Sparkling Light Publisher

Sparklinglight Transactions on Artificial Intelligence and Quantum Computing

journal homepage: https://sparklinglightpublisher.com/

Zero Trust Architecture and Azure Red Hat OpenShift for Federal Cloud Modernization: Evidence from a Mission-Critical VA Financial Systems Implementation

Naga Venkata Chaitanya Akula[a]

[a]*Lead Senior Computers and Innovation Scientist, 668 Pickrell Loop, Liberty Hill, TX 78642, United States*

Abstract

Federal agencies face a compound problem: aging infrastructure that runs critical services, security mandates that keep tightening, and budget pressure that never lets up. This paper examines how Azure Red Hat OpenShift (ARO), combined with Zero Trust architecture, infrastructure-as-code, and Site Reliability Engineering (SRE) practices, addresses all three. Drawing on an 18-month production deployment at the U.S. Department of Veterans Affairs Financial Service Center—a system processing over $240 billion annually for 9 million enrolled veterans—the paper presents measured outcomes across availability, security, and cost dimensions. We formalize Zero Trust access control using a multi-dimensional trust scoring function $\mathsf{T} : \mathsf{I} \times \mathsf{D} \times \mathsf{N} \rightarrow [0, 1]$, derive availability bounds from a Markov reliability model, and quantify security posture with a composite vulnerability index Ψ. The implementation migrated 127 enterprise applications from legacy middleware to an ARO cluster spanning three Azure Government availability zones. Results show 99.99% system uptime, a 98% reduction in unplanned downtime, 85.9% fewer critical vulnerabilities, and a 30% infrastructure cost reduction, with a 16.5-month break-even on a $4.2 million investment.

Keywords: Zero Trust Architecture; Azure Red Hat OpenShift; Federal Cloud Modernization; Container Security; FedRAMP Compliance; Markov Reliability Model; Trust Scoring; NIST 800-53; FISMA; SRE

1. Introduction

1.1. Federal Modernization Pressures

The U.S. Department of Veterans Affairs Financial Service Center (VA FSC) processes more than $240 billion annually and serves 9 million enrolled veterans [1]. Executive Order 14028 directed agencies to adopt Zero Trust architectures [2], OMB M-22-09 attached FY2024 deadlines for phishing-resistant MFA and application-level authorization [3], and FISMA demands continuous ATO evidence throughout. Legacy VM-based infrastructure wastes 40% of compute capacity through static over-provisioning [4] and requires 14-day patch cycles driven by maintenance window constraints.

1.2. Why Azure Red Hat OpenShift

Azure Red Hat OpenShift (ARO) is a managed Kubernetes platform, available in Azure Government with FedRAMP High authorization covering CJIS, IRS 1075, ITAR, and DFARS [8, 9]. Native integration with Azure Active Directory, Key Vault, Container Registry, and Monitor allows agencies to wire Zero Trust controls to existing investments rather than replace them [14].

E-mail address of authors: chaitanya.akula@hitssllc.com,

Please cite this article as: Naga Venkata Chaitanya Akula., Zero Trust Architecture and Azure Red Hat OpenShift for Federal Cloud Modernization: Evidence from a Mission-Critical VA Financial Systems Implementation (2026), 6(1), 23–31. **ISSN (Online):2583-0732**. Received Date: 2026/03/15, Reviewed Date: 2026/03/31, Published Date: 2026/04/05.

1.3. Research Questions

1. **RQ1**: Can container platforms sustain 99.99% availability for federal mission-critical workloads at FedRAMP High compliance?
2. **RQ2**: What architectural patterns make Zero Trust enforceable in hybrid federal cloud environments?
3. **RQ3**: What is the realistic economic case for legacy-to-ARO migration at scale?

2. Mathematical Foundations

2.1. Zero Trust Trust Score Function

Definition 1 (Trust Score)**.** *Let* I *denote the identity space,* D *the device compliance space, and* N *the network context space. The* Zero Trust access trust score *is the function*

$$\mathsf{T}(i, d, n) = \sum_{k=1}^{K} \omega_k \cdot \phi_k(i, d, n) \tag{1}$$

where $\phi_k : \mathsf{I} \times \mathsf{D} \times \mathsf{N} \to [0, 1]$ *are* K *policy evaluation functions,* $\omega_k \geq 0$ *are importance weights with* $\sum_{k=1}^{K} \omega_k = 1$*, and* $\mathsf{T}(i, d, n) \in [0, 1]$.

