JOURNAL LA MULTIAPP

VOL. 07, ISSUE 03 (747-759), 2026
DOI: 10.37899/journallamultiapp.v7i3.3220

Comparative Evaluation of Three-Tier and Cisco ACI Data Center Architectures

Andi Muhammad Bintang Fadhilah[1], Ari Ramadhan[1]

[1]*Politeknik IDN Bogor, Indonesia*

**Corresponding Author: Andi Muhammad Bintang Fadhilah*
E-mail: bintang@idn.ac.id

Article Info

Article history:
Received 2 April 2026
Received in revised form 8 May 2026
Accepted 3 August 2026

Keywords:
Data Center
Network Architecture
Scalability
Security
Performance

Abstract

This study aims to conduct a comprehensive comparative evaluation of traditional three-tier and application-centric data center network architectures within an enterprise environment. The research adopts a descriptive qualitative approach using a comparative case study design combined with design evaluation techniques. Empirical data were collected from a real-world data center operating both architectures concurrently across multiple sites, supported by documentation analysis, operational monitoring, and direct observation. The analysis is structured around five key capability dimensions, namely scalability, security, reliability, operational simplicity and agility, as well as network visibility and virtualization awareness. The findings demonstrate that the application-centric architecture significantly outperforms the traditional model in scalability, security, operational efficiency, and visibility. It supports a higher number of endpoints and nodes with lower resource utilization, achieves reduced network latency, enables fine-grained security segmentation, and provides centralized policy-based management with integrated monitoring capabilities. Although the traditional architecture shows slightly higher aggregate availability, this is primarily influenced by operational maturity rather than architectural superiority. The results indicate that the application-centric approach offers a more adaptive, scalable, and secure solution for modern enterprise data center environments. This study contributes practical insights for network engineers and decision-makers by providing an evidence-based evaluation framework and highlighting key considerations for infrastructure transformation strategies.

Introduction

The exponential development of digital technologies and the transformation of business processes in the era of cloud computing have positioned data center infrastructure as a critical component supporting modern organizational operations (Maloo & Nikolov, 2022; Maloo, Nikolov, & Ahmed, 2023; Mujib & Sari, 2020; Nunes, Mendonca, Nguyen, Obraczka, & Turletti, 2014). Data centers are no longer passive repositories of computational resources; rather, they function as integrated environments that support mission-critical services across sectors such as finance, government, e-commerce, and cloud computing ecosystems (Janovic, 2023; Rout, Sahoo, Patra, Sahoo, & Puthal, 2021). The rapid growth of artificial intelligence workloads, distributed computing, and microservices-based applications has significantly increased the complexity of network demands, requiring infrastructures that are not only reliable but also highly scalable, secure, agile, and capable of providing comprehensive network visibility (Frank & Beard, 2025; Maloo & Nikolov, 2022; Osundare & Bolatito Ige, 2024).

ISSN: 2716-3865 (Print), 2721-1290 (Online)

For several decades, the Three-Tier architecture, consisting of access, distribution (aggregation), and core layers, has been widely adopted as the standard model for data center network design (Ahmadi, 2021; Bairy, 2021; Janovic, 2023). This hierarchical model was primarily engineered to support north-south traffic patterns, where communication flows between external clients and internal servers. However, the widespread adoption of virtualization, containerization, and microservices has fundamentally shifted traffic patterns toward predominantly east-west communication, where server-to-server interactions dominate within the data center fabric (Frank & Beard, 2025; Maloo & Nikolov, 2022). Studies indicate that east-west traffic may account for approximately 70–80% of total data center traffic, thereby challenging the fundamental assumptions underlying the traditional Three-Tier design (Maloo & Nikolov, 2022). As a result, the architecture becomes less efficient due to increased latency, suboptimal routing paths, and potential bottlenecks.

One of the principal limitations of the Three-Tier architecture lies in its reliance on the Spanning Tree Protocol (STP) to prevent Layer 2 loops. While STP ensures network stability, it inherently disables redundant links, resulting in inefficient bandwidth utilization and limited scalability (Chiesa et al., 2021). Additionally, traditional network management depends heavily on manual configuration via command-line interfaces on a per-device basis, which introduces operational inefficiencies and increases the likelihood of human error (Maloo et al., 2023; Prasetyo, Prihandi, Rifqi, & Budiarto, 2024). This approach is particularly unsuitable for modern data centers that require rapid provisioning and dynamic configuration adjustments. Furthermore, the security model of the Three-Tier architecture is predominantly perimeter-based, assuming implicit trust within internal network segments. Such an approach fails to provide granular inspection of east-west traffic, thereby creating vulnerabilities that facilitate lateral movement in the event of a security breach (Mujib & Sari, 2020; Wonor, Musa, & Osazuwa, 2025; Xu et al., 2025).

