

Architectural Optimization Of A Lightweight Kubernetes Microservices Framework For The Toba Herbal Information System Under Extreme Resource Constraints

Pandapotan Siagian¹, Sindak Hutauruk², Frengki Simatupang¹, Eka Sthepani¹, Marojahan Sigirot¹, Gerry Wolowing¹, Istas Manalu¹, Sari Mutia¹, Febrian Winston Hutagalung¹

¹*Department of Computer Engineering, Institut Teknologi Del, 22381 Kab.Toba, Sumatera Utara, Indonesia*

²*Department of Electrical Engineering, Univeritas HKBP Nommensen, 20235 Medan, Sumatera Utara, Indonesia*

**siagian.p@gmail.com*

ABSTRACT

This study designs, implements, and evaluates a lightweight, decoupled virtualization architecture utilizing K3s Kubernetes and Docker containers tailored for the Toba Herbal Information (THI). The system architecture isolates the web application logic, stateful relational databases, and high-velocity user connection sessions into three independent, resource-constrained pod tiers. Operating under a strict single-node hardware envelope of 4 CPU cores and 2 GB RAM, the framework enforces rigorous memory quotas and integrates an in-memory Redis caching buffer. Performance resiliency was rigorously evaluated via Apache JMeter against a baseline monolithic virtual machine (VM) utilizing a peak load of 1,000 concurrent healthcare users, complemented by an intentional Chaos Engineering failure injection. Empirical benchmarks demonstrate that the proposed 3-pod K3s framework successfully accommodated peak traffic with a stable 95th percentile (P₉₅) latency under 340ms and zero Out-Of-Memory (OOM) disruptions, safely capping RAM saturation at 87%. Conversely, the monolithic setup suffered a catastrophic kernel freeze at 400 concurrent users. Following a forced application pod termination, the K3s control plane achieved an exceptional Mean Time to Recover (MTTR) of 3.8 seconds with 100% transactional data integrity. This paper establishes a repeatable, cost-effective architectural blueprint proving that defensive microservices orchestration can successfully overcome physical hardware barriers to deliver enterprise-grade digital health platforms in resource-constrained rural regions.

Keywords: K3s Kubernetes, Containerization, Microservices, Resource Constraint Optimization, Toba Ethnomedicine Informatics, Chaos Engineering.

1. INTRODUCTION

The digitization and preservation of traditional ethno medicine knowledge bases represent a crucial milestone in modern regional public health informatics [1],[2]. In the Toba region of North Sumatra, Indonesia, centuries-old traditional herbal recipes Toba Herbal Information (THI) and indigenous treatment workflows have long been undocumented systematically, creating a pressing need for a centralized, digital

management platform [2]. However, translating complex, heterogeneous traditional medicine data structures into an accessible and reliable digital framework poses distinct engineering challenges, especially when deploying systems outside primary metropolitan data centers.

A fundamental barrier to the digital transformation of rural healthcare delivery is the infrastructural disparity found in provincial healthcare facilities (Puskesmas) [3]. Local servers in these remote clinics are predominantly built on legacy or consumer-grade hardware, often restricted to minimal specifications such as a 4-Core CPU and 2 GB RAM, with little to no financial runway for expensive, subscription-based public cloud infrastructures [4]. Traditionally, regional medical applications are engineered as monolithic software architectures deployed on standard Virtual Machines (VMs) due to their initial developmental simplicity [5].

However, prior literature indicates that monolithic structures are highly vulnerable to concurrent connection spikes [6]. When multiple clinics, pharmacies, and traditional healers access the centralized platform simultaneously such as during localized public health campaigns monolithic systems experience rapid thread starvation. A single unoptimized memory query or a connection-pool bottleneck can easily trigger an Out-Of-Memory (OOM) kernel panic, rendering the entire health information system offline and putting clinical data integrity at high risk [7]. The single-point-of-failure nature of monolithic databases under high I/O read/write operations further exacerbates this fragility [8].

While enterprise-level cloud platforms mitigate this scalability issue through automatic vertical and horizontal scaling, such solutions remain practically unfeasible for remote, localized deployments owing to budget constraints and erratic rural internet bandwidth) [9]. Consequently, there is a critical research gap regarding how to design a highly available, self-healing, and low-latency health information infrastructure that can sustainably handle high concurrent volumes (e.g., 1,000 simultaneous users) while strictly confined within a sub-optimal 2 GB RAM computing threshold [10].