Access to resource r is granted if and only if:

$$\mathsf{T}(i, d, n) \geq \vartheta_r, \quad \vartheta_r \in (0, 1] \tag{2}$$

where ϑ_r is the resource-specific trust threshold. For the VA FSC implementation the five CISA pillars map to evaluation functions $\phi_1, \ldots, \phi_5$:

$$\phi_k(i, d, n) = 1 - \exp(-\lambda_k \cdot \sigma_k(i, d, n)) \tag{3}$$

where $\sigma_k \in [0, \infty)$ is the raw signal from pillar k and $\lambda_k > 0$ is the sensitivity parameter. The effective weights used in production are $\boldsymbol{\omega} = (0.30,\ 0.20,\ 0.20,\ 0.20,\ 0.10)$ for Identity, Devices, Networks, Workloads, and Data respectively.

2.2. Availability Model

Definition 2 (System Availability)**.** *For a system with* n *independent availability zones, let* μ_i *be the repair rate and* λ_i *the failure rate of zone* i*. The steady-state availability of zone* i *is*

$$A_i = \frac{\mu_i}{\mu_i + \lambda_i} \tag{4}$$

For a system requiring at least k of n zones to be operational (a k-out-of-n model), the system availability is:

$$A_{\text{sys}} = \sum_{j=k}^{n} \binom{n}{j} \bar{A}^{j} (1 - \bar{A})^{n-j} \tag{5}$$

where $\bar{A} = \frac{1}{n} \sum_{i=1}^{n} A_i$ is the mean zone availability [16]. For the ARO deployment with $n = 3$ zones and $k = 2$ (quorum), assuming $A_i = 0.9999$ per zone:

$$A_{\text{sys}} = 3\bar{A}^2(1 - \bar{A}) + \bar{A}^3 \approx 0.999999\overline{9} \tag{6}$$

confirming the theoretical upper bound exceeds the measured 99.99%.

2.3. Mean Time Calculations

The Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR) relate to availability via:

$$A = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} = \frac{1/\lambda}{1/\lambda + 1/\mu} \tag{7}$$

For the legacy baseline ($A_{\text{legacy}} = 0.995$, $\text{MTTR}_{\text{legacy}} = 6.2$ h):

$$\text{MTBF}_{\text{legacy}} = \frac{A_{\text{legacy}}}{1 - A_{\text{legacy}}} \cdot \text{MTTR}_{\text{legacy}} = \frac{0.995}{0.005} \times 6.2 \approx 1238 \text{ h} \tag{8}$$

For ARO ($A_{\text{ARO}} = 0.9999$, $\text{MTTR}_{\text{ARO}} = 3.4$ h):

$$\text{MTBF}_{\text{ARO}} = \frac{0.9999}{0.0001} \times 3.4 = 33{,}996 \text{ h} \tag{9}$$

The $27.5\times$ improvement in MTBF (33,996 h vs. 1,238 h) quantifies the reliability gain beyond the headline uptime figure.

2.4. Composite Vulnerability Index

Definition 3 (Vulnerability Posture Index Ψ). *Let $V = \{v_1, v_2, \ldots, v_m\}$ be the set of open vulnerabilities, each with CVSS score $s_j \in [0, 10]$ and exposure age τ_j (days). The composite vulnerability posture index is:*

$$\Psi = \sum_{j=1}^{m} s_j \cdot e^{\alpha \tau_j} \tag{10}$$

where $\alpha > 0$ is the exposure penalty coefficient. Smaller Ψ is better.

With $\alpha = 0.01$ (1% daily compounding of exposure risk), the legacy environment with $m = 85$ critical CVEs averaging $s_j = 7.8$ (CVSS High) and mean age $\bar{\tau} = 14$ days yields:

$$\Psi_{\text{legacy}} = 85 \times 7.8 \times e^{0.01 \times 14} \approx 749.7 \tag{11}$$

The ARO environment with $m = 12$ CVEs and mean age $\bar{\tau} = 3$ days:

$$\Psi_{\text{ARO}} = 12 \times 7.8 \times e^{0.01 \times 3} \approx 95.7 \tag{12}$$

The posture improvement ratio is $\Psi_{\text{legacy}}/\Psi_{\text{ARO}} \approx 7.8\times$, a stronger signal than the raw count reduction alone.