In response to these structural limitations, Software-Defined Networking (SDN) has emerged as a transformative paradigm that separates the control plane from the data plane, enabling centralized network intelligence, programmability, and policy-driven management (Singh & Jha, 2017). SDN introduces a higher level of abstraction, allowing administrators to define network behavior through software-based policies rather than device-specific configurations. Among the various SDN implementations, Cisco Application Centric Infrastructure (ACI) represents one of the most mature and comprehensive solutions for enterprise data centers, integrating hardware components such as Cisco Nexus 9000 series switches with a centralized controller known as the Application Policy Infrastructure Controller (APIC) (Ahmadi, 2021).

Unlike conventional SDN approaches based on OpenFlow, Cisco ACI adopts a declarative and application-centric policy model in which administrators define connectivity and security requirements at the application level (Zilberman et al., 2015; Ajayi et al., 2026). These high-level policies are automatically translated into detailed configurations that are distributed across the network fabric. The spine-leaf topology employed by ACI ensures predictable latency, typically limited to two hops between endpoints, while also enabling multiple equal-cost paths that eliminate the need for STP and maximize bandwidth utilization (Ahmadi, 2021; Systems, 2024a). This architectural approach enhances scalability, resilience, and operational efficiency compared to traditional hierarchical designs.

Despite the increasing adoption of Cisco ACI in enterprise environments, existing literature provides limited comprehensive comparative analyses between Three-Tier and ACI architectures, particularly those grounded in real-world implementations. Prior studies tend to focus on isolated aspects such as performance (Hamidi Younessi, 2025), security through micro-segmentation [18], or automation efficiency (Prasetyo et al., 2024), without integrating these dimensions into a unified evaluation framework. Moreover, the discrepancy between vendor-recommended best practices and actual enterprise implementations is rarely examined

ISSN: 2716-3865 (Print), 2721-1290 (Online)

systematically, even though such analysis is essential for practical decision-making (Systems, 2024a)

Therefore, this study aims to address these gaps by conducting a comprehensive comparative evaluation of the Three-Tier architecture and Cisco ACI based on five key capability dimensions: scalability, security, reliability, operational simplicity and agility, as well as network visibility and virtualization awareness. The analysis is grounded in a real-world enterprise case study involving a data center that operates both architectures concurrently during an ongoing migration process (Harkiolakis et al., 2026; Palli, 2023; By combining empirical operational data with a structured analytical framework, this research seeks to contribute both academically and practically by providing evidence-based insights and actionable recommendations for network engineers, architects, and IT decision-makers engaged in data center transformation initiatives.

Methods

This study adopts a descriptive qualitative approach, implemented through a comparative case study combined with design evaluation techniques. This methodological choice is considered appropriate because the research focuses on analyzing and comparing two complex sociotechnical systems namely, data center network architectures based on their operational characteristics and design attributes rather than on statistical hypothesis testing (Creswell & Creswell, 2022). The comparative case study approach enables an in-depth and contextually grounded examination within a real-world enterprise environment, while the design evaluation perspective provides a structured analytical framework to assess architectural performance against predefined criteria (Systems, 2024a). The overall research process is organized into five interconnected stages, beginning with a systematic literature review and benchmarking to establish theoretical foundations and best practice references, followed by case description and empirical data collection, gap analysis between actual implementation and vendor-recommended design, comparative evaluation across key capability dimensions, and finally synthesis leading to practical recommendations. To enhance the robustness of the findings, data triangulation was employed by integrating documentary evidence, operational data, and observational insights (Systems, 2024b).

The empirical investigation was conducted at Data Center X (DC X), an enterprise-scale infrastructure operated by a major Islamic banking institution in Indonesia. This site was selected due to its unique operational condition in which both Three-Tier Legacy architecture and Cisco Application Centric Infrastructure (ACI) are simultaneously deployed across three geographically distributed locations, namely DC1, DC2, and a Disaster Recovery Center (DRC). Such a coexistence environment provides a rare opportunity for direct comparison under equivalent organizational, regulatory, and operational conditions. At the time of observation, the Three-Tier architecture consisted of approximately 65 switching nodes distributed across DC1 and DRC, supporting around 5,000 endpoints and managed through decentralized CLI-based configurations. In contrast, the Cisco ACI deployment comprised more than 100 leaf switches and six spine switches across three sites, supporting over 7,000 endpoints and centrally managed through a cluster of APIC controllers integrated via Cisco Nexus Dashboard Orchestrator in a multi-site architecture. The financial services context further introduces stringent requirements related to security compliance, availability, and operational reliability, thereby strengthening the analytical depth of the study. All institutional identifiers have been anonymized in accordance with a Non-Disclosure Agreement, and all observations were conducted with informed consent from relevant stakeholders without disrupting operational processes.

