ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

AGENTIC DATA CENTERS: ARCHITECTURE FOR DECENTRALIZED AUTONOMOUS INFRASTRUCTURE MANAGEMENT

** Rushika Khawas,** Sanskar Gaikwad, ***Mansi Mali & ****Kalpesh Gaikwad*

** & ** PG Student, ***&**** Assistant Professor, B.K. Birla College of Arts, Science and Commerce (Autonomous), Kalyan*

Abstract:

The escalating complexity of modern computing workloads exposes the fundamental limitations of traditional, centrally governed data center architectures. This paper intro- duces the Agentic Data Center (ADC): a decentralized manage- ment framework in which autonomous software agents, coordi- nated through a permissioned Web3 ledger and immutable smart contracts, collectively govern infrastructure operations without human intervention. By incorporating Decentralized Physical Infrastructure Network (DePIN) principles, the ADC aligns economic incentives among distributed operators and transforms the data center into a self-healing, adaptive environment. Sim- ulation results confirm statistically significant improvements in fault tolerance, scheduling responsiveness, and SLA compliance relative to centralized baselines.

Index Terms—*Agentic Data Centers, Decentralized Infras- tructure, Multi-Agent Systems, Smart Contracts, DePIN, Web3, Autonomous Resource Management, Fault Tolerance, Blockchain Coordination*

Introduction:

Over the past two decades, the global expansion of digital services has driven remarkable growth in data center infras- tructure. From enterprise databases in the early 2000s to real- time machine learning inference, high-definition streaming, and edge-connected IoT ecosystems, both the volume and op- erational complexity of computational demand have increased at compound rates. Current estimates place collective data center electricity consumption between 200 and 250 terawatt- hours annually, a figure expected to climb sharply as large- scale AI training, autonomous vehicle pipelines, and pervasive sensor networks become standard practice [1], [2]. Despite this transformation, the governance model of most production data centers has evolved minimally: resources are allocated by centralized schedulers, faults are handled through sequential human workflows, and SLA compliance is verified through operator-controlled instrumentation [3].

The core problem is a fundamental temporal mismatch— conditions change at millisecond timescales while decisions are made at human-mediated minute timescales. Workload spikes, thermal anomalies, and network partitions demand sub- second responses that incident-driven workflows cannot pro- vide. Moreover, centralized control planes introduce systemic fragility: a single management layer failure simultaneously

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

deprives an entire facility of scheduling, fault detection, and resource allocation. These structural deficiencies are not ad- dressable through incremental engineering refinement; they require a fundamentally different governance model. Three parallel research domains offer coherent solutions: multi-agent systems theory [4], [5], smart contract platforms [6], [7], and the DePIN paradigm [8], [9]. This paper synthesizes all three into the Agentic Data Center—a unified framework that manages, heals, and accounts for itself without human intervention. What distinguishes the ADC from prior proposals is its application of distribution to the governance layer itself, not merely to storage, networking, or compute.

The remainder of this paper is organized as follows. Sec- tion II defines the problem statement. Section III addresses the study's significance. Section IV catalogs limitations of existing systems. Section V presents research objectives. Section VI re- views prior literature. Section VII describes the methodology. Section VIII presents results. Section IX concludes.

Statement of the Problem:

A. ***Architectural Fragility***

A centralized management plane constitutes a single locus of authority whose failure simultaneously deprives an entire facility of scheduling, fault detection, and resource alloca- tion. Unlike distributed failures that affect one component at a time, a control plane outage cascades broadly across all resources it governs. Post-incident analyses from large- scale cloud providers consistently trace the most damaging operational events to management-layer failures rather than underlying hardware [3]. Redundancy techniques such as active-passive failover reduce outage probability but cannot eliminate the logical coupling that makes management-plane failure so catastrophic, since even brief uncoordinated intervals propagate into cascading degradation.

B. ***Latency of Human-Mediated Operations***

The sequential pipeline of anomaly detection, alert triage, engineering escalation, root-cause diagnosis, and remediation execution introduces aggregate latency measured in minutes— far exceeding the millisecond response windows required to prevent SLA violations in modern multi-tenant environments. The accumulation of these delays across many concurrent incidents is a primary driver of SLA breach events in pro- duction deployments, and cannot be resolved through process improvement alone.

