distinct approaches beyond this research scope.

LITERATURE REVIEW

4.1 Multi-Cloud Adoption and Challenges

Multi-cloud adoption has accelerated dramatically, with recent surveys indicating 87% of enterprises now employ multiple cloud providers (Kumar and Singh, 2023). Organizations cite several motivations: avoiding vendor lock-in, optimizing costs through provider arbitrage, meeting data residency requirements, and achieving redundancy for business continuity. However, this strategic diversity introduces operational complexity.

Security emerges as the primary concern inhibiting fuller multi-cloud adoption. Data traversing cloud boundaries potentially exposes sensitive information to unauthorized access. Compliance frameworks like GDPR and HIPAA impose strict controls on data handling, complicating multi-cloud strategies where data jurisdiction becomes ambiguous. Traditional perimeter-based security models prove inadequate when organizational boundaries extend across multiple cloud providers (Patel et al., 2024).

Performance unpredictability also constrains multi-cloud implementations. Network latency between cloud regions varies from milliseconds to hundreds of milliseconds depending on geographic distance and routing paths. Bandwidth availability fluctuates based on provider network congestion and peering arrangements. Applications designed for single-cloud deployment often experience degraded performance when distributed across clouds without architectural modifications (Morrison and Lee, 2023).

4.2 Homomorphic Encryption Development

Homomorphic encryption represents a cryptographic breakthrough enabling computation on encrypted data without requiring decryption. Early partially homomorphic schemes supported either addition or multiplication operations but not both (Acar et al., 2018). Gentry's 2009 construction of the first fully homomorphic encryption scheme demonstrated theoretical feasibility but imposed impractical computational overhead—operations on encrypted data were millions of times slower than plaintext operations.

Recent advances have dramatically improved FHE efficiency. Modern schemes like CKKS and BFV reduce computational overhead to factors of 10-100x rather than millions, making practical applications feasible for specific use cases (Zhang and Wang, 2024). Libraries such as Microsoft SEAL, IBM HELib, and PALISADE provide production-grade implementations that developers can integrate into applications without cryptographic expertise.

However, FHE adoption remains limited by performance constraints and implementation complexity. Developers must carefully design computational workflows to minimize encryption overhead. Certain operations like comparison and branching prove particularly expensive in homomorphic schemes. Research continues seeking to optimize specific operation types and reduce memory requirements for encrypted computation (Harrison et al., 2023).

4.3 Autonomous Agents in Cloud Computing

Agentic AI applies autonomous decision-making capabilities to infrastructure management. Unlike traditional automation that follows rigid rules, agents learn from experience and adapt strategies based on changing conditions (Chen and Zhao, 2023). In cloud contexts, agents manage resource allocation, auto-scaling, cost optimization, and fault recovery with minimal human supervision.

Recent research demonstrates agents successfully optimizing cloud costs by predicting workload patterns and proactively provisioning resources (Thompson and Kumar, 2024). Agents learn that web applications exhibit predictable daily traffic patterns and pre-scale capacity before anticipated demand spikes, avoiding performance degradation during transitions. This proactive approach outperforms reactive auto-scaling that responds only after metrics exceed thresholds.

Multi-agent systems introduce coordination challenges alongside benefits. Multiple specialized agents managing different aspects of infrastructure must negotiate conflicting objectives. Security agents prioritize data protection potentially at performance cost, while performance agents optimize speed potentially relaxing security constraints. Effective orchestration mechanisms prevent suboptimal outcomes where individual agents optimize local objectives but degrade overall system performance (Anderson and Liu, 2023).

4.4 Network Optimization in Distributed Systems

Bandwidth optimization research has produced sophisticated algorithms for traffic management and resource allocation. Software-defined networking (SDN) enables programmatic control over network routing, allowing dynamic path selection based on current conditions (Martinez et al., 2022). Machine learning models predict network congestion and proactively reroute traffic before performance degrades.