Data collection was carried out through a combination of primary and secondary sources to ensure both empirical depth and theoretical rigor. Primary data were obtained through four

ISSN: 2716-3865 (Print), 2721-1290 (Online)

main techniques. First, design documentation analysis was conducted on High-Level Design (HLD), Low-Level Design (LLD), topology diagrams, and ACI policy configurations, including Tenant, VRF, Bridge Domain, Endpoint Group, and Contract structures. Second, operational data were extracted from monitoring systems, covering metrics such as availability, CPU utilization, network latency, and packet loss across a three-month period from December 2025 to February 2026, involving 147 legacy nodes and 238 ACI nodes. Third, direct observation of network operations was performed through authorized sessions, enabling documentation of provisioning workflows, troubleshooting processes, and incident response practices. Fourth, informal expert consultations with network engineers and architects provided contextual insights into design decisions and operational trade-offs. Secondary data were used to establish benchmarks and theoretical grounding, including Cisco's official design guides, certification materials, technical whitepapers, and peer-reviewed academic literature from recognized databases, as well as industry standards such as TIA-942-B and Uptime Institute classifications (Al-Fares, Loukissas, & Vahdat, 2008; Benson, Akella, & Maltz, 2010; Kreutz et al., 2015).

The comparative evaluation in this study is structured around five key capability dimensions, namely scalability, security, reliability, operational simplicity and agility, and network visibility with virtualization awareness. These dimensions were identified through a combination of practitioner insights and literature review, where they consistently emerge as the principal criteria for evaluating data center network architectures (Haji et al., 2021; Hamidi Younessi, 2025; Singh & Jha, 2017). Each dimension is operationalized through specific indicators. Scalability refers to the ability of the architecture to accommodate growth in devices, endpoints, and traffic volume, including aspects such as modular expansion and bandwidth capacity. Security encompasses segmentation models, policy granularity, Zero Trust implementation, and resistance to lateral movement threats. Reliability is assessed through redundancy design, failover mechanisms, convergence time, and availability levels. Operational simplicity and agility relate to management models, automation capabilities, provisioning speed, and adaptability to continuous deployment environments. Network visibility and virtualization awareness involve monitoring depth, endpoint tracking, telemetry capabilities, and integration with virtualized or containerized environments.

The analytical process combines descriptive, comparative, and gap analysis techniques. Descriptive analysis is used to construct a detailed representation of each architecture, including topology, operational behavior, and policy structure. Comparative analysis is then conducted through systematic side-by-side evaluation based on the defined capability dimensions, supported by both qualitative observations and quantitative operational metrics such as availability percentages, CPU utilization, and latency measurements. Gap analysis is applied to identify discrepancies between the actual implementation of Cisco ACI at DC X and Cisco's recommended best practices, focusing on aspects such as tenant design, contract configuration, VRF placement, and multi-tenancy strategy. Comparative references from other institutions implementing ACI in alignment with best practices are used to contextualize these deviations. Finally, the results are synthesized to identify patterns, trade-offs, and strategic implications, which form the basis for developing practical recommendations aligned with established migration frameworks (Ahmadi, 2021; Systems, 2024a).

To ensure the trustworthiness of the findings, this study applies data triangulation by integrating multiple data sources and methods, thereby reducing potential bias and increasing validity. Researcher reflexivity is also acknowledged, particularly in relation to the investigator's professional background in network engineering, which may influence interpretation. Ethical considerations are strictly observed through compliance with Non-Disclosure Agreements, anonymization of sensitive information, and the use of informed consent for all contributing stakeholders.

ISSN: 2716-3865 (Print), 2721-1290 (Online)

Results and Discussion

Operational Performance Comparison

Operational data were collected from 385 monitored nodes, comprising 147 Three-Tier Legacy nodes and 238 Cisco ACI nodes across DC1, DC2, and the Disaster Recovery Center (DRC) from December 2025 to February 2026. The architectural description separately reports approximately 65 switching nodes in the Three-Tier environment and more than 106 nodes in the Cisco ACI environment. The main operational performance results are summarized in Table 1.

Table 1. Operational Performance Comparison between Three-Tier Legacy and Cisco ACI

Metric	Three-Tier Legacy	Cisco ACI	Difference	Superior
Endpoints Supported	~5,000 (2 sites)	~7,000+ (3 sites)	+40%	ACI
Switch/Node Count	~65 nodes	106+ nodes	+63%	ACI
Mean CPU Utilization	10.26%	3.33%	−68%	ACI
Mean Network Latency	7.21 ms	3.22 ms	−55%	ACI
Median Network Latency	3.50 ms	1.00 ms	−71%	ACI
ICMP Average Latency	4 ms	4 ms	None	Equal
ICMP Maximum Latency	13 ms	11 ms	−15%	ACI
Packet Loss	0%	0%	None	Equal
Traffic Distribution	Asymmetric (55–92%)	Balanced (3–5%)	—	ACI

Cisco ACI recorded substantially lower mean CPU utilization than the Three-Tier Legacy architecture, at 3.33% compared with 10.26%, representing an approximately 68% reduction. Two legacy nodes reached peak CPU utilization of approximately 92–99%, indicating substantially higher resource utilization within parts of the legacy environment.