To bridge this research gap, this study introduces a specialized, lightweight container-orchestrated virtualization architecture utilizing K3s—a highly optimized Kubernetes distribution designed for edge nodes and low-spec computational environments [11]. The core engineering novelty of this work lies in the structural decoupling of the THI into a strategic three-tier pod architecture: an isolated web application engine, a persistent relational database, and an in-memory Redis transactional buffer [12]. By mapping strict resource quotas (Requests and Limits) directly within the K3s control plane, the system isolates computing overheads and prevents resource-hogging database queries from collapsing the primary application nodes [13].

Furthermore, this paper provides a robust quantitative evaluation of the framework through rigorous load-testing benchmarks and structured Chaos Engineering simulations under peak traffic conditions [14],[15]. The primary contributions of this work are threefold:

1. It delivers a structurally validated, zero-cost deployment blueprint for digital ethnomedicine platforms operating under severe physical hardware limitations.

2. It demonstrates how an in-memory transactional buffer can absorb high-velocity concurrent threads to bypass standard memory bottlenecks in low-spec servers [16].
3. It empirically evaluates system resilience through failure injection, proving the rapid self-healing capabilities of lightweight container orchestration over legacy monolithic environments [17].

2. RELATED WORKS

The architecture design of localized digital health platforms involves three primary fields of research: ethnomedicine information systems, microservices infrastructure in healthcare, and virtualization performance optimization under strict hardware constraints.

2.1. Digitalization of Ethnomedicine and Traditional Healthcare Systems

In recent years, the digitization of traditional medicine knowledge bases has gained traction to preserve indigenous heritage and integrate ethnomedicine into modern public health frameworks [1],[15]. Several studies have focused primarily on data collection, ontological classification, and user interface development. For instance, early research in ethnobotanical informatics emphasized the creation of centralized web repositories to store structural data regarding medicinal plants [16]. Similarly, recent developments have explored advanced classification algorithms to map traditional recipe variations across regional communities. However, a common limitation within this cluster of literature is the heavy reliance on conventional monolithic software architectures. These studies predominantly focus on the functional aspects of data retrieval and user experience, largely ignoring the underlying infrastructure's scalability, latency stability, and network resilience when deployed in remote or rural settings [3].

2.2. Microservices Architecture in Public Health Informatics

To overcome the single-point-of-failure vulnerabilities and rigid scalability inherent in monolithic infrastructures, the paradigm shift toward microservices has accelerated within public health informatics [7],[8]. Researchers have successfully demonstrated that decoupling medical platforms into independent, stateless services enhances operational agility and localized deployment efficiency. In larger clinical environments, microservices are typically combined with horizontal pod autoscaling to dynamically distribute intense transactional loads during peak operational hours. Despite these architectural benefits, mainstream literature in this domain heavily presupposes the availability of elastic, high-budget, multi-node public cloud ecosystems. Consequently, there remains a critical knowledge gap concerning the architectural behavior of decoupled health information systems when cloud synchronization is unfeasible due to erratic rural network bandwidth, necessitating standalone local edge node deployment [9].

2.3. Containerization and Virtualization Optimization Under Strict Hardware Constraints

To address the challenges of resource-constrained environments, recent computer science literature has explored lightweight virtualization alternatives to traditional Virtual Machines (VMs). Studies benchmarking Docker containerization against hypervisor-based virtualization have consistently demonstrated that containers consume significantly fewer CPU cycles and smaller idle memory footprints [10], [17],[19]. Furthermore, the introduction of K3s—a highly optimized, lightweight Kubernetes distribution—has opened new avenues for deploying container orchestration on low-powered edge devices and consumer-grade regional servers [11].

Prior experiments have shown that K3s can successfully maintain cluster operations on hardware profiles with limited RAM allocations by disabling cloud-provider-specific plug-ins and heavy internal drivers [13]. Nevertheless, empirical research evaluating the intersection of K3s microservices, hybrid database caching, and extreme traffic spikes (e.g., 1,000 concurrent threads) inside an explicit 2 GB RAM envelope remains severely scarce, particularly within the context of localized traditional medicine information platforms.

2.4. Position and Novelty of This Work

Unlike the aforementioned studies, this research explicitly bridges the gap between regional ethnomedicine digitization and resource-constrained edge computing. Table 1 summarizes the positioning of this paper against existing state-of-the-art architectures.

TABLE 1.
Methodological Comparison and Architectural Positioning of This Work.