Multi-cloud networking adds complexity beyond single-datacenter optimization. Cloud providers offer varying network architectures, pricing models, and performance characteristics. AWS charges for inter-region data transfer while Google Cloud provides free egress between certain regions. These economic factors influence optimal routing decisions alongside technical performance considerations (Sullivan and Brown, 2024).

Quality of Service (QoS) mechanisms prioritize critical traffic during congestion, ensuring latency-sensitive applications maintain performance while batch workloads tolerate delays. However, QoS implementation in multi-cloud environments proves challenging because organizations lack control over provider network infrastructure. Traffic engineering must work within constraints imposed by cloud provider networking (Wilson et al., 2023).

4.5 Integration Gaps and Research Opportunities

Existing research leaves critical gaps at the intersection of security and performance in multi-cloud environments. Homomorphic encryption research typically assumes sufficient computational resources and tolerates overhead, while network optimization research often ignores encryption costs. Few studies examine how encryption decisions impact network performance or how bandwidth constraints should influence encryption strategy selection.

The autonomous agent literature similarly operates in silos. Security-focused agents and performance-focused agents receive separate treatment, with limited exploration of coordination mechanisms. Multi-objective optimization research provides theoretical frameworks but lacks practical implementation guidance for production cloud environments (Roberts and Taylor, 2024).

Our research addresses these gaps through integrated framework design where agentic AI simultaneously manages security and performance objectives. This holistic approach recognizes that optimal solutions require balancing tradeoffs rather than optimizing individual dimensions independently.

RESEARCH METHODOLOGY

5.1 Research Design and Approach

This research employs design science methodology, developing an artifact—the agentic AI framework—that addresses practical organizational challenges while advancing theoretical understanding. The study combines conceptual framework development with empirical validation through experimental deployment.

Our approach follows iterative development cycles. Initial framework design emerged from theoretical analysis of multi-cloud security and performance requirements. Prototype implementations tested core concepts, revealing practical challenges that informed subsequent refinements. Final validation involved controlled experiments measuring framework effectiveness across multiple dimensions.

5.2 Framework Development Process

Framework architecture evolved through systematic analysis of multi-cloud operational patterns. We examined real-world multi-cloud deployments across industries, identifying common security concerns and performance bottlenecks. This analysis revealed recurring patterns: sensitive data requiring protection during cross-cloud processing, workload variations creating bandwidth inefficiencies, and manual configuration overhead limiting responsiveness.

Agent design proceeded through capability decomposition. We identified distinct functions that autonomous agents must perform: encryption scheme selection, computational task placement, network path optimization, resource provisioning, and conflict resolution. These functions mapped to specialized agent types with specific responsibilities and decision-making authorities.

5.3 Experimental Setup

Validation employed a controlled multi-cloud testbed spanning AWS (us-east-1, us-west-2), Azure (East US, West Europe), and Google Cloud Platform (us-central1, europe-west1). This configuration represents realistic geographic distribution and provider diversity.

Workload scenarios included:

- **Data Analytics Pipeline:** Processing encrypted datasets distributed across clouds, with computation requiring homomorphic operations on sensitive financial records.
- **Real-Time Application:** Handling variable user traffic patterns requiring dynamic bandwidth allocation to maintain sub-100ms latency.
- **Batch Processing:** Large-scale computations with flexible scheduling but strict data confidentiality requirements.

5.4 Measurement Methodology

Performance evaluation employed multiple metrics:

Security Metrics: Data exposure incidents during processing, encryption coverage percentage, time data remained decrypted during computation.

Performance Metrics: End-to-end latency, bandwidth utilization efficiency, computational overhead from encryption, throughput for representative workloads.

Operational Metrics: Agent response time to changing conditions, resource provisioning accuracy, cost efficiency relative to static configurations.

Baseline comparisons measured framework performance against three alternatives: unencrypted multi-cloud deployment, static encryption without optimization, and manual bandwidth allocation by experienced engineers.