Network latency measurements also favored Cisco ACI. Mean latency decreased from 7.21 ms in the legacy architecture to 3.22 ms in ACI, while median latency decreased from 3.50 ms to 1.00 ms. A separate ICMP test using 200 packets per endpoint showed an identical average latency of 4 ms for both architectures; however, the maximum latency was lower in ACI at 11 ms compared with 13 ms in the legacy environment. No packet loss was recorded for either architecture.

Traffic distribution showed a further difference between the two architectures. The ACI environment maintained load variance of approximately 3–5 percentage points across active uplinks, whereas the legacy environment showed considerably more asymmetric traffic distribution, with variance ranging from 55–92 percentage points across blocked and active paths.

Site-Level Availability Analysis

Availability was further examined at both aggregate and site levels because the ACI environment included DC2, a newly commissioned site that remained in an active stabilization phase during the observation period. The results are presented in Table 2.

ISSN: 2716-3865 (Print), 2721-1290 (Online)

Table 2. Aggregate and Site-Level Availability

Location	Three-Tier Legacy	Cisco ACI	Comparative Result
DC1	99.9992%	100.0000%	ACI slightly higher
DC2	Not reported	99.6489%	ACI site under stabilization
DRC	99.9998%	99.9994%	Approximately equal
Aggregate	99.9995%	99.8169%	Legacy numerically higher

At the aggregate level, the Three-Tier Legacy architecture recorded availability of 99.9995%, compared with 99.8169% for Cisco ACI. However, the ACI aggregate value was substantially affected by DC2, where 124 nodes were part of a newly commissioned site undergoing operational stabilization and recorded availability of 99.6489%.

Site-level comparison showed a different pattern. At DC1, Cisco ACI achieved 100.0000% availability compared with 99.9992% for the legacy architecture. At DRC, legacy and ACI availability were 99.9998% and 99.9994%, respectively, representing a difference of only 0.0004 percentage points. Thus, the observed aggregate availability difference was primarily associated with the lower availability recorded at the newly commissioned DC2 ACI site.

Gap Analysis: Implementation versus Recommended Design

The Cisco ACI implementation at DC X was evaluated against Cisco's recommended design approach and reference implementations at DC Y and DC Z. The analysis identified several differences between the implemented architecture and the recommended configuration, particularly in Tenant architecture, Contract configuration, and VMM Integration, as summarized in Table 3.

Table 3. Gap Analysis of Cisco ACI Implementation at DC X

Evaluation Aspect	Recommended Design	DC X Implementation	Identified Implication
Tenant Architecture	Separate Tenants according to business unit or application domain	Core objects, including L3Out configurations, Bridge Domains, route maps, and Contracts, are consolidated within the Common Tenant; EPGs are located in user-defined Tenants	Simplifies current operation but may constrain future multi-tenancy development
Contract Configuration	Contract configuration aligned with Tenant and application segmentation	Contract scope configured at the VRF level	Supports prevention of unintended cross-Tenant communication
VMM Integration	Integration with virtualization platforms for automated endpoint and policy management	VMM Integration is not activated	Limits real-time virtual endpoint tracking and hypervisor-integrated policy provisioning

The most substantial architectural deviation concerned Tenant design. DC X centralized core infrastructure objects, including L3Out configurations, Bridge Domains, route maps, and Contracts, within the Common Tenant, while Endpoint Groups were maintained in user-

ISSN: 2716-3865 (Print), 2721-1290 (Online)

defined Tenants such as *serverfarm*, *shared-service*, and *dmz*. This configuration was adopted to simplify operational management and reduce configuration navigation requirements.

Contract scope was configured at the VRF level rather than at the Global level. This configuration remained consistent with the operational requirement to prevent unintended communication across Tenants. However, the current Tenant structure may require architectural modification if the organization later expands the ACI environment to accommodate subsidiaries or third-party entities requiring stronger administrative and policy isolation.

A further gap concerned the non-activation of VMM Integration. DC X deliberately maintained separate operational responsibilities between the network engineering and server administration teams. As a result, native ACI integration with virtualization platforms was not implemented, limiting automatic port-group provisioning, real-time virtual endpoint tracking, and hypervisor-integrated policy management.

Comparative Evaluation across Five Key Capability Dimensions

The two architectures were subsequently compared across the five capability dimensions defined in the study: scalability, security, reliability, operational simplicity and agility, and network visibility and virtualization awareness. The comparative findings are summarized in Table 4.

Table 4. Comparative Evaluation across Five Key Capability Dimensions

Dimension	Three-Tier Legacy	Cisco ACI	Outcome	Margin
Scalability	STP-limited; cascade expansion; ~5,000 endpoints across 2 sites	Modular expansion; 7,000+ endpoints across 3 sites	ACI	High
Security	Perimeter-based; implicit permit; limited east-west visibility	EPG/Contract whitelist; micro-segmentation; Service Graph	ACI	High
Reliability	STP reconvergence; mature aggregate availability of 99.9995%	ECMP active-active; sub-second failover; DC1 availability of 100%	ACI*	Moderate
Operational Simplicity and Agility	Per-device CLI; manual configuration; no centralized API	APIC, NDO, NDI; REST API; intent-based automation	ACI	High
Visibility and Virtualization Awareness	Per-device SNMP; fragmented network view; no native VM awareness	Fabric-wide telemetry; endpoint tracking; health scoring	ACI	High

The numerical aggregate availability advantage of the legacy architecture was influenced by the stabilization of the newly commissioned DC2 ACI site.