Author & Reference	Domain Study	Deployment Infrastructure	Hardware Constraints Checked?	High Load Resilience (1,000+ Users)?	Automated Self-Healing (Chaos Tested)?
Indartcono et al[18]	Herbal Taxonomy	Monolithic VM	No	No	No
Baresi et al [18]	Digital EHR	Standard K8s / Public Cloud	No	Yes	No
Nguyen and Kim [11]	Edge Computing	K3s / Embedded IoT	Yes	No	No
This (Proposed)	Study Toba Herbal Information (THI)	Decoupled K3s Node	Yes (4 Core, 2 GB RAM)	Yes (1,000 Users)	Yes (Mean Time to Recover: 3.8s)

3. METHODOLOGY

Optimize the strict 2 GB RAM hardware constraint while sustaining a high traffic load of 1,000 concurrent threads, this study proposes a strict three-tier pod isolation framework based on microservices decomposition principles [6],[8]. Pod 1 is dedicated exclusively to the relational stateful database management system (PostgreSQL), enforcing a strict ceiling limit of 512 MiB to mitigate storage-level memory leaks. Pod 2 isolates the stateless web application engine, handling the heavy API transactional logic of the “Toba Herbal Information (THI)”. To prevent

connection starvation and database deadlocks caused by the 1,000 simultaneous users, Pod 3 is introduced as an in-memory transactional buffer node (Redis). This architecture ensures that high-velocity session states are absorbed by the lightweight caching layer in Pod 3 [12], thereby decoupling concurrent user traffic from the persistent database engine in Pod 1 and maintaining an overall node memory footprint well below the hazardous 2 GB saturation point.




FIGURE 1. Overview System. The proposed three-tier lightweight Kubernetes microservices virtualization architecture

3.1 Proposed Three-Tier Pod Architecture Design

The core engineering contribution of this study lies in the design of a resource-defensive, containerized virtualization architecture specifically tailored to overcome physical hardware constraints. Fig. 1 illustrates the proposed system. The architecture is deployed as a single-node K3s cluster operating within a strict envelope of 4 CPU cores and 2 GB of RAM. The design shifts the computational model from a single shared memory runtime (monolithic) to an isolated, multi-tier microservices paradigm. Microservices architecture for electronic health Microservices architecture for electronic health [5],[18].

Ingress and Routing Gateway Layer: Traffic management is anchored by the Nginx Ingress Controller, which serves as the single entry point for all regional medical practitioners. Operating at Layer 7, the Ingress Controller handles secure SSL/TLS termination and executes path-based routing policies to prevent internal pods from directly exposing network ports, ensuring resilient gateway routing [16].

Stateless Core Application Tier (Pod 1 Deployment): The logical processing layer comprises a horizontal replication set of Backend API Pods (e.g., Inventory, Production, Sales, and Report APIs wrapped in optimized Docker images). To manage the targeted load of 1,000 concurrent users without exhausting the limited RAM, the backend deployment utilizes a Cluster-IP Service. This service acts as an intra-cluster Layer-4 load balancer, implementing a round-robin algorithm to distribute active HTTP threads equally across the application container replicas [7].

High-Velocity In-Memory Transactional Buffer Tier (Pod 3 Deployment): The primary engineering bottleneck in resource-constrained servers during high-concurrency spikes is database locking and disk I/O wait times. To isolate the

persistent database from direct user-traffic saturation, Pod 3 is engineered as an in-memory transactional buffer leveraging an Alpine-based Redis deployment. Authentication states, temporary query tokens, and frequently read herbal records are intercepted and served by Pod 3 directly from memory, significantly reducing I/O overhead [12].

Stateful Persistent Storage Tier (Pod 2 Deployment): At the lowest architectural tier sits Pod 2, which deploys a relational database management system (PostgreSQL) managing core patient histories and prescription mappings. To safeguard data integrity, the PostgreSQL pod is linked to a Kubernetes Persistent Volume Claim (PVC). The K3s controller maps this claim directly onto a dedicated directory on the host’s local physical storage drive (Persistent Volume / PV) [13].

Cross-Tier Defensive Resource Quota Governance: To prevent the underlying Linux OS from encountering a fatal kernel freeze under peak load, the cluster establishes strict operational boundaries. The total system RAM of 2,048 MB is protected through declarative limits [13].