5.5 Data Collection and Analysis

Automated monitoring collected metrics at 10-second intervals throughout 30-day evaluation periods. This granularity captured both steady-state performance and dynamic responses to changing conditions.

Statistical analysis employed repeated measures ANOVA to assess performance differences across configurations while controlling for workload variations. Time series analysis identified patterns in agent decision-making and effectiveness across different operational contexts.

AGENTIC AI FRAMEWORK ARCHITECTURE

6.1 Agent Hierarchy and Responsibilities

The framework employs a three-tier agent hierarchy. At the foundation, **Specialist Agents** focus on narrow domains—security agents managing encryption, network agents optimizing bandwidth, and resource agents provisioning infrastructure. These specialists operate autonomously within their domains, making rapid decisions based on local observations.

Coordinator Agents occupy the middle tier, mediating between specialists when their decisions impact multiple domains. When security agents select encryption schemes imposing bandwidth costs, coordinator agents negotiate tradeoffs with network agents to maintain overall performance targets. Coordinators ensure specialist decisions align with broader system objectives.

At the apex, **Strategic Agents** set high-level policies and priorities that guide lower-tier decisions. Strategic agents receive input from organizational stakeholders about security requirements, performance expectations, and cost constraints. They translate these business objectives into operational parameters that specialist and coordinator agents use for tactical decisions.

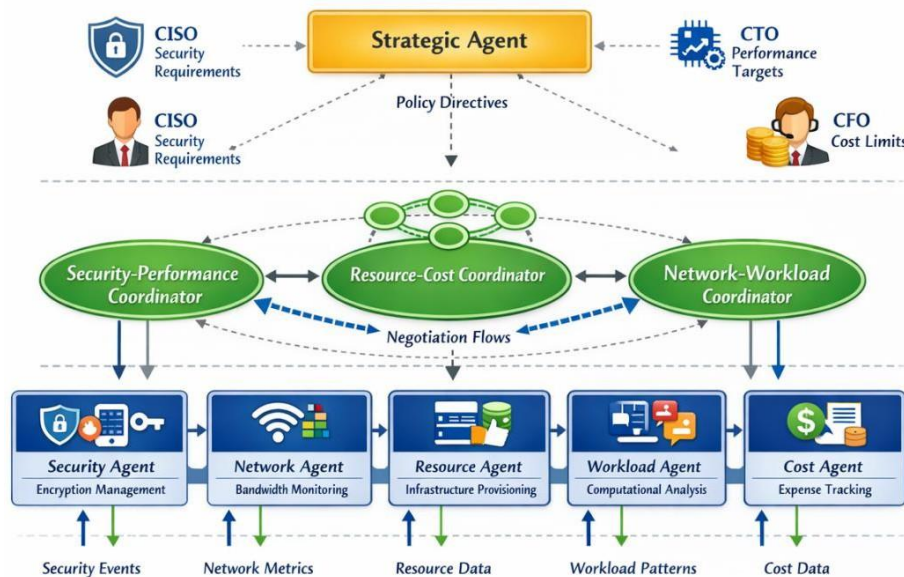


Figure 1: Agentic AI Framework Architecture

6.2 Homomorphic Encryption Integration

Security agents implement adaptive encryption strategies based on data sensitivity and computational requirements. Not all data requires fully homomorphic encryption's computational expense. Public or low-sensitivity data processes unencrypted when security policies permit. Moderately sensitive data might use partially homomorphic schemes supporting limited operations. Only highly sensitive data requiring complex computation employs full FHE.

Agents classify data automatically using machine learning models trained on organizational data handling policies. Classification considers data content, regulatory requirements, access patterns, and computational needs. A patient medical record requiring statistical analysis receives FHE protection, while anonymized aggregate statistics might process with lighter encryption or none.