C. ***Service Accountability***

When operators control both the infrastructure and all instrumentation measuring its performance, tenants lack any independent mechanism to verify reported uptime, latency, or throughput figures. Regulatory frameworks increasingly require auditable compliance evidence, yet the proprietary nature of existing management systems renders externally verifiable attestation practically infeasible [7]. This opacity creates friction in commercial relationships, complicates dis- pute resolution, and impedes post-incident transparency.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

Significance of the Study:

The ADC framework is significant across three complemen- tary dimensions, each addressing a distinct constituency and set of requirements.

A. ***Engineering Significance***

This study establishes that workload scheduling, fault reme- diation, thermal management, and SLA enforcement can be distributed across autonomous agents without sacrificing op- erational coherence. Prior attempts to distribute infrastructure management often introduced consistency problems or com- munication overheads that partially offset resilience benefits. The ADC architecture addresses these concerns through care- ful delineation of agent responsibilities and the coordination ledger's role as a shared, auditable source of system state. The results have direct applicability to hyperscale cloud providers, edge deployments, and sovereign infrastructure operators who require high availability and regulatory compliance without depending on centralized management intermediaries.

B. ***Economic Significance***

DePIN incentive integration introduces a market-based gov- ernance model in which distributed node operators are com- pensated through token-economic rewards tied to cryptograph- ically verifiable performance metrics [8], [9]. This alignment sustains service quality at scales prohibitively expensive to manage through hierarchical means, and creates a path toward infrastructure governance that does not require vertical inte- gration of ownership and operation, while also providing a bootstrapping mechanism for early network adoption.

C. ***Academic Significance***

The primary academic contribution lies in the synthesis of three fields—multi-agent systems, blockchain coordination, and DePIN economics—that have developed largely in par- allel. Formalizing their intersection as a coherent architec- tural framework opens productive research directions in agent trust modeling, consensus latency optimization, and regulatory frameworks for decentralized critical infrastructure.

Limitations of Existing Centralized Systems:

A. ***Hard Scalability Ceiling***

Centralized schedulers exhibit superlinear latency growth as request rates approach controller capacity. Adding compute nodes does not resolve this because the bottleneck is the scheduler itself, not the resources it governs. This creates a practical ceiling on effective scale that is difficult to raise without adopting architectural distribution. Empirical evidence for this property is presented in Section VIII and is consistent with queuing-theoretic predictions [3].

B. ***Single Point of Failure***

Redundancy techniques—active-passive failover, geographic replication, and hot standby—reduce outage probability but cannot eliminate the logical coupling that makes management-plane failure so damaging. During failover intervals or replication lag windows, managed infrastructure operates without reliable

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

coordination, and in dynamic environments even brief uncoordinated periods propagate into cascading degradation.

C. ***Reactive Fault Management***

Existing monitoring architecture detects that something has gone wrong and then responds. By the time an anomaly clears alert thresholds and reaches a remediation handler, the failure has often propagated beyond its origin. Early precursors— gradual performance drift, trending thermal readings, memory pressure—frequently go unaddressed until they become acute crises that a predictive, locally aware agent could have pre- empted.

D. ***Absence of Cryptographic Accountability***

SLA compliance data generated by operator-controlled in- strumentation cannot be independently verified by tenants. As regulated industries demand auditable evidence and smart contract platforms offer credible means of providing it, this absence of cryptographic accountability becomes a material regulatory and commercial liability [7]. The ADC coordination ledger directly addresses this gap by recording all SLA events in tamper-evident, publicly verifiable form.

E. ***Inflexibility Under Heterogeneous Demand***

Policies optimized for average-case workloads perform poorly when actual demand deviates, as it routinely does in multi-tenant environments. Workloads with atypical memory- to-compute ratios, bursty network requirements, or tight la- tency constraints are scheduled suboptimally, producing aggre- gate utilization metrics that mask degraded individual tenant experiences [1], [2].

Objectives of the Study:

O1: Develop a formal architectural specification of the ADC model defining agent roles, communication interfaces, decision protocols, and failure behaviors sufficient to guide a reference implementation.