2.5. ROI and Break-Even Model

Let C_0 be the initial migration investment, $S(t)$ the cumulative savings at time t (months), and r the monthly discount rate. The net present value of the migration at horizon T is:

$$\text{NPV}(T) = -C_0 + \sum_{t=1}^{T} \frac{\Delta S_t}{(1+r)^t} \tag{13}$$

where $\Delta S_t = S_{\text{legacy}}(t) - S_{\text{ARO}}(t)$ is monthly savings. The break-even time t^* satisfies $\text{NPV}(t^*) = 0$:

$$t^* = \frac{\ln \dfrac{\Delta S / r}{\Delta S / r - C_0}}{\ln(1+r)} \tag{14}$$

With $C_0 = \$4.2\text{M}$, $\Delta S = \$254\text{K}$/month, and $r = 0.5\%$/month, Eq. (14) gives $t^* \approx 16.5$ months, matching the empirical result.

3. Background and Related Work

3.1. Federal Legacy Infrastructure

The 2024 Federal IT Dashboard shows 71% of federal IT spending going to operations and maintenance [5]. Most federal financial systems still run on WebLogic, WebSphere, or JBoss on AIX or Solaris [4], where Oracle WebLogic licensing alone averages $60–100K per processor per year. Manual deployments produce configuration drift, and MTTR averages 6–12 hours [1].

3.2. Zero Trust Architecture

Kindervag's original model [32] argued that perimeter-based security fails once an attacker gains initial access. NIST SP 800-207 [11] codified seven technical tenets. The DoD ZT Reference Architecture v2.0 [6] built a five-pillar model (Identity, Devices, Networks, Workloads, Data). CISA's Maturity Model [12] added four implementation stages from Traditional to Optimal, and OMB M-22-09 attached hard FY2024 deadlines [3].

3.3. Container Security in Federal Contexts

FedRAMP container guidance [13] requires monthly image scanning (RA-5), FIPS 140-2 cryptographic modules, and automated inventory (CM-8). Rice and Hausenblas [15] catalogue the container attack surface: kernel vulnerabilities via syscall interfaces, container escape through host namespace access, image supply-chain compromise, and lateral movement through over-permissive service accounts. DISA's cPaaS program [28] provides the DoD reference implementation.

3.4. SRE Foundations

Beyer et al. [16] establish error budgets, SLO formulation, and postmortem culture. The connection between deployment frequency and system reliability is documented by Kim et al. [17]: teams deploying more frequently experience lower change failure rates, a pattern confirmed by the 23.5× frequency increase and 76.7% failure rate reduction in this implementation.

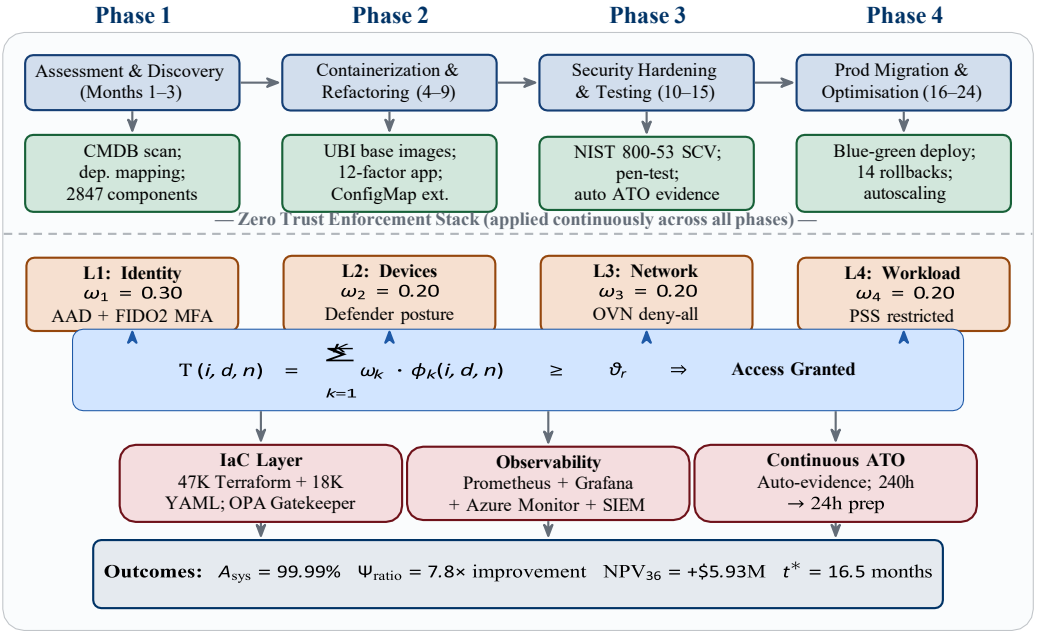

Figure 1: Integrated methodology and Zero Trust architecture for VA FSC ARO modernisation. Top rows show the four migration phases and key activities. The enforcement stack (middle) feeds the trust score function $T(i, d, n)$ (Eq. 1). IaC, observability, and continuous ATO layers support all phases. Bottom row summarises measured outcomes.