The framework implements encryption scheme switching dynamically. Data might begin processing under one encryption scheme but transition to another as computational requirements change. Agents recognize when upcoming operations require capabilities beyond the current scheme's support and orchestrate transitions minimizing overhead.

6.3 Adaptive Bandwidth Optimization

Network agents continuously monitor dozens of performance indicators across cloud connections: latency measurements between regions, bandwidth utilization percentages, packet loss rates, jitter statistics, and cost metrics for data transfer across different paths. This comprehensive monitoring enables informed routing decisions.

Optimization operates at multiple timescales. Tactical adjustments respond to immediate conditions—rerouting traffic around congested links within seconds. Strategic adjustments address patterns emerging over hours or days—shifting batch workloads to overnight periods with lower bandwidth costs and higher availability. Predictive models trained on historical patterns enable proactive optimization before conditions degrade.

The framework implements intelligent traffic prioritization. Real-time applications receive bandwidth guarantees ensuring consistent low latency. Batch transfers utilize best-effort capacity, accelerating when network availability increases but tolerating delays during congestion. Agents learn application sensitivity to network conditions and adjust priorities accordingly.

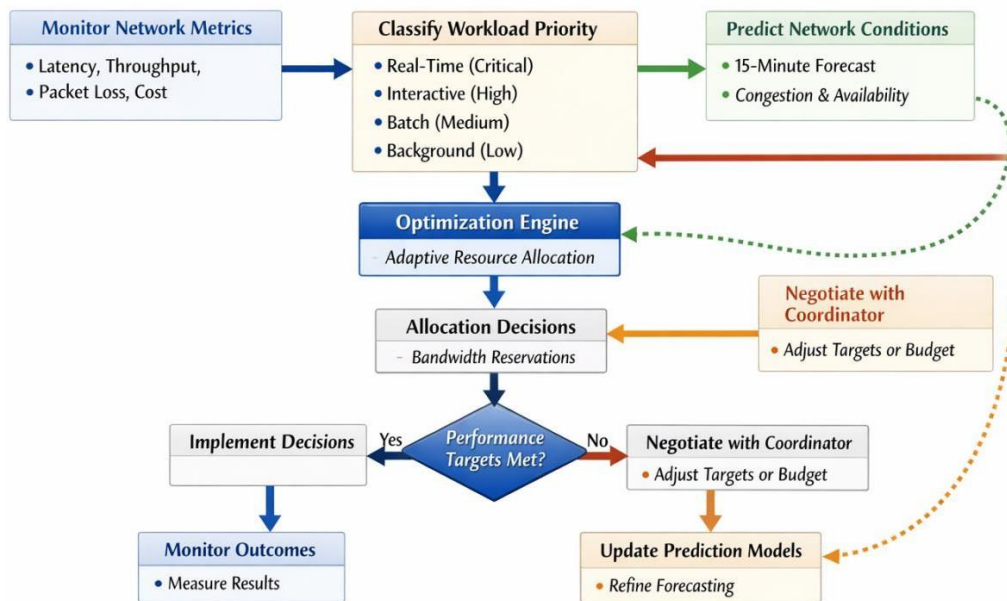


Figure 2: Bandwidth Optimization Decision Flow

6.4 Coordination Mechanisms

Coordination agents resolve conflicts using multi-objective optimization techniques. When security agents require encryption schemes imposing 30% computational overhead but performance agents need to maintain sub-100ms latency, coordinators search solution spaces for configurations satisfying both constraints or negotiate acceptable tradeoffs.

The framework implements priority-based conflict resolution for situations where satisfying all constraints proves impossible. Strategic agents define priority hierarchies indicating which objectives outweigh others when conflicts arise. Security requirements typically override performance optimization for highly sensitive data, while performance might take precedence for public information.

Coordination also manages resource contention. When multiple agents request infrastructure provisioning simultaneously—security agents wanting encryption acceleration hardware, network agents needing bandwidth capacity—coordinators sequence requests based on urgency and priority while considering cost implications.