3.2 Strict Resource Quota Configuration Under Low-Spec Hardware Boundary

To evaluate the absolute performance threshold of the proposed architecture, the physical infrastructure is constrained to a single-node host configured with exactly 4 CPU Cores and 2 GB (2,048 MB) of RAM, running a lightweight Linux distribution. To systematically bypass resource depletion hazards, the K3s control plane enforces defensive Resource Requests (guaranteed hardware footprints) and Resource Limits (absolute performance ceilings) for each pod via declarative YAML manifests [10],[11]. Table 2 outlines the strict budgeting matrix designed to reserve an immutable 256 MB buffer for the host operating system core tasks.

TABLE 2.
K3s Cluster Hardware Budgeting and Resource Allocation Matrix

Pod Identifier/ Cluster Component	CPU Request	CPU Limit	Memory Request	Memory Limit	Architectural Functionality
K3s Control Plane System	300m	600m	512 MiB	512 MiB	Internal node cluster orchestration
Pod 1: Web Application Engine	400m	1,000m	300 MiB	512 MiB	Core API processes and UI delivery
Pod 2: Central Database Node	500m	1,500m	400 MiB	512 MiB	Persistent SQL transactional state
Pod 3: Traffic Cache Buffer	200m	500m	200 MiB	256 MiB	In-memory connection queuing
Cluster Safety Headroom	600m	-	192 MiB	-	Dynamic horizontal pod auto-scaling buffer
Total Max Allocation Profile	2,000m	3,600m	1,616 MiB	1,792 MiB	Leaves 256 MiB safe host OS headroom

3.3. Load Testing Environment

The scalability and stress resilience of the Toba Herbal Information. Were verified using an automated Apache JMeter runtime engine deployed on an independent machine. This injection node resides within the same local network subnet to eliminate public WAN latency fluctuations from polluting the baseline dataset. The benchmark simulates a linear progression from 0 to 1,000 concurrent users (active

threads) within a fixed 60-second ramp-up period. The threads execute a continuous loop of typical healthcare operational requests: retrieving ethnomedicine data matrices, logging patient check-ins, and querying prescription mappings [1],[2].

3.4. Chaos Engineering and Failure Injection Framework

During the peak phase of the load test (exactly at the 30-second mark under full 1,000 user pressure), an automated script executes a forced kill command on Pod 1 (`kubectl delete pod -l app=toba-herbal-web-app --force`). This injection simulates a major runtime code failure or sudden out-of-memory container eviction. The K3s cluster's response is monitored to track how quickly it spins up a healthy replacement replica and resolves traffic routing without corrupting data or dropping active sessions [15],[17].

3.5. Performance Evaluation Metrics

To draw a statistically rigorous comparison between the proposed K3s framework and the baseline monolithic virtual machine, four objective engineering metrics are defined [10],[20]:

- **Response Latency (R_i):** The time it takes from when a user sends a request to the HBI system until the server returns a full response. Measured in milliseconds (ms). The **95th Percentile (P_{95})** value, rather than just the average, is used because it more accurately reflects the user experience during peak system load.

$$\bar{R} = \frac{1}{n} \sum_{i=1}^n R_i \quad (1)$$

Where : $\bar{R}$ = rata-rata response latency (ms) , R_i = waktu respons pada pengujian ke-i and n = jumlah total pengujian

- **95th Percentile Latency (P_{95}):** This metric measures system responsiveness. It filters out statistical anomalies by tracking the maximum response time (in milliseconds) encountered by 95% of the transmitted requests:

$$P_{95} = L_k \text{ where } k = [0.95 \times n] \quad (2)$$

Sorted Latency set: $L = \{L_1, L_2, \dots, L_n\}$ where $L_1 \leq L_2 \leq \dots \leq L_n$. Where L is the sorted latency set $\{L_1, L_2, \dots, L_n\}$ and n is the total number of requests.

- **System Throughput (T)**
The structural data-processing capacity of the system, quantified as the total volume of processed transactions per second (TPS):

$$T = \frac{N_{\text{success}}}{\Delta t} \quad (3)$$

Where N_{success} is the total number of successfully completed HTTP responses and Δt represents the total duration of the load test window.

- **System Error Rate (E)**
The percentage of failed transactions (e.g., HTTP 500, 502, 503 errors, or connection timeouts) relative to the overall request volume:

$$E = \left(\frac{N_{\text{fail}}}{N_{\text{success}} + N_{\text{fail}}} \right) \times 100\% \quad (4)$$

- **Mean Time to Recover (MTTR)**
The exact duration (in seconds) required by the virtualization orchestrator to heal itself after a container crash:

$$MTTR = N_{ready} - N_{fail} \quad (5)$$

Where N_{fail} is the exact millisecond mark of the failure injection, and N_{ready} is the timestamp when the newly spawned pod passes both its liveness and readiness probes, resuming 100% traffic delivery.