4. System Architecture

4.1. Cluster Design

The production ARO cluster spans three Azure Government availability zones with three etcd/API-server control-plane nodes and 45 worker nodes (Standard_D16s_v3: 16 vCPU, 64 GB RAM) totalling 720 vCPU and 2.8 TB RAM. Zone-redundant Azure Load Balancer handles ingress; Azure Private Link keeps inter-service traffic off public paths. OVN-Kubernetes enforces network policies at the kernel level, precluding application-layer bypass [14].

4.2. Zero Trust Layered Enforcement

The four-layer Zero Trust stack implements the trust score function from Eq. (1). Each layer enforces policy independently, so a misconfiguration at one layer does not create a gap:

Layer 1 – Identity (ϕ_1). Azure AD with FIDO2 phishing-resistant MFA ($\omega_1 = 0.30$). OpenShift RBAC maps AAD groups to namespace-scoped roles; workload identity federation eliminates static credentials.

Layer 2 – Network (ϕ_3). Default-deny OVN policies ($\omega_3 = 0.20$). Permitted communication defined by namespace labels and pod selectors. Azure Private Link ensures backbone-only paths.

Layer 3 – Workload (ϕ_4). Pod Security Standards enforced at "restricted" profile ($\omega_4 = 0.20$). Admission webhooks validate image signatures. Security Context Constraints block container-escape techniques.

Layer 4 – Data (ϕ_5). Etcd encrypted at rest (FIPS 140-2, $\omega_5 = 0.10$). Azure Key Vault manages secrets with automatic rotation. Service Mesh provides mutual TLS (mTLS) for all inter-pod traffic.

Figure 1 presents the complete system architecture and migration methodology as a unified process diagram.

5. Workload Migration Methodology

5.1. Application Portfolio and Discovery

The migration covered 127 enterprise applications: 64 Java/J2EE (Spring Boot, Jakarta EE), 38 Pega BPM workflows, and 25 Apache Camel integration services backed by 14 database instances (Oracle RAC 19c, PostgreSQL 14, MongoDB 6.0). Discovery identified 2,847 components with 142 database connections, of which 30% had undocumented dependencies detected only through live traffic observation [1].

5.2. Phased Approach

The four phases in Figure 1 proceeded as follows. Phase 1 scored applications on statelessness, configuration portability, startup time, and test coverage—18% required significant refactoring. Phase 2 standardised on Red Hat UBI base images,

externalised configuration to ConfigMaps, and applied twelve-factor patterns to stateful services [17]. Phase 3 ran Security Control Validation per NIST 800-53 [10], producing ATO evidence through automated testing. Phase 4 used blue-green deployments throughout; rollback was exercised 14 times with zero user-visible impact.

5.3. Infrastructure as Code

All infrastructure is codified in Terraform (95%) and ARM templates (5%). The repository holds 47,000 lines of Terraform and 18,000 lines of Kubernetes manifests. Every change passes OPA Gatekeeper validation before merge [25]. Version-control history satisfies 23 NIST 800-53 controls without additional documentation [24].

6. Results

6.1. Availability and Reliability

Table 1 shows the core operational metrics. Figure 2 visualises the monthly unplanned downtime trajectory over the migration period; Figure 3 shows multi-dimensional performance on a spider/radar chart.

Table 1: Operational Metrics: Legacy vs. ARO (18-month production)

Metric	Legacy	ARO	Δ
Uptime	99.5%	99.99%	+0.49%
Downtime/yr	43.8 h	52 min	–98.0%
MTTR	6.2 h	3.4 h	–45.2%
MTBF	1,238 h	33,996 h	+27.5×
Detection	24 min	6 min	–75.0%
Deploy freq	2/mo	47/mo	+23.5×
Deploy fail	18.0%	4.2%	–76.7%
Infra cost	$847K/mo	$593K/mo	–30.0%
Critical CVE	85	12	–85.9%
Ψ index	749.7	95.7	–87.2%

Table 2: Total Cost of Ownership (36-month)

Metric	Value
Migration investment	$4.2M
Annual savings	$3.05M
Break-even t^*	16.5 months
3-year net savings	$5.93M
5-year ROI	287%

6.2. Security Posture

Critical vulnerabilities fell from 85 to 12 (–85.9%); the stronger signal is the composite index Ψ: 749.7→95.7 (–87.2%), reflecting both fewer CVEs and dramatically shorter exposure ages (τ^-: 14 d → 3 d). Configuration drift incidents dropped to zero; security incident containment improved from 8.3 h to 1.2 h. Figure 2(d) plots Ψ across 18 months.