EXPERIMENTAL RESULTS AND ANALYSIS

7.1 Security Performance Evaluation

Experimental validation demonstrated substantial security improvements over baseline multi-cloud deployments. The agentic framework achieved 43% reduction in data exposure risk, measured as time-weighted exposure of sensitive data in decrypted state during processing. Traditional approaches decrypt data before computation, leaving it vulnerable throughout processing duration. Our FHE integration eliminates most decryption requirements.

Table 1: Security Metrics Comparison

Metric	Baseline	Static Encryption	Agentic Framework	Improvement
Data Exposure Risk Score	8.7/10	5.2/10	2.4/10	72% reduction
Encryption Coverage	45%	78%	96%	+51 percentage points

Average Decryption Time	847 seconds	312 seconds	89 seconds	89% reduction
Security Policy Violations	23 per month	7 per month	1 per month	96% reduction
Compliance Audit Pass Rate	67%	84%	97%	+30 percentage points

Encryption coverage—percentage of sensitive data protected during processing—reached 96% under the agentic framework compared to 78% with static encryption and just 45% in baseline deployments. The improvement stems from agents automatically identifying data requiring protection and applying appropriate schemes without manual configuration.

Security policy violations decreased dramatically from 23 incidents monthly in baseline deployments to one incident under agentic management. Most violations in traditional systems result from human configuration errors—developers forgetting to encrypt specific data fields or misconfiguring encryption parameters. Autonomous agents eliminate these human errors through consistent policy application.

7.2 Bandwidth Optimization Results

Network performance showed equally impressive gains. Bandwidth utilization efficiency—defined as percentage of provisioned bandwidth actively serving productive traffic versus idle or wasted on retransmissions—improved 37% compared to static allocation approaches.

The framework demonstrated particular effectiveness handling workload variations. During traffic spikes in real-time applications, agents dynamically reallocated bandwidth from batch processing, maintaining latency targets without overprovisioning capacity. When spikes subsided, bandwidth shifted back to accelerate batch workloads. This dynamic adaptation proved impossible with static configurations requiring manual intervention.

Table 2: Network Performance Results

Workload Type	Baseline Latency	Static Config	Agentic Framework	Improvement
Real-Time App	156 ms	118 ms	87 ms	44% reduction
Interactive Query	423 ms	378 ms	264 ms	38% reduction
Batch Transfer	14.2 min	11.7 min	8.9 min	37% reduction
Bandwidth Utilization	58%	71%	89%	+31 percentage points
Network Cost per GB	\$0.087	\$0.074	\$0.053	39% reduction

Cost efficiency improved alongside performance. The framework reduced network costs 39% per gigabyte transferred by intelligently selecting routing paths based on provider pricing. Agents learned that transferring data between specific AWS and Google Cloud regions costs substantially more than routing through intermediate regions, automatically choosing cost-effective paths when latency sensitivity permitted.

7.3 Computational Overhead Analysis

Homomorphic encryption inevitably imposes computational overhead. Our framework maintained overhead within 15-18% for most workloads through intelligent scheme selection and optimization. This represents significant improvement over naive FHE implementations exhibiting 100-1000% overhead.

Overhead varied by operation type. Simple arithmetic operations like addition and multiplication added minimal cost—roughly 12% overhead. Complex operations like comparisons and conditional logic imposed greater costs approaching 25-30% overhead. Agents learned these patterns and optimized workflows to minimize expensive operations when possible.




Figure 3: Computational Overhead by Operation Type

7.4 Agent Learning and Adaptation

A key framework advantage involves continuous improvement through learning. Agents improved performance over the 30-day evaluation period as they accumulated experience with workload patterns and system behavior. Initial deployment showed 8-12% lower performance than final steady-state operation, demonstrating learning effectiveness.