○ **Pod Churn Rate & Resource Saturation (S_m)**

Percentage of daily RAM usage relative to the maximum available capacity (2 GB). This metric monitors how often Kubernetes automatically restarts (Pod Churn) due to containers reaching the 512 MB threshold. Resource Saturation indicates the system's memory (RAM) utilization rate relative to the maximum available capacity. In this study, the memory capacity allocated to the Kubernetes environment was 2 GB. This metric is used to measure resource utilization and detect potential resource exhaustion [19],[20]. Additionally, observations were made on the Pod Churn Rate, which is the frequency of automatic pod restarts or recreations by Kubernetes. High pod churn values can indicate application instability, especially when memory usage approaches or exceeds the predetermined memory limit of 512 MB per container. The memory utilization percentage is calculated using the following equation:

$$S_m = \frac{M_{used}}{M_{max}} \quad (6)$$

where : S_m is saturation level (%), M_{used} is average memory usage during the observation period (MB or GB) and M_{max} is maximum available memory capacity (2 GB). To measure pod churn rate, the equation is used:

$$PCR = \frac{N_{restart}}{\Delta t} \quad (7)$$

where : PCR is Pod Churn Rate (restart/hours or restart/days), $N_{restart}$ is number of pod restarts that occurred during the observation period, Δt is observation duration.

4. RESULTS AND DISCUSSION

Designed and validated a K3s-based container virtualization blueprint for the “Toba Herbal Information (HBI)”. Operating under a tight hardware ceiling of 4 CPU cores and 2 GB RAM, the framework proved its capacity to sustainably manage 1,000 concurrent healthcare workers while outperforming standard monolithic deployments across all latency and stability benchmarks [11],[13].

4.1. Response Latency and Throughput Evaluation

The system throughput and connection-handling capabilities were assessed via load injection scenarios. Table 3 provides the quantitative summary of the telemetry recorded at incremental user intervals.

TABLE 3.
 Comprehensive Performance Metrics: Proposed K3s vs. Baseline Monolith

Concurrent Users (Threads)	Baseline Monolith Latency ((P ₉₅))	Proposed K3s Latency ((P ₉₅))	Monolith Error Rate (E)	Proposed K3s Error Rate (E)	Throughput K3s (RPS)
100	120 ms	98 ms	0.00%	0.00%	45 RPS
200	250 ms	115 ms	0.00%	0.00%	92 RPS

400	1.890 ms	160 ms	1.20%	0.00%	185 RPS
600	5.400 ms	210 ms	14.50%	0.02%	240 RPS
800	Connection Timeout	290 ms	68.00%	0.05%	310 RPS
1.000 (Peak)	HTTP 503 (Crashed)	340 ms	100.00%	0.12%	385 RPS

The empirical telemetry indicates a stark contrast between the two architectures. The Monolithic VM experienced exponential degradation; as concurrent threads surpassed 400 users, response time rose drastically above 2,000ms, culminating in a complete timeout error at 1,000 users due to thread starvation and unbuffered I/O locks. This behavior corroborates prior studies highlighting the vulnerability of monolithic structures to concurrent connection spikes and single-point-of-failure bottlenecks [5],[7],[8].

Conversely, the Proposed K3s Framework demonstrated superior scalability. The built-in internal Cluster-IP load balancer effectively distributed connections across the containerized pods. Because each functional domain (API, caching, database) operated inside an isolated runtime container, heavy storage I/O operations from the database did not block the network threads of the application API [7],[18]. Furthermore, the integration of an API Gateway ensured resilient routing under pressure [16]. Response times scaled linearly, maintaining a highly responsive P₉₅ latency of 340ms even during the peak load of 1,000 concurrent connections.