Scalability

The Three-Tier architecture at DC X showed scalability constraints associated with its hierarchical structure and STP-based operation. Expansion of port capacity required additional Access and Aggregation switches and increased configuration complexity as the environment expanded. The legacy architecture supported approximately 5,000 endpoints across two sites.

ISSN: 2716-3865 (Print), 2721-1290 (Online)

In comparison, the Cisco ACI environment supported more than 7,000 endpoints across three sites and operated with substantially lower mean CPU utilization. The ACI spine-leaf architecture also allowed Leaf switches to be added for port expansion and Spine switches to be added for backbone capacity without restructuring the existing fabric. The observed scale of 106+ nodes and more than 7,000 endpoints demonstrates the larger operational capacity of the ACI environment.

Security

The legacy architecture employed a perimeter-oriented, VLAN-based segmentation model. Traffic within the same VLAN could traverse the access network without firewall inspection, resulting in limited control and visibility over east-west communication.

Cisco ACI applied an EPG and Contract-based model in which communication between application groups was governed by explicit policies. The architecture also supported micro-segmentation and Policy-Based Redirect for steering selected traffic through firewall or intrusion-prevention services. NDO enabled consistent propagation of policies across the three ACI sites. These characteristics provided more granular segmentation and policy control than the legacy architecture.

Reliability

The aggregate availability results numerically favored the Three-Tier Legacy environment. However, this result was substantially influenced by the newly commissioned DC2 ACI site. At operationally mature sites, the difference was minimal. ACI achieved 100.0000% availability at DC1 compared with 99.9992% for the legacy architecture, while the difference at DRC was only 0.0004 percentage points.

The architectures also differed in their redundancy mechanisms. The legacy environment relied on STP-based path management, whereas ACI used ECMP active-active forwarding. The ACI management and data planes were also separated through the APIC-based architecture, allowing forwarding operations to continue based on the existing programmed hardware state even when controller connectivity was unavailable.

Operational Simplicity and Agility

Management of the Three-Tier environment was performed primarily through individual CLI access across approximately 65 devices. Configuration changes therefore required device-level administration and created greater dependence on manual operational procedures.

Cisco ACI centralized network management through APIC, while NDO supported policy orchestration across multiple sites and NDI provided additional operational monitoring and assurance capabilities. The availability of REST APIs also enabled integration with automation platforms and Infrastructure-as-Code tools such as Ansible, Terraform, and Python. Therefore, the ACI environment provided a more centralized and automation-oriented management model than the legacy architecture.

Network Visibility and Virtualization Awareness

The Three-Tier Legacy environment relied primarily on SNMP polling, syslog, and device-level CLI commands. Consequently, operational information was distributed across individual network devices and required manual correlation.

Cisco ACI provided fabric-wide endpoint tracking, hierarchical health scoring, and integrated flow telemetry. Endpoint information, including MAC address, IP address, and Leaf attachment point, could be tracked across the fabric. However, the full virtualization-awareness capability of ACI was not utilized at DC X because VMM Integration had not been activated. The implementation therefore demonstrated substantially greater network visibility

ISSN: 2716-3865 (Print), 2721-1290 (Online)

than the legacy architecture, although its virtualization integration capabilities remained only partially utilized.

The findings of this study provide robust empirical support for the proposition that Cisco ACI constitutes a fundamentally more capable architecture than Three-Tier Legacy for modern enterprise data center environments across four of the five evaluated dimensions. The scalability, security, operational simplicity, and visibility advantages are not marginal improvements; they reflect qualitatively distinct design paradigms. The elimination of STP-imposed bandwidth and domain-size constraints, the transition from perimeter-centric to micro-segmented Zero Trust security, the shift from per-device CLI management to centralized intent-based policy, and the move from fragmented device-level monitoring to fabric-wide telemetry each represent substantive advancements addressing documented structural limitations of the Three-Tier model (Bairy, 2021; Hamidi Younessi, 2025).

The single dimension on which the legacy architecture retains a numerical advantage aggregate availability must be interpreted with considerable caution (Wang et al., 2025; Sebastian, 2025; Pepe et al., 2026). The 0.18 percentage point gap is predominantly driven by DC2, a newly commissioned ACI site still in active stabilization, whose inclusion in the aggregate figure depresses ACI's mean. Site-level comparisons at mature deployments demonstrate that ACI is comparable or superior in availability, and the structural superiority of ECMP active-active failover sub-second versus 1–50 seconds for STP reconvergence constitutes a more deterministic reliability model for mission-critical financial services workloads. These findings are consistent with the reliability analysis of Hamidi Younessi (2025) and the broader SDN reliability literature surveyed by Kreutz et al (2015).