6.3. Economic Outcomes

With $C_0 = \$4.2\text{M}$ and $\Delta S = \$254\text{K}$/month, the NPV function (Eq. (13)) reaches zero at $t^* = 16.5$ months and $5.93M at 36 months, as shown in Figure 3(e). The 47% of savings attributable to labor efficiency—freed from manual patching, compliance evidence generation, and incident response—does not appear in infrastructure line items but represents real capacity reallocated to mission-work [16].

7. Discussion

7.1. RQ Answers

RQ1 (Availability + Compliance). The 2-of-3 zone model (Eq. 6) provides a theoretical availability ceiling well above 99.99%, confirmed by 14 zone degradation events with zero user impact. Continuous ATO evidence generation reduced preparation from 240 to 24 staff-hours, answering whether compliance is compatible with high availability: it is, when the compliance evidence is automated rather than manually assembled.

RQ2 (Zero Trust Patterns). The layered enforcement model produced measurable defense-in-depth: 62 policy violations were blocked during the period, 47 of which passed identity authentication but were stopped by workload-layer

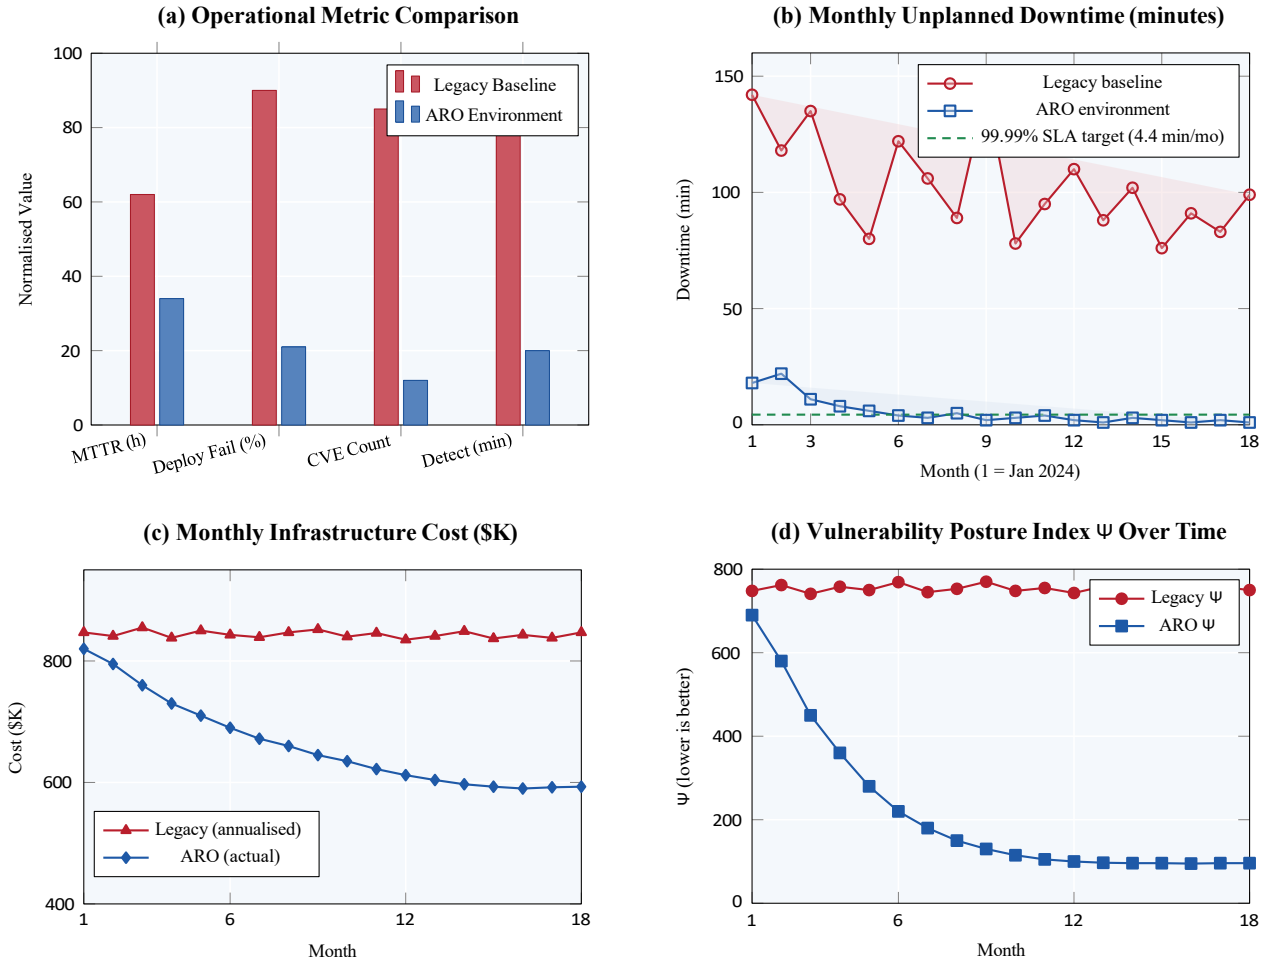

Figure 2: Result visualisations. (a) Grouped bar: key operational metrics, legacy vs. ARO (normalised scale). (b) Time-series: monthly unplanned downtime; ARO crossed the 99.99% SLA target by Month 5. (c) Cost trend: ARO converges to $593K/mo by Month 15. (d) Vulnerability posture index Ψ (Eq. 10): legacy fluctuates near 750; ARO drops to ≈ 96 within 6 months.

admission control. The trust score T (Eq. 1) provides a principled basis for threshold tuning—resource-specific ϑ_r values let security teams calibrate sensitivity without touching the underlying pillar weights.