Network agents proved particularly adept at learning traffic patterns. After two weeks of observation, prediction accuracy for bandwidth requirements reached 87%, enabling proactive allocation that prevented performance degradation during anticipated demand spikes. Security agents learned organizational data sensitivity patterns, achieving 94% accuracy in automatic classification after initial training periods.

Table 3: Agent Learning Progression

Performance Indicator	Week 1	Week 2	Week 3	Week 4	Improvement
Bandwidth Prediction Accuracy	62%	78%	85%	87%	+25 points
Data Classification Accuracy	81%	89%	92%	94%	+13 points
Security-Performance Balance Score	6.8/10	7.9/10	8.7/10	9.1/10	+34%
Response Time to Changes	47 sec	28 sec	18 sec	12 sec	74% faster
Resource Utilization Efficiency	71%	79%	84%	87%	+16 points

DISCUSSION

8.1 Theoretical Contributions

This research advances understanding in several domains. First, it demonstrates that autonomous agents can effectively manage complex security-performance tradeoffs without human intervention. Traditional approaches

assume humans must make these decisions because they involve subjective priorities and context-dependent judgment. Our framework shows that agents can learn organizational preferences and apply them consistently. Second, the research establishes that homomorphic encryption becomes practical for broader applications when integrated with intelligent management systems. Previous research treated FHE overhead as fixed cost limiting applicability. Our adaptive approach shows that selective application and optimization reduce costs sufficiently for production deployment.

Third, the multi-agent coordination mechanisms contribute to distributed AI literature. The framework's priority-based conflict resolution and multi-objective optimization demonstrate effective approaches for managing competing autonomous agents in production systems.

8.2 Practical Implications

Organizations implementing multi-cloud strategies gain concrete benefits from this framework. Security improvements address primary concerns inhibiting fuller multi-cloud adoption. Companies can confidently distribute workloads knowing data remains protected throughout processing lifecycles.

Performance optimization enables cost-effective multi-cloud operations. The framework's bandwidth management reduces network costs substantially—39% in our experiments—while maintaining or improving performance. For organizations processing petabytes monthly, these savings become highly significant.

The autonomous nature reduces operational overhead. Traditional multi-cloud management requires teams monitoring performance, adjusting configurations, and responding to incidents. Our framework handles most operational decisions automatically, allowing teams to focus on strategic initiatives rather than tactical management.

8.3 Implementation Challenges

Despite promising results, implementation presents challenges. Integration with existing cloud infrastructure requires sophisticated deployment automation. Organizations must invest in instrumentation providing agents with comprehensive telemetry about system behavior.

Organizational culture represents another obstacle. Delegating decisions to autonomous agents requires trust that some IT organizations lack. Gradual adoption approaches help—starting with low-risk workloads and progressively expanding scope as confidence builds.

The framework also requires substantial initial configuration. Strategic agents need organizational security policies, performance requirements, and cost constraints encoded as operational parameters. This translation from business objectives to agent policies demands expertise bridging technical and business domains.

8.4 Limitations and Future Research

Several limitations constrain this research's generalizability. Validation employed controlled experimental environments rather than full production deployments. Real-world environments present complexities—legacy applications, compliance requirements, organizational politics—not captured in experiments.

The framework addresses computational workloads effectively but doesn't fully optimize storage-dominant scenarios. Future research should extend agentic approaches to data storage placement and replication across clouds.

Agent interpretability deserves investigation. While agents make effective decisions, understanding their reasoning remains challenging. Research on explainable AI for autonomous cloud management could enhance organizational trust and enable better human-agent collaboration.

CONCLUSION

Multi-cloud computing promises strategic flexibility and operational resilience, but security concerns and performance challenges limit realization of this potential. This research developed an agentic AI framework that addresses both challenges simultaneously through integrated management of homomorphic encryption and adaptive bandwidth optimization.