The gap analysis at DC X illustrates a recurrent tension in enterprise SDN deployments between architectural purity and operational pragmatism. The Common Tenant object consolidation strategy is a defensible simplification for a single-organization context at current scale, but introduces constraints on future multi-tenancy evolution that architects should explicitly document and plan for. The non-activation of VMM Integration, while organizationally motivated, foregoes a significant portion of ACI's differentiated value real-time virtual endpoint tracking, hypervisor-integrated policy provisioning, and end-to-end overlay-to-underlay visibility. As Indonesian banking technology organizations evolve toward integrated network-cloud operations models (Osundare & Bolatito Ige, 2024), this represents a high-value, tractable optimization pathway.

The empirical findings of this study both corroborate and extend prior literature. The scalability advantages of spine-leaf topology over Three-Tier hierarchical design, as articulated by Bairy and Maru are confirmed by the DC X operational data showing that the ACI fabric supports 40% more endpoints and 63% more nodes at 68% lower CPU utilization. The security findings extend the micro-segmentation literature of Mujib & Sari (2020) and Wonor et al (2025) from experimental settings to a production financial services environment, contextualizing them within a comprehensive architectural comparison that illustrates the structural impossibility of equivalent micro-segmentation under VLAN-based Three-Tier design. The operational simplicity findings extend the automation efficiency evidence of Prasetyo et al (2024) by situating provisioning time reductions within a broader five-dimension evaluation framework and identifying proactive operational assurance capabilities NDI anomaly detection, pre-change analysis, firmware advisory scanning not examined in prior work. Collectively, this study contributes the first holistic, real-world enterprise case study evaluating all five capability dimensions simultaneously in a three-site production deployment with financial sector regulatory requirements (Aydin, 2026; Omar & Badar, 2026).

ISSN: 2716-3865 (Print), 2721-1290 (Online)

For practitioners, the findings carry several actionable implications. Network engineers adopting Cisco ACI should recognize that effective deployment requires a paradigm shift from device-centric to policy-centric thinking; organizations that retain CLI-centric habits on an ACI platform risk negating a substantial portion of its agility benefit. Investment in naming standards, object governance policies, and automation toolchain integration at the outset of deployment is strongly recommended. Network architects should design tenant structures with future multi-tenancy requirements in mind even when immediate organizational context does not demand it, and should treat VMM Integration activation as a near-term optimization priority rather than an optional capability. For IT decision-makers in similar Indonesian enterprise contexts, the study provides quantitative evidence 68% CPU reduction, 55% latency reduction, and elimination of STP-blocked bandwidth supporting the business case for migration, alongside security posture improvements directly relevant to Indonesian banking regulatory compliance requirements (Nasution, 2023; Makmur, 2024; Arifin et al., 2024).

Based on the operational insights derived from DC X's active migration process, a six-phase migration framework is proposed: (1) Assessment comprehensive inventory of workloads, VLANs, gateways, application dependencies, and L4–L7 service chains; (2) Fabric Build ACI spine-leaf commissioning, APIC cluster validation, and L3Out-based coexistence connectivity; (3) Incremental Workload Migration per-application-group migration with structured validation cycles; (4) Gateway Migration transfer of default gateway functions to ACI Border Leaf with careful attention to ARP behavior and forwarding continuity; (5) L4–L7 Integration and Policy Optimization Service Graph migration, Contract consolidation, and NDI-based compliance monitoring; (6) Finalization and Legacy Decommission documentation update, monitoring baseline refresh, and legacy infrastructure removal only after sustained traffic stability validation. Each phase should incorporate predefined rollback criteria to mitigate migration risk in production financial services environments.

Conclusion

This study presented a comprehensive comparative evaluation of Three-Tier Legacy and Cisco ACI architectures across five key capability dimensions, grounded in three months of operational data from 385 nodes distributed across three sites in an enterprise Islamic banking data center in Indonesia. The findings demonstrate that Cisco ACI outperforms Three-Tier Legacy in four of five dimensions with decisive advantage: scalability (40% more endpoints, 63% more nodes, 68% lower CPU utilization), security (micro-segmented EPG/Contract whitelist versus perimeter-only VLAN control), operational simplicity and agility (centralized APIC/NDO/NDI management and full API programmability versus per-device CLI), and network visibility and virtualization awareness (fabric-wide endpoint tracking and hierarchical health scoring versus fragmented per-device monitoring). On reliability, the legacy architecture's superior aggregate availability reflects operational maturity and a deployment maturity artifact rather than architectural design superiority; at site-equivalent mature comparisons, ACI is comparable or superior, and its ECMP active-active failover model provides structurally more deterministic convergence than STP-dependent reconvergence.