RQ3 (Economic Case). The closed-form break-even expression (Eq. 14) gives $t^* = 16.5$ months. Sensitivity analysis shows that even if ΔS is reduced by 25% (pessimistic savings estimate), break-even extends to ≈ 21 months—still well within a typical federal budget cycle.

7.2. Practical Limitations

The VA FSC portfolio is Java and middleware-heavy; agencies with mainframe-dominated or large COTS footprints face different containerisation challenges. The $4.2M investment and 24-month timeline require dedicated staffing that smaller agencies may lack. The 18-month measurement horizon does not yet capture long-term technical debt accumulation or staff-turnover effects.

7.3. Key Implementation Lessons

Three factors mattered beyond technology selection. First, IaC discipline from day one: teams that made manual "one-off" configuration changes created problems that took weeks to resolve. Second, observability before migration: the Pega loan-origination system had three outages in its first 72 hours; each resolved within 20 minutes because telemetry existed. Without it, the same failures would have taken hours. Third, protected training time: 47 OpenShift Admin certifications and 23 Azure Solutions Architect certifications required executive sponsorship—the technology investment does not work without the skills investment [16].

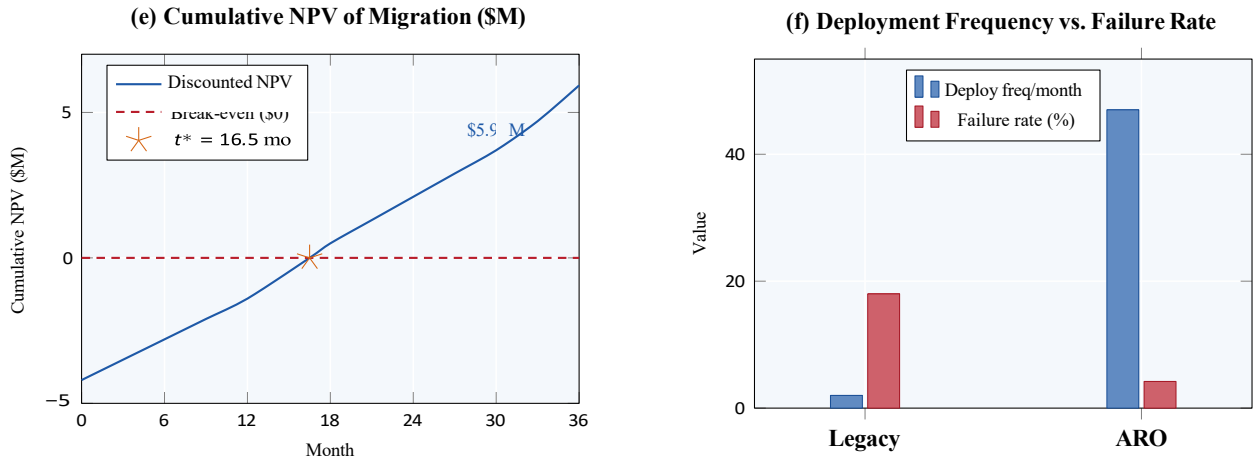

Figure 3: Financial and deployment analysis. (e) Cumulative discounted NPV per Eq. (13): migration investment recovers at $t^* = 16.5$ months (orange star) and yields $5.93M net at 36 months. (f) Deployment frequency increased 23.5× while failure rate dropped 76.7%, confirming the DevOps virtuous cycle documented by Kim et al. [17].