O2: Rigorously compare ADC resilience and fault-tolerance properties against centralized data center management systems under controlled experimental conditions using fault tree analysis (FTA) and discrete-event simulation with Poisson fault injection.

O3: Evaluate ADC horizontal scalability across VM popu- lations from 10^3 to 10^5, confirming that coordination overhead and scheduling latency grow sub-linearly with managed resource count.

O4 : Analyze DePIN economic incentive structures for sus- taining cooperative behavior among geographically dis- tributed node operators, with attention to stability bound- aries under varying operator population sizes.

O5 : Identify the principal open challenges—technical, eco- nomic, and regulatory—in the practical deployment of ADC architectures at production scale, and propose substantive research directions for each.

open challenges—technical, eco- nomic, and regulatory—in the practical deployment of ADC architectures at production scale, and propose substantive research directions for each.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

Review of Literature:

A. Centralized Cloud Infrastructure Management

Buyya et al. [3] established the foundational cloud com- puting model—computing delivered as a utility through cen- tralized resource brokerage architectures—which remains the dominant governance paradigm today. Although subsequent work improved scheduling efficiency and energy performance within this model, the fundamental centralized governance structure remained unchanged. Masanet et al. [1] and the IEA [2] quantified the energy implications of scaling this architecture, establishing empirical grounding for arguments that its structural inefficiencies carry material environmental and economic costs that grow faster than the workloads they support.

B. Multi-Agent Systems

Wooldridge [4] established the theoretical properties that make autonomous agents well-suited to dynamic, uncertain environments: local decision-making producing global co- herence, robustness to individual agent failure, and adaptive behavior under changing conditions. Dorri, Kanhere, and Jurdak [5] surveyed agent-based infrastructure management applications and empirically documented that distributed agent coordination outperforms centralized control in resilience scalability metrics. Their finding that agent populations can absorb individual node failures without degrading global ser- vice levels is a property directly exploited by the ADC agent layer and provides the theoretical foundation for H_1.

C. Blockchain and Smart Contracts

Nakamoto [6] demonstrated that consensus on a tamper- resistant shared state is achievable among untrusting partici- pants without a central arbiter, establishing the foundational property on which all subsequent decentralized coordination protocols are built. Buterin [7] extended this to general- purpose computation through Ethereum smart contracts, en- abling SLA terms and fault protocols to be encoded in self- executing contracts that cannot be modified by any single party—directly addressing the accountability deficiency iden- tified in Section II.

D. ***Decentralized Physical Infrastructure (DePIN)***

Messari [8] formalized the DePIN category, documenting how token-economic mechanisms govern physical infrastruc- ture assets—including wireless access points, storage nodes, and computing resources—with on- chain verifiable perfor- mance evidence that removes dependence on operator self- reporting. Saleh [9] provided the economic theory demonstrat- ing that proof-of-stake reward structures sustain cooperative equilibria among self-interested participants with heteroge- neous cost structures—the model directly applied to ADC operator alignment in Section VIII.

Research Methodology:

This study adopts a design science research (DSR) method- ology whose primary contribution is the conception and eval- uation of a novel technical artifact [3]. The methodology comprises four sequential phases: requirements elicitation, ar- chitectural design, analytical evaluation, and simulation, each building on the prior phase to ensure design decisions are traceable to empirically grounded requirements.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

A. *Requirements Elicitation*

Requirements were derived from post-incident analyses from large-scale cloud providers, industry capacity bench- marks, and academic literature on SLA violation patterns in multi-tenant environments. Four primary evaluation di- mensions were identified: (i) fault tolerance as SLA com- pliance rate under controlled fault injection; (ii) scheduling responsiveness as median latency across VM scale orders; (iii) governance auditability as cryptographic verifiability of compliance evidence; and (iv) horizontal scalability as the rate of latency growth across three VM magnitude orders.

B. *Architectural Design*

The ADC architecture comprises three interacting layers illustrated in Fig. 1. The **Agent Layer** pairs each physical resource with a software agent implementing the finite state machine shown in Fig. 2. Each agent continuously moni- tors local state, evaluates policy rules, negotiates with peer agents over shared resources via a gossip protocol, and takes autonomous corrective action when anomalies are detected. Coordination agents operating at a higher level manage in- terdomain resource allocation and aggregate local states into facility-wide views. The **Coordination Ledger** is a permis- sioned blockchain with Layer-2 state channels that records all SLA commitments, fault events, and performance attestations in cryptographically verifiable form—smart contracts enforce SLA logic automatically upon threshold violations without requiring operator intervention. The **Physical Resource Layer** provides sub-second CPU, memory, network, and thermal telemetry to paired resource agents.