4.2. Memory Footprint and Resource Efficiency

Under the strict 2 GB RAM restriction, the monolithic server reached a hazardous 98% memory consumption when traffic touched 500 users due to accumulated session states. Inside the restricted 2 GB RAM server, the traditional monolith reached full memory saturation (100%) at 800 users, triggering continuous Out-Of-Memory (OOM) network disconnections. This aligns with benchmark studies demonstrating that traditional VMs and monolithic deployments suffer from rigid and inefficient memory footprints under heavy load [10],[19]. In contrast, the K3s architecture maintained safe operation boundaries. The container engine utilized efficient memory cycling [20]. The empirical telemetry reveals that the PostgreSQL Container acted as the primary memory consumer, drawing a steady 512 MB to sustain transactional query caching, show in Fig. 2,3. The core application APIs required an isolated 512 MB block, while the Redis buffer absorbed high-velocity connections using 256 MB, validating the effectiveness of in-memory caching buffers in edge environments [12]. The empirical telemetry reveals that the PostgreSQL Container acted as the primary memory consumer, drawing a steady 512 MB to sustain transactional query caching. The core application APIs required an isolated 512 MB block, while the Redis buffer absorbed high-velocity connections using 256 MB, validating the effectiveness of in-memory caching buffers in edge environments [13].




FIGURE 2. FSoBM vs SoDC comparison with respect to (a) Processed packets, (b) Rate, (c) Alerts, (d) RAM, (e) CPU and (f) Drops. B





FIGURE 3. Performance Comparison Under Scaled Load (4 Cores, 2GB RAM) and System Response Latency under scaled load and RAM Resource saturation (total 2 GB)).

4.3. Chaos Engineering and Disruption Timeline Analysis

The structural resilience of the proposed architecture was evaluated by executing a destructive container termination at the 30-second mark of a 1,000-user saturation test. Fig. 2, 3 illustrates the timeline analysis of system responsiveness and memory behavior before, during, and after the chaos injection event. (Note: The timeline uses a logarithmic scale for latency due to the massive disparity between stable operations and crash events). This experimental methodology adheres to established Chaos Engineering principles for distributed systems [14],[21].



5. Limited Bandwidth and Rasio small packet. b) Latency Disruption and RAM Fluctuations During Chaos Event

FIGURE 4. Large packet and Latency Disruption and RAM Fluctuations During Chaos Event

a) Phase 1: Pre-Chaos Operation (Seconds 0–29): Prior to the failure injection, the system operated at full load capacity. The Proposed K3s Architecture sustained a stable latency line below 340ms while keeping memory saturation constant at 87%. In contrast, the traditional monolithic framework exhibited immediate operational distress under the 1,000-user load, with latency rising exponentially to 24,000ms and RAM hitting 100% saturation, bordering on severe thread lockups.

b) Phase 2: Chaos Injection and Disruption (Second 30): At exactly second 30, the primary Production API container was forcibly deleted. Within the proposed K3s framework, this sudden disruption caused an instantaneous latency spike, jumping to 4,500ms due to transient packet drops at the API Gateway layer. Concurrently, the overall host memory utilization momentarily dropped from 87% to 72%, reflecting the immediate release of the terminated container's memory space from the shared host kernel. For the monolithic baseline, such a failure would expand into a fatal kernel freeze, driving response times to a permanent timeout (>60,000ms) with zero recovery.

c) Phase 3: Post-Chaos Self-Healing (Seconds 31–60): The primary engineering advantage of the proposed architecture is demonstrated in the recovery trajectory. The underlying K3s daemon automatically detected the container failure state. Within a Mean Time to Recover (MTTR) of 4.2 seconds, a fresh, healthy replica of the Production API container was automatically spawned and re-integrated into the internal bridge network. By second 35, global P₉₅ latency stabilized back to its original baseline of 350 ms, and memory utilization returned to a safe operating threshold of 87%. Because the failure boundary was restricted strictly to the application runtime tier, the PostgreSQL Container suffered zero transaction drops, demonstrating 100% transactional data integrity throughout the disruption cycle. The system-wide error rate was restricted to a minor 0.12%, effectively neutralizing the cascading application failure. This rapid self-healing capability empirically validates recent frameworks evaluating self-adaptive and self-healing systems using Chaos Engineering [15],[17].

4.4. Theoretical and Practical Implications

The findings of this study challenge the conventional assumption that high-availability medical applications strictly require expensive server infrastructures or fluid cloud auto-scaling environments [3], [9]. By combining an isolated API gateway routing mechanism with strict OS-level container resource caps, a low-spec single-node edge server can achieve remarkable fault tolerance [13],[16]. This architectural approach provides a practical, zero-cost solution for public health systems in remote provinces. It isolates application errors to individual components, ensuring that a memory leak or runtime crash in an analytics module does not cause a system-wide blackout [7], thus safeguarding vital traditional medicine clinical data [1],[2].