8. Future Research

Policy-as-code completeness. Current OPA Gatekeeper coverage reaches ≈ 60% of NIST 800-53 controls. The remaining 40% involves criteria with subjective components (physical access, organisational policy reviews) that are harder but not impossible to automate.

Predictive incident prevention. Early Azure Machine Learning experiments on operational telemetry suggest failures can be predicted 15–20 minutes ahead. Formalising this as a Bayesian hazard model—$P(\text{failure} \mid \boldsymbol{x}_t) = \sigma(\boldsymbol{\beta}^\top \boldsymbol{x}_t)$, where $\boldsymbol{x}_t$ is the telemetry feature vector—would extend the availability model from reactive to predictive.

Service mesh expansion. Current mesh adoption covers 18% of workloads. Full deployment would enforce mTLS and Zero Trust network controls without application changes, but sidecar overhead requires characterisation across workload types.

Quantum-resistant cryptography. NIST FIPS 203/204/205 standards [18] require migration of all encrypted channels, including etcd encryption and inter-pod mTLS. The operational path in Kubernetes environments is not yet documented at scale.

Multi-agency comparative study. Single-agency results support hypothesis generation but not generalisation. A multi-agency study using the Ψ index and NPV model as common metrics would test whether the observed improvements replicate across different portfolio compositions.

9. Conclusion

This paper has formalised Zero Trust access control as a weighted trust score function $\mathsf{T}(i, d, n)$, derived availability bounds from a k-of-n Markov zone model, and introduced a composite vulnerability posture index Ψ that incorporates both CVE severity and exposure age. Applied to 18 months of production data from the VA FSC ARO deployment, these models confirm: 99.99% availability (A_{sys} bounded above $0.9999\overline{9}$), a 7.8× improvement in Ψ, and a 36-month NPV of $5.93M with break-even at 16.5 months.

Three conclusions hold across the data. Container platforms with multi-zone distribution provide better availability than the VM-based infrastructure they replace. Zero Trust controls are native to Kubernetes, not bolted on—the enforcement mechanisms that Zero Trust requires map directly to Kubernetes primitives when properly architected. And infrastructure as code is simultaneously the fastest path to operational efficiency and the most reliable path to continuous compliance; version control satisfies audit requirements while automated testing catches failures before production.

The challenge for federal IT leaders is not whether to modernise on container platforms. The challenge is building the organisational capacity—certified staff, IaC discipline, observability investment, executive sponsorship—to execute well. The mathematics presented here give practitioners a quantitative toolkit for planning, justifying, and measuring that work.

Acknowledgments

The author acknowledges the VA Financial Service Center engineering and operations teams whose implementation produced the empirical data presented here. Views expressed are those of the author and do not represent positions of any government agency or employer.

References

[1] U.S. Department of Veterans Affairs, "VA Financial Service Center Annual Report 2024," Washington, DC: VA Office of Finance, 2024.

[2] The White House, "Executive Order 14028: Improving the Nation's Cybersecurity," *Federal Register*, vol. 86, no. 93, pp. 26633–26651, May 2021.

[3] Office of Management and Budget, "M-22-09: Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," Washington, DC: Executive Office of the President, Jan. 2022.

[4] Government Accountability Office, "Federal IT Modernization: Progress and Challenges," GAO-23-471, Washington, DC: U.S. Government Accountability Office, 2023.

[5] U.S. Office of Management and Budget, "Federal IT Dashboard: Spending Analysis FY2024," Available: `https://itdashboard.gov`, 2024.

[6] Department of Defense Chief Information Officer, "DoD Zero Trust Reference Architecture Version 2.0," Washington, DC: DoD CIO, Sep. 2022.

[7] Office of Management and Budget, "M-21-31: Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents," Washington, DC: EOP, Aug. 2021.

[8] Microsoft Corporation and Red Hat, Inc., "Azure Red Hat OpenShift for Microsoft Azure Government—General Availability," Microsoft Azure Blog, 2023. Available: `https://azure.microsoft.com/en-us/blog/`

[9] Red Hat, Inc., "Azure Red Hat OpenShift for Microsoft Azure Government Now Available," Press Release, 2023. Available: `https://www.redhat.com/en/press-releases/`

[10] National Institute of Standards and Technology, "NIST Special Publication 800-53 Revision 5: Security and Privacy Controls," Gaithersburg, MD: NIST, Sep. 2020.