The gap analysis identified the Common Tenant consolidation strategy as a pragmatic deviation from Cisco's multi-tenant design recommendation defensible at current organizational scale but introducing constraints on future architectural evolution and the non-activation of VMM Integration as a consequential capability gap that should be addressed as operational team boundaries evolve. Taken together, the findings support the conclusion that migrating from Three-Tier to Cisco ACI constitutes a strategically sound infrastructure investment for enterprise data centers facing demands for increased scale, granular security enforcement, operational automation, and comprehensive network visibility. The six-phase

ISSN: 2716-3865 (Print), 2721-1290 (Online)

migration framework proposed in this study provides practitioners with a structured, empirically grounded roadmap for managing this transition in high-availability production environments.

Future research should extend the observation period to six to twelve months and incorporate multiple case study sites across different industry sectors to strengthen generalizability. Quantitative total cost of ownership (TCO) analysis, comparative evaluation of VMM Integration activation impact on operational efficiency, and investigation of ACI's performance characteristics under AI/ML and GPU-cluster east-west traffic at terabit scale represent important directions for subsequent investigation.

References

Ahmadi, A. (2021). *CCNP data center application centric infrastructure 300-620 DCACI official cert guide*. Cisco Press.

Ajayi, O. F., Kuyoro, A. O., Izang, A. A., Ayankoya, F. Y., Akande, O., Udosen, A. A., ... & Mensah, Y. A. (2026). A Graph Neural Network Framework for Automated Intent-to-Policy Translation and Enforcement in Software-Defined Networks. *Journal homepage: http://iieta. org/journals/isi*, *31*(2), 567-577.

Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. In *Proceedings of the ACM SIGCOMM 2008 conference on data communication* (pp. 63–74). Association for Computing Machinery. https://doi.org/10.1145/1402958.1402967

Arifin, R., Zulfa, E. A., Hanita, M., & Runturambi, A. J. S. (2024). Unveiling Indonesia's migration and border governance: Challenges and imperatives post-pandemic. *Social Sciences & Humanities Open*, *10*, 101202. https://doi.org/10.1016/j.ssaho.2024.101202

Aydin, K. (2026). Strategic Roadmap for Hydrogen Mobility in Türkiye: Implementation Phases, Export Potential, and Regional Hub Development. *New Energy Exploitation and Application*, *5*(2), 51-73. https://doi.org/10.54963/neea.v5i2.2350

Bairy, V. (2021). Spine and leaf architectures for next-generation data centers: Performance, scalability, and resilience. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5189881

Benson, T., Akella, A., & Maltz, D. A. (2010). Network traffic characteristics of data centers in the wild. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement* (pp. 267–280). Association for Computing Machinery. https://doi.org/10.1145/1879141.1879175

Chiesa, M., Kamisiński, A., Rak, J., Rétvári, G., & Schmid, S. (2021). A survey of fast-recovery mechanisms in packet-switched networks. *IEEE Communications Surveys & Tutorials*, *23*(2), 1253-1301.

Cisco. (2024). *Cisco Application Centric Infrastructure (ACI) design guide*. https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric-infrastructure-design-guide.html

Cisco. (n.d.). *Cisco Application Centric Infrastructure (ACI) operations & troubleshooting (DCACIO)*. Cisco U. https://u.cisco.com/paths/cisco-application-centric-infrastructure-operations-troubleshooting-174

Creswell, J. W., & Creswell, J. D. (2022). *Research design: Qualitative, quantitative, and mixed methods approaches* (6th ed.). SAGE Publications.

ISSN: 2716-3865 (Print), 2721-1290 (Online)

Haji, S. H., Zeebaree, S. R. M., Saeed, R. H., Ameen, S. Y., Shukur, H. M., Omar, N., Sadeeq, M. A. M., Ageed, Z. S., Ibrahim, I. M., & Yasin, H. M. (2021). Comparison of software defined networking with traditional networking. *Asian Journal of Research in Computer Science, 9*(2), 1–18. https://doi.org/10.9734/AJRCOS/2021/v9i230216

Hamidi Younessi, A. (2025). *Evaluating the advantages of Cisco ACI over traditional three-tier architecture* [Bachelor's thesis, Häme University of Applied Sciences]. Theseus.

Harkiolakis, N., Dasgupta, A., Ameya, A., Aneja, A., Agarwal, K., Bhardwaj, M., ... & Senthilkumar, H. (2026). A case study on integrating real-world data infrastructure development into an undergraduate data science curriculum through internships. *International Journal of Teaching and Case Studies*, *16*(3), 226-243. https://doi.org/10.1504/IJTCS.2026.155418

Janovic, J. (2023). *Cisco ACI: Zero to hero: A comprehensive guide to Cisco ACI design, implementation, operation, and troubleshooting*. Apress. https://doi.org/10.1007/978-1-4842-8838-2

Knight Frank. (2025). *Global data centres report 2025*. Knight Frank.