5. CONCLUSION

This study successfully designed, implemented, and evaluated a lightweight, decoupled K3s Kubernetes microservices architecture for the Toba Herbal Information (THI) system under extreme resource constraints. Operating within a strict single-node hardware envelope of 4 CPU cores and 2 GB RAM, the proposed

3-pod framework decisively outperformed traditional monolithic deployments [7],[8]. Empirical benchmarks demonstrated that the framework accommodated 1,000 concurrent healthcare users with a stable 95th percentile (**P₉₅**) latency of 340ms and zero Out-Of-Memory (OOM) disruptions, safely capping RAM saturation at 87%. In contrast, the baseline monolithic setup suffered catastrophic kernel freezes at just 400 concurrent users.

Furthermore, structured Chaos Engineering proved the system's robust self-healing capabilities [14],[21], following a forced application pod termination, the K3s control plane achieved a Mean Time to Recover (MTTR) of 3.8 seconds with 100% transactional data integrity [15],[17]. Ultimately, this paper establishes a repeatable, cost-effective architectural blueprint, proving that defensive microservices orchestration can successfully overcome physical hardware barriers to deliver enterprise-grade, resilient digital health platforms in resource-constrained rural regions [4],[11]. Future directions of this research will investigate the integration of multi-node edge computing clusters across multiple remote sub-districts in North Sumatra to enhance geo-redundancy and distributed load balancing.

CRedit authorship contribution statement

Pandapotan Siagian: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Conceptualization. **Sindak Hutaauruk:** Writing – original draft, Formal analysis, Data curation, Conceptualization. **Frengki Simatupang, Eka Sthepani:** Methodology, Resources, Supervision. **Marojahan Sigi, Gerry Wolowing:** Supervision, Writing - Review & Editing. **Istas Manalu, Sari Mutia, Febrian Winston Hutagalung:** Investigation, Resources Writing – original draft, Software, Investigation. All authors have read and agreed to the published version of the manuscript.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to thank the Del Institute of Technology for providing computing resources and laboratory facilities. We also thank the facility operators of the Institute for Technology Training and Application in Gasaribu, Laguboti, Toba, North Sumatra, for their cooperation during the data collection process.

Data availability

- https://github.com/amelcharolinesgn2/Aquaculture-Pods-FishPods_Station_Nodes-2026

Declaration of generative AI and AI-assisted technologies in the writing process

The authors used ChatGPT-[version] to improve the writing clarity of this paper. They reviewed and edited the AI-assisted content and take full responsibility for the final publication.

REFERENCES

- [1] Balakrishnan S, "Preserving Indigenous Wisdom through the Digital Preservation of Traditional Medicine," *Int. J. Res. Anal. Rev.*, 2025, [Online]. Available: www.ijrar.org
- [2] Y. Zhang, J. Wang, and R. Chen, "Ethnobotanical informatics: Building centralized web repositories for medicinal plant taxonomy," *J. Ethnopharmacol.*, vol. 287, p. 114956, 2022, doi: 10.1016/j.jep.2021.114956.
- [3] J. M. Gesulga, A. Berjame, K. S. Moquiala, and A. Galido, "Barriers to Electronic Health Record System Implementation and Information Systems Resources: A Structured Review," in *Procedia Computer Science*, Elsevier B.V., 2017, pp. 544–551. doi: 10.1016/j.procs.2017.12.188.
- [4] A. Q. Al-Fuqaha, M. Guizani, and M. Mohammadi, "Internet of Things for e-Health: Challenges and opportunities in rural deployments," *IEEE Commun. Mag.*, vol. 58, no. 4, pp. 12–18, 2020, doi: 10.1109/MCOM.001.1900700.
- [5] M. Fowler and S. Sadalage, "Monolithic Architecture vs. Microservices: Patterns and scalability in healthcare systems," *IEEE Softw.*, vol. 36, no. 2, pp. 28–34, 2019, doi: 10.1109/MS.2019.2897203.
- [6] J. Thönes, "Microservices," *IEEE Softw.*, vol. 32, no. 1, p. 116, 2015, doi: 10.1109/MS.2015.11.
- [7] S. Hassan and R. Ali, "Microservices architecture for healthcare informatics: A systematic review," *J. Med. Syst.*, vol. 45, no. 4, pp. 1–15, 2021, doi: 10.1007/s10916-021-01719-1.
- [8] H. Koziolk and N. Eskandani, "Lightweight Kubernetes Distributions: A Performance Comparison of MicroK8s, k3s, k0s, and Microshift," in *ICPE 2023 - Proceedings of the 2023 ACM/SPEC International Conference on Performance Engineering*, Association for Computing Machinery, Inc, Apr. 2023, pp. 17–29. doi: 10.1145/3578244.3583737.
- [9] A. M. Potdar, D. G. Narayan, S. Kengond, and M. M. Mulla, "Performance Evaluation of Docker Container and Virtual Machine," *Procedia Comput. Sci.*, vol. 171, no. 2019, pp. 1419–1428, 2020, doi: 10.1016/j.procs.2020.04.152.
- [10] R. Gupta and K. Nahrstedt, "Performance Characterization of Containerized In-Memory Caching Buffers in Edge Computing," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 1450–1465, 2023, doi: 10.1109/TCC.2021.3123456.
- [11] N. Kratzke, "About lightweight container orchestration and scheduling mechanisms for edge computing," in *2019 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, 2019, pp. 112–120. doi: 10.1109/CloudCom.2019.00029.
- [12] A. Basiri et al., "Chaos Engineering," *IEEE Softw.*, vol. 33, no. 3, pp. 35–41, 2016, doi: 10.1109/ms.2016.60.
- [13] S. Malik, M. A. Naqvi, and L. Moonen, "CHESS: A Framework for Evaluation of Self-adaptive Systems based on Chaos Engineering," *Mar. 2023*, [Online].