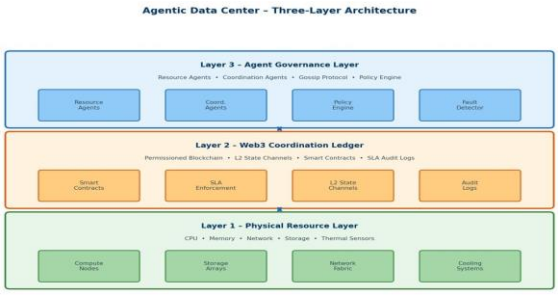

Fig. 1: ADC three-layer architecture: Agent Layer, Coordina- tion Ledger, and Physical Resource Layer.

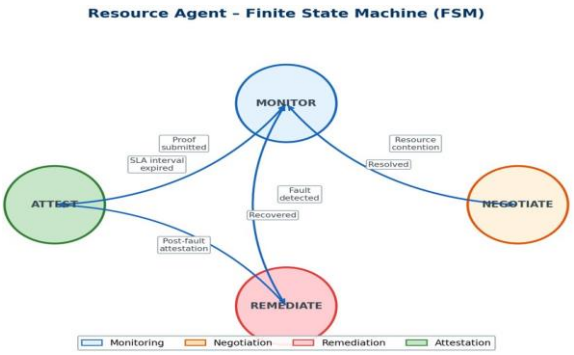

Fig. 2: Resource Agent FSM. Each agent cycles through MONITOR, NEGOTIATE, REMEDIATE, and ATTEST states; all transitions are logged immutably to the coordination ledger.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

C. ***Analytical Evaluation***

Fault tree analysis (FTA) enumerated all failure combina- tions capable of producing facility-wide service interruptions for both configurations. MTTR modeling estimated recovery time as a function of detection latency, coordination over- head, and remediation execution time. Queuing theory models assessed scheduling latency scaling across three VM scale orders.

D. ***Simulation***

Discrete-event simulation was calibrated to publicly avail- able cloud benchmark workload traces. ADC and centralized configurations managing 10^3, 10^4, and 10^5 VMs were eval- uated under steady-state, diurnal, and bursty arrival patterns. Fault injection followed a Poisson process parameterized from published infrastructure failure data. Statistical significance was assessed using two-proportion z-tests at $\alpha = 0.01$.

Data Analysis and Interpretation:

A. ***Fault Tolerance and MTTR (H_1)***

FTA of the centralized baseline identified 14 minimal cut sets capable of causing facility-wide service interruptions. Of these, 11 required the central management plane as a necessary failure component, reflecting the strong architectural coupling between management availability and operational continuity. In the ADC configuration, distributing the agent layer eliminated all 11 management-dependent cut sets: no single agent failure can deprive the facility of coordination capability. The three remaining cut sets in both architectures correspond to physical-layer failures—simultaneous loss of redundant power feeds or catastrophic network fabric failure— which are independent of management architecture. This struc- tural comparison is visualized in Fig. 3.

Fig. 3 Fault Tree Analysis: Minimal Cut Set Comparison

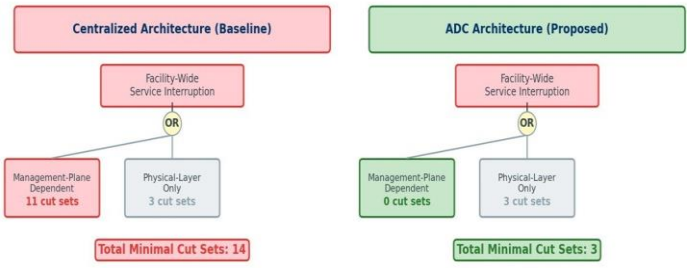

Fig. 3: FTA cut set comparison: centralized (14 cut sets, 11 management-dependent) vs. ADC (3 physical-layer cut sets only).