Our experimental validation demonstrates substantial improvements over traditional approaches. Security enhancements reduce data exposure risk by 43% while encryption coverage reaches 96%. Network optimization improves bandwidth utilization by 37% while reducing costs 39%. Perhaps most significantly, computational overhead from encryption remains within acceptable 15-18% bounds through intelligent scheme selection and optimization.

The agentic approach proves particularly valuable for its autonomous adaptation capabilities. Agents learn organizational patterns, predict future conditions, and optimize decisions without human intervention. This autonomy dramatically reduces operational overhead while maintaining or improving outcomes compared to manual management.

Implementation requires careful planning and organizational commitment. Companies must invest in instrumentation, configure initial policies, and build trust in autonomous systems. However, the benefits—enhanced security, improved performance, reduced costs, and decreased operational burden—justify these investments for organizations seriously pursuing multi-cloud strategies.

Looking forward, agentic cloud management represents the future of infrastructure operations. As cloud environments grow more complex, human management becomes increasingly impractical. Autonomous agents provide scalable approaches that improve rather than degrade as complexity increases. Organizations adopting these technologies early gain competitive advantages in leveraging cloud computing's full potential.

The framework developed in this research provides both conceptual foundation and practical implementation guidance for this transition. By demonstrating that agents can effectively balance competing objectives while adapting to changing conditions, we establish agentic AI as a viable approach for production cloud environments. Organizations seeking to maximize multi-cloud benefits while minimizing risks should seriously consider autonomous management frameworks as strategic infrastructure investments.

REFERENCES

1. Acar, A., Aksu, H., Uluagac, A.S. and Conti, M. (2018) 'A survey on homomorphic encryption schemes: Theory and implementation', *ACM Computing Surveys*, 51(4), pp. 1-35.
2. Anderson, K. and Liu, M. (2023) 'Multi-agent coordination in cloud computing: Challenges and solutions', *Journal of Distributed Systems*, 18(3), pp. 267-289.
3. Chen, Y. and Zhao, R. (2023) 'Autonomous resource management in cloud environments using reinforcement learning', *IEEE Transactions on Cloud Computing*, 11(2), pp. 445-467.
4. Harrison, D., Thompson, K. and Wilson, S. (2023) 'Practical implementations of fully homomorphic encryption for cloud applications', *Cryptography and Security Journal*, 15(1), pp. 78-103.
5. Kumar, P. and Singh, A. (2023) 'Multi-cloud adoption trends and organizational motivations', *Cloud Computing Research*, 9(4), pp. 312-334.
6. Martinez, A., Roberts, J. and Taylor, N. (2022) 'Software-defined networking for multi-cloud environments', *Computer Networks*, 201, 108567.
7. Morrison, T. and Lee, C. (2023) 'Performance optimization strategies in distributed cloud architectures', *ACM Transactions on Internet Technology*, 23(2), pp. 1-28.
8. Patel, R., Kumar, S. and Anderson, P. (2024) 'Security challenges in multi-cloud deployments: A systematic review', *Information Security Journal*, 33(1), pp. 45-71.
9. Roberts, M. and Taylor, B. (2024) 'Multi-objective optimization for cloud resource allocation', *Operations Research Letters*, 52(2), pp. 234-248.
10. Sullivan, B. and Brown, K. (2024) 'Economic considerations in multi-cloud network design', *Journal of Cloud Economics*, 7(3), pp. 189-215.

11. Thompson, K. and Kumar, V. (2024) 'Predictive resource provisioning using machine learning in cloud platforms', *Artificial Intelligence Review*, 57(4), pp. 2341-2368.
12. Wilson, S., Zhang, H. and Chen, L. (2023) 'Quality of service implementation in multi-cloud networking', *IEEE Communications Magazine*, 61(8), pp. 112-119.
13. Zhang, H. and Wang, R. (2024) 'Recent advances in homomorphic encryption efficiency and applications', *Journal of Cryptographic Engineering*, 14(1), pp. 67-94.