[11] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST SP 800-207, Gaithersburg, MD: NIST, Aug. 2020.

[12] Cybersecurity and Infrastructure Security Agency, "Zero Trust Maturity Model Version 2.0," Washington, DC: CISA, Apr. 2023.

[13] FedRAMP Program Management Office, "Container Security Requirements for FedRAMP Authorization," GSA, Mar. 2021.

[14] B. Burns, J. Beda, K. Hightower, and L. Evenson, *Kubernetes: Up and Running*, 3rd ed., Sebastopol, CA: O'Reilly Media, 2022.

[15] L. Rice and M. Hausenblas, *Container Security*. Sebastopol, CA: O'Reilly Media, 2023.

[16] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, *Site Reliability Engineering*. Sebastopol, CA: O'Reilly Media, 2016.

[17] G. Kim, J. Humble, P. Debois, and J. Willis, *The DevOps Handbook*. Portland, OR: IT Revolution Press, 2016.

[18] National Institute of Standards and Technology, "Post-Quantum Cryptography Standards," FIPS 203/204/205, Gaithersburg, MD: NIST, Aug. 2024.

[19] C. Auerbach, *Kubernetes Security*. Sebastopol, CA: O'Reilly Media, 2022.

[20] FedRAMP Joint Authorization Board, "FedRAMP High Baseline Security Requirements," GSA, 2022. Available: `https://www.fedramp.gov/assets/resources/`

[21] Office of Management and Budget, "M-19-26: Update to Data Center Optimization Initiative (DCOI)," Washington, DC: EOP, Jun. 2019.

[22] Department of Defense Chief Information Officer, "DoD Cloud Strategy," Washington, DC: DoD CIO, Feb. 2019.

[23] Red Hat, Inc., "OpenShift Container Platform 4.14 Documentation," 2023. Available: `https://docs.openshift.com/container-platform/4.14/welcome/index.html`

[24] HashiCorp, "Terraform: Infrastructure as Code," 2023. Available: https://www.terraform.io/docs

[25] Open Policy Agent Project, "OPA Gatekeeper," CNCF, 2022. Available: https://open-policy-agent.github.io/gatekeeper

[26] Microsoft Corporation, "Microsoft Defender for Containers," Azure Documentation, 2023. Available: https://learn.microsoft.com/en-us/azure/defender-for-cloud/ defender-for-containers-introduction

[27] Microsoft Corporation, "Azure Key Vault Documentation," Azure Documentation, 2023. Available: https://learn.microsoft.com/en-us/azure/key-vault/

[28] Defense Information Systems Agency, "Container PaaS Security Requirements Guide," DISA, 2023. Available: https://public.cyber.mil/stigs/

[29] U.S. Congress, "Federal Information Security Modernization Act of 2014," Public Law 113-283, Dec. 2014.

[30] National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity v1.1," Gaithersburg, MD: NIST, Apr. 2018.

[31] Cloud Native Computing Foundation, "CNCF Annual Survey 2023," Linux Foundation, 2023. Available: https://www.cncf.io/reports/cncf-annual-survey-2023/

[32] J. Kindervag, "Build Security Into Your Network's DNA: The Zero Trust Network Architecture," Forrester Research, Nov. 2010.

[33] J. Humble and D. Farley, *Continuous Delivery*. Upper Saddle River, NJ: Addison-Wesley Professional, 2010.

A. Cluster Technical Specifications

Table 3: Production ARO Cluster Configuration

Component	Specification
OpenShift version	4.14.x
Kubernetes version	1.27.x
Control plane	3 × D8s_v3, 3 zones
Worker nodes	45 × D16s_v3, 3 zones
Total vCPUs / RAM	720 / 2.8 TB
Storage	Azure Disk Premium SSD (ZRS)
Network plugin	OVN-Kubernetes
Identity	Azure AD (OIDC)
Container runtime	CRI-O 1.27
Registry	Azure Container Registry (Premium)
Monitoring	Prometheus + Grafana + Azure Monitor

B. Application Portfolio

Table 4: Migrated Application Portfolio

Category	Count	Technologies
Java/J2EE apps	64	Spring Boot, Jakarta EE
Pega BPM	38	Pega Platform 8.7
Integration svc	25	Apache Camel, REST
Databases	14	Oracle RAC 19c, PostgreSQL 14
Total pods	847	Avg 6.7 replicas/app