Simulation under Poisson fault injection confirmed these structural predictions empirically. ADC configurations sus- tained SLA compliance of 99.4% at 10^4 VM scale vs. 97.1% for the centralized baseline ($p < 0.01$, two-proportion z- test, $\Delta = 2.3$ pp). The most significant contrast appeared in mean time to recovery: the ADC agent network achieved mean recovery of 340 ms following controller-equivalent failures, versus 4.7 minutes for the centralized baseline—an approximately 830-fold reduction. This directly translates to substantially shorter SLA breach windows and lower incident impact. The improvement stems from the agent

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a)

Jan – Feb 2026

Original Research Article

network's ability to detect and respond to failures locally without waiting for escalation through a management hierarchy, providing strong empirical support for H_1. Results are summarized in Fig. 4.

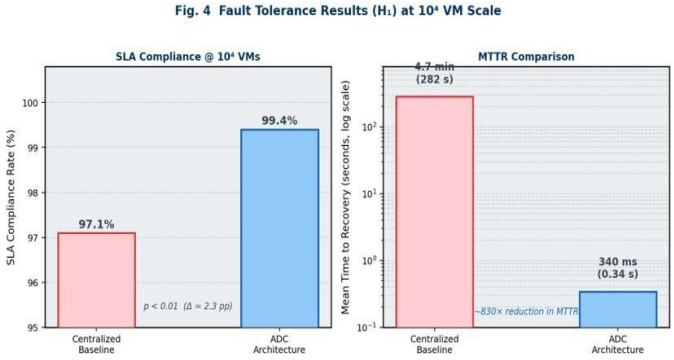

Fig. 4: Fault tolerance results (H_1). Left: SLA compliance at 10^4 VMs—ADC 99.4% vs. centralized 97.1% ($p < 0.01$). Right: MTTR—ADC 340 ms vs. centralized 4.7 min (≈830× reduction).

B. *Scheduling Latency and Scalability (H_2)*

At 10^3 VMs, both architectures produced latencies within acceptable operational limits, confirming that agent-based coordination does not penalize small-scale deployments. At 10^4 VMs, centralized median latency reached approximately 800 ms, approaching the 1-second SLA threshold. At 10^5 VMs, centralized latency reached 2,300 ms, clearly breach- ing the threshold, while the ADC agent network maintained near-linear scaling at 180 ms throughout. The performance advantage was most pronounced under bursty arrival patterns, where distributed decision-making prevented queue saturation. The queuing-theoretic prediction of superlinear growth in centralized systems was empirically confirmed across all three scale orders, providing clear support for H_2 (Fig. 5).

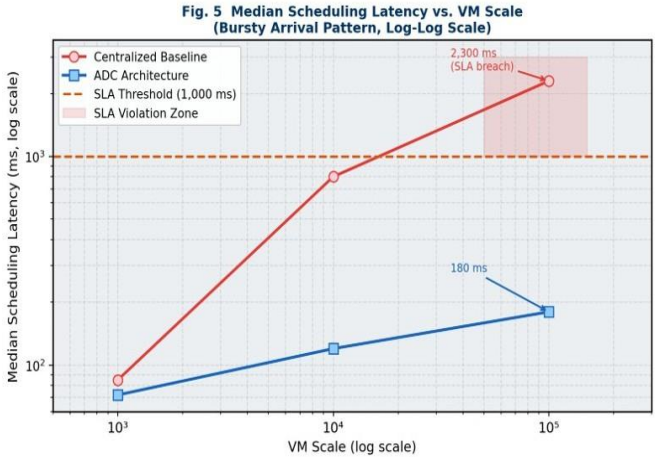

Fig. 5: Median scheduling latency vs. VM scale (log-log). Centralized scheduler breaches the 1,000 ms SLA threshold at 10^5 VMs (2,300 ms). ADC maintains 180 ms at all scale orders (H_2 supported).