Kreutz, D., Ramos, F. M. V., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE, 103*(1), 14–76. https://doi.org/10.1109/JPROC.2014.2371999

Makmur, K. L. (2024). Why only scrutinise formal finance? Money laundering and informal remittance regulations in Indonesia. *Journal of Economic Criminology*, *6*, 100111. https://doi.org/10.1016/j.jeconc.2024.100111

Maloo, S., & Nikolov, I. (2022). *Cisco data center fundamentals*. Cisco Press.

Maloo, S., Nikolov, I., & Ahmed, F. (2023). *CCNP and CCIE data center core DCCOR 350-601 official cert guide* (2nd ed.). Cisco Press.

Mujib, M., & Sari, R. F. (2020). Performance evaluation of data center network with network micro-segmentation. In *2020 12th International Conference on Information Technology and Electrical Engineering (ICITEE)* (pp. 27–32). IEEE. https://doi.org/10.1109/ICITEE49829.2020.9271749

Nasution, E. R. (2023). The role of regulatory authority in maintaining financial security in the Indonesian banking sector: A legal framework analysis. *The International Journal of Politics and Sociology Research*, *11*(3), 386-397. https://doi.org/10.35335/ijopsor.v11i3.209

Nunes, B. A. A., Mendonca, M., Nguyen, X.-N., Obraczka, K., & Turletti, T. (2014). A survey of software-defined networking: Past, present, and future of programmable networks. *IEEE Communications Surveys & Tutorials, 16*(3), 1617–1634. https://doi.org/10.1109/SURV.2014.012214.00180

Omar, M. N., & Badar, M. I. M. A. (2026). An Islamic Juristic Framework for Data and Algorithmic Monetisation in the Digital Economy. *International Journal of Islamic Finance and Sustainable Development*, *18*(2), 107-127. https://doi.org/10.55188/ijifsd.v18i2.1255

Osundare, O. S., & Ige, A. B. (2024). Transforming financial data centers for fintech: Implementing Cisco ACI in modern infrastructure. *Computer Science & IT Research Journal, 5*(8), 1806–1816. https://doi.org/10.51594/csitrj.v5i8.1397

Palli, S. S. (2023). Real-time data integration architectures for operational business intelligence in global enterprises. *International Journal of Scientific Research in*

ISSN: 2716-3865 (Print), 2721-1290 (Online)

Computer Science, Engineering and Information Technology (IJSRCSEIT), 9 (1), 361-371.

Pepe, M., Crisan, A., Maravelakis, E., Palumbo, D., Dewedar, A. K. H., & Klapa, P. (2026, March). Mapping 3D Digital Heritage at Scale: A ChatGPT-Assisted Analysis of Sketchfab's "Cultural Heritage & History" Models. In *Informatics* (Vol. 13, No. 3, p. 36). MDPI. https://doi.org/10.3390/informatics13030036

Prasetyo, L., Prihandi, I., Rifqi, M., & Budiarto, R. (2024). Implementation of automation configuration of enterprise networks as software defined network. *Computer Science and Information Technologies, 5*(2), 99–111. https://doi.org/10.11591/csit.v5i2.p99-111

Rout, S., Sahoo, K. S., Patra, S. S., Sahoo, B., & Puthal, D. (2021). Energy efficiency in software defined networking: A survey. *SN Computer Science, 2*(4), Article 308. https://doi.org/10.1007/s42979-021-00659-9

Sebastian, D. (2025). Modernizing Credit Risk with Data Mesh: A Large Bank's Transformation to Real-Time Credit Intelligence. *Journal of Computer Science and Technology Studies*, *7*(10), 650-664.

Singh, S., & Jha, R. K. (2017). A survey on software defined networking: Architecture for next generation network. *Journal of Network and Systems Management, 25*(2), 321–374. https://doi.org/10.1007/s10922-016-9393-9

Wang, Y., Jin, C., Wang, T., & Xu, D. (2025). The interpretation of historical layer evolution laws in historic districts from the perspective of the historic urban landscape: a case study in Shenyang, China. *Land*, *14*(5), 1029. https://doi.org/10.3390/land14051029

Wonor, I. A., Musa, M. O., & Osazuwa, C. M. (2025). Zero trust and micro-segmentation: Strengthening network security. *The American Journal of Management and Economics Innovations, 7*(10), 56–70. https://doi.org/10.37547/tajmei/Volume07Issue10-05

Xu, D., Gondal, I., Yi, X., Susnjak, T., Watters, P., & McIntosh, T. R. (2025). The erosion of cybersecurity zero-trust principles through generative ai: A survey on the challenges and future directions. *Journal of Cybersecurity and Privacy*, *5*(4), 87. https://doi.org/10.3390/jcp5040087

Zilberman, N., Watts, P. M., Rotsos, C., & Moore, A. W. (2015). Reconfigurable network systems and software-defined networking. *Proceedings of the IEEE*, *103*(7), 1102-1124. https://doi.org/10.1109/JPROC.2015.2435732

ISSN: 2716-3865 (Print), 2721-1290 (Online)