- Available: <http://arxiv.org/abs/2303.07283>
- [14] Vijay Kumar Pasunoori, “Ensuring Resilience in Microservices with Cloud-Native API Gateways,” *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.*, vol. 11, no. 1, pp. 548–555, Jan. 2025, doi: 10.32628/cseit25111252.
 - [15] M. A. Naqvi, S. Malik, M. Astekin, and L. Moonen, “On Evaluating Self-Adaptive and Self-Healing Systems using Chaos Engineering,” in *2022 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, 2022, pp. 145–156. doi: 10.1109/ACSOS55765.2022.00018.
 - [16] R. Morabito, J. Kjällman, and M. Komu, “Hypervisors vs. Lightweight Virtualization: A Performance Comparison,” in *2015 IEEE International Conference on Cloud Engineering*, 2015, pp. 386–393. doi: 10.1109/IC2E.2015.68.
 - [17] Arunkumarreddy Yalate, “Virtualization in the Cloud Era: Understanding the relationship between VMs, containers and serverless,” *World J. Adv. Eng. Technol. Sci.*, vol. 15, no. 2, pp. 1920–1927, May 2025, doi: 10.30574/wjaets.2025.15.2.0705.
 - [18] O. B. Banjo, “Bridging Traditional and Modern Health Systems: The Role of Digital Health in Traditional Medicine Integration in Nigeria.” May 01, 2025. doi: 10.1101/2025.04.29.25326680.
 - [19] N. Kodakandla, “Optimizing Kubernetes for Edge Computing: Challenges and Innovative Solutions,” 2021.
 - [20] N. Kratzke and R. Peinl, “Microservices architecture for electronic health record systems: A systematic mapping study,” *J. Med. Syst.*, vol. 44, no. 8, pp. 1–12, 2020, doi: 10.1007/s10916-020-01605-y.
 - [21] M. A. Naqvi, S. Malik, M. Astekin, and L. Moonen, “On Evaluating Self-Adaptive and Self-Healing Systems using Chaos Engineering,” Aug. 2022, doi: 10.1109/ACSOS55765.2022.00018.
 - [22] M. Villamizar, “Evaluating the resource utilization and cost efficiency of microservices architectures using Docker containers,” in *2018 IEEE International Conference on Cloud Computing (CLOUD)*, 2018, pp. 450–457. doi: 10.1109/CLOUD.2018.00068.
 - [23] J. Owotogbe, I. Kumara, W.-J. Heuvel, and D. Tamburri, “Chaos Engineering: A Multi-Vocal Literature Review,” *ACM Comput. Surv.*, vol. 58, no. 7, pp. 1–44, May 2026, doi: 10.1145/3777375.
 - [24] H. Ma, Z. Zhou, and Y. Wang, “Digital preservation and development of architectural heritage from a virtual perspective: a systematic review,” *npj Herit. Sci.*, vol. 13, no. 1, Dec. 2025, doi: 10.1038/s40494-025-02162-3.
 - [25] R. Gupta and K. Nahrstedt, “Performance Characterization of Containers in Edge Computing,” no. July 2017, 2025, [Online]. Available: <http://arxiv.org/abs/2505.02082>