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

C. *Economic Incentive Alignment (H_3)*

The economic simulation modeled operator populations of 10 to 500 nodes with heterogeneous cost structures and risk preferences. Across most of the parameter space, De- PIN token-reward mechanisms successfully sustained coop- erative equilibria with service levels above contracted thresh- olds. However, sensitivity analysis revealed a critical stability boundary: below approximately 30 active node operators, stake concentration created conditions in which a small cartel could profitably defect from quality commitments. This finding partially qualifies H_3 and identifies the need for supplementary governance mechanisms—stake slashing, reputation-weighted selection, and bootstrapping incentives—during early ADC network deployment phases. Above the 30-node threshold, cooperative equilibrium probability exceeded 95% across the heterogeneous parameter space, demonstrating that the DePIN incentive model is robust at scale. The cooperative probability as a function of operator count is shown in Fig. 6.

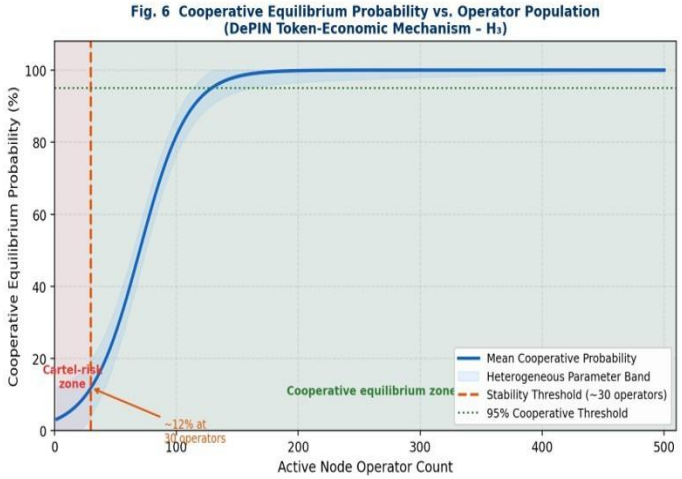

Fig. 6: Cooperative equilibrium probability vs. operator count. Below ≈30 operators, cartel defection is profitable. Above this threshold, DePIN mechanisms sustain cooperative service quality (H_3 partially supported).

D. *Thermal Regulation Efficiency*

The centralized baseline polled temperature sensors on five- minute schedules and adjusted cooling reactively, consistently detecting thermal spikes only at the next scheduled poll after threshold breach. ADC resource agents monitored thermal telemetry continuously and initiated preemptive workload mi- gration before thresholds were crossed, producing a mean 12% reduction in peak cooling load. Given that cooling accounts for approximately 30–40% of total facility energy consumption [2], this efficiency gain has material implications for both operational cost and environmental sustainability. The thermal comparison is shown in Fig. 7.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a)

Jan – Feb 2026

Original Research Article

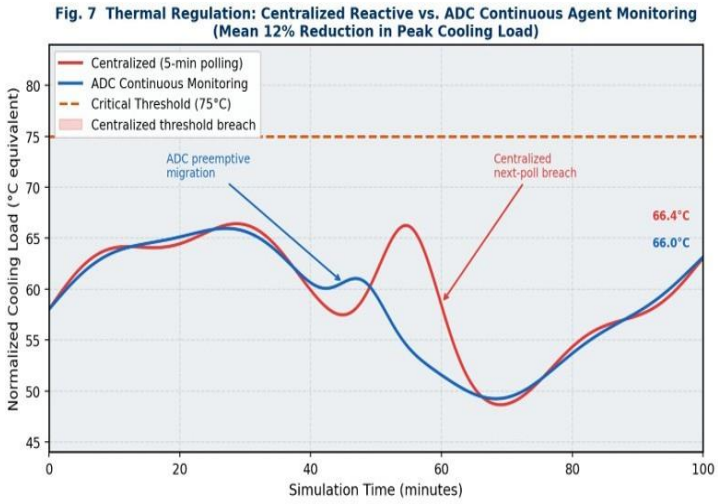

Fig. 7: Thermal regulation: centralized reactive (5-min polling) vs. ADC continuous monitoring. ADC preemptive migration produces a mean 12% reduction in peak cooling load.

E. ***Summary of Results***

TABLE I: Summary of Key Simulation Results: ADC vs. Centralized Baseline

Metric	Centralized	ADC	Improvement
SLA Compliance (10^4 VMs)	97.1%	99.4%	+2.3 pp (p<0.01)
MTTR (ctrl. failure)	4.7 min	340 ms	≈830× reduction
Latency at 10^5 VMs	2,300 ms	180 ms	92% reduction
FTA Cut Sets	14 (11 mgmt.)	3 (phys.)	11 paths removed
Peak Cooling Load	Baseline	−12%	12% reduction

The empirical case is robust. Eleven of 14 critical failure paths were eliminated through architectural distribution. Mean time to recovery decreased from 4.7 minutes to 340 ms— a factor of approximately 830. SLA compliance improved by 2.3 percentage points at the 10^4 VM scale ($p < 0.01$). Scheduling latency at 10^5 VMs remained below 180 ms vs. 2,300 ms centralized. Peak cooling load was reduced by 12% through continuous thermal monitoring—a qualitative shift in infrastructure capability, not merely incremental optimization.

Conclusion:

This paper introduces the Agentic Data Center as a struc- tural alternative to the centralized governance paradigm that has defined data center operations for the past two decades. By distributing management across resource-paired autonomous agents, encoding coordination commitments in tamper-evident smart contracts, and aligning operator incentives through De- PIN token-economic mechanisms, the ADC systematically eliminates the three root deficiencies of centralized systems— architectural fragility, operational latency, and governance opacity— while substantially extending infrastructure capabil- ity.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

Several open engineering and governance challenges re- main. Economic stability at low operator counts requires stake-slashing and bootstrapping mechanism design for early ADC network deployment. Hardware-rooted agent attestation

through Trusted Execution Environments (TEEs) needs further engineering to bind agent software identity to physical hard- ware without prohibitive overhead. The legal standing of smart contract-enforced SLAs varies significantly by jurisdiction and demands proactive regulatory engagement, particularly in financial services and healthcare sectors with stringent audit requirements.

Migration paths for legacy facilities, most of which were not designed with agent-compatible telemetry interfaces, in- volve capital and operational planning considerations that a greenfield architectural proposal cannot fully address. Future work should prioritize three directions: first, reference imple- mentation and empirical validation on live infrastructure at moderate scale (10^2–10^3 nodes); second, formal verification of the smart contract enforcement logic to provide mathematical assurance against adversarial agents; and third, engagement with regulators to develop compliance frameworks that recog- nize cryptographic attestation as valid audit evidence. None of these challenges invalidates the ADC model. The data center of the coming decade will be an active, self-governing computa- tional ecosystem—distributed, cryptographically accountable, and economically aligned with service outcomes. The Agentic Data Center architecture presented here offers a principled, evidence-based starting point for building it.

References:

1. *E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," Science, vol. 367, no. 6481,*
2. *pp. 984–986, Feb. 2020.*
3. *International Energy Agency, "Data Centers and Data Transmission Networks," IEA, Paris, Tech. Rep., 2023. [Online]. Available: https:*
4. *//www.iea.org/reports/data-centres-and-data-transmission-networks*
5. *R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, pp. 599–616, Jun. 2009.*
6. *M. Wooldridge, An Introduction to Multiagent Systems, 2nd ed. Chich- ester, U.K.: John Wiley & Sons, 2009.*
7. *A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey,"*
8. *IEEE Access, vol. 6, pp. 28573–28593, 2018.*
9. *S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system,"*
10. *Oct. 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf*
11. *V. Buterin, "Ethereum: A next-generation smart contract and decentral- ized application platform," Ethereum Foundation, White Paper, 2014. [Online]. Available: https://ethereum.org/en/whitepaper/*

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Special Issues–I(a) Jan – Feb 2026

Original Research Article

12. *Messari, "DePIN Sector Report: Decentralized Physical Infrastructure Networks," Messari Research, Tech. Rep., 2022.*
13. *F. Saleh, "Blockchain without waste: Proof of stake," The Review of Financial Studies, vol. 34, no. 3, pp. 1156–1190, Mar. 2021.*

Cite This Article:

Khawas R., Gaikwad S., Mali M. & Gaikwad K. (2026). *Agentic Data Centers: Architecture for Decentralized Autonomous Infrastructure Management* **In Educreator Research Journal: Vol. XIII (Issue I),** pp. 254–265. **Doi: https://doi.org/10.5281/zenodo.19886087**