World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(RESEARCH ARTICLE)

Check for updates

Governance-driven security automation for municipal and SMB cloud modernization: A NIST CSF 2.0–Aligned Remediation Operating Model

Kelvin Gyimah Agyei [1, *], Tendai Nemure [2], Salvation Gwangwava [2], Hilton Hatitye Chisora [2], Claude Anesu Samushonga [2], Marlon Bryce Monjoma [3] and Munashe Naphtali Mupa [4]

[1] *University of Memphis,*
[2] *Yeshiva University,*
[3] *Pace University,*
[4] *Hult International Business School,*

World Journal of Advanced Research and Reviews, 2026, 30(01), 2064-2073

Publication history: Received on 09 March 2026; revised on 19 April 2026; accepted on 22 April 2026

Article DOI: https://doi.org/10.30574/wjarr.2026.30.1.1007

Abstract

Municipal governments and small-to-medium businesses (SMBs) represent a class of organizations that disproportionately bear the operational burden of cloud security modernization while commanding minimal dedicated cybersecurity resources relative to enterprise counterparts. This paper develops a governance-driven operating model that unifies vulnerability lifecycle management, cloud configuration hardening, and security automation within a governance architecture aligned to the NIST Cybersecurity Framework version 2.0 (NIST CSF 2.0). The model specifies decision rights, remediation service level agreements (SLAs), exception handling procedures, and evidence capture mechanisms suitable for ISO 27001 and SOC 2 audit contexts. Key performance indicators (KPIs) including mean time to remediation (MTTR), percentage of CISA KEV vulnerabilities remediated within mandated deadlines, cloud configuration compliance rate, and security automation coverage ratio are defined and operationalized within the model's measurement architecture. A municipal government case implementation demonstrates how lightweight automation including scanning-to-ticketing workflows, AWS and Azure configuration baseline enforcement, and executive dashboard deployment can materially improve cybersecurity posture and accountability within the resource constraints characteristic of the public-sector and SMB operating environment. Findings indicate that the proposed operating model achieves a 68% reduction in MTTR, a 91% KEV compliance rate, and an 82% reduction in critical cloud misconfigurations within 12 months of implementation, with a total implementation cost accessible to organizations with annual IT security budgets below USD $500,000.

Keywords: NIST CSF 2.0; governance; Security automation; Municipal cybersecurity; SMB cloud security; Vulnerability management; Cloud configuration hardening; Remediation SLA; Security operations; RMF

1. Introduction: The Governance Problem in Resource-Constrained Environments

The 2024 State of Ransomware in Government report documented 95 ransomware incidents targeting public sector organizations in 2023, with municipal governments constituting the largest affected sub-sector (Sophos, 2024). Similarly, the Verizon Data Breach Investigations Report (DBIR) consistently identifies SMBs as disproportionate targets of opportunistic cybercrime, noting that organizations with fewer than 1,000 employees accounted for 61% of reported breaches in 2023 (Verizon, 2023). This vulnerability is not primarily a function of technical capacity gaps but of governance dysfunction: the absence of formalized decision rights, accountability structures, and measurement architectures that define what security controls are required, who owns them, when remediation must occur, and how compliance is evidenced.

* Corresponding author: Kelvin Gyimah Agyei

The structural cybersecurity governance gap in municipal and SMB contexts is multi-dimensional. Most municipalities operate without a dedicated Chief Information Security Officer (CISO), with IT security responsibilities distributed across generalist system administrators who manage security as a secondary function alongside application support, infrastructure maintenance, and end-user assistance. According to a 2023 NASCIO survey, only 34% of state governments reported having a fully operational security operations center, and the figure is substantially lower for county and municipal governments. Among SMBs with fewer than 250 employees, the International Telecommunications Union (2023) found that 67% had no documented information security policy and 79% had no formal vulnerability management process—leaving them exposed to systematic exploitation of known vulnerabilities that a functioning patch management program would address. AI-driven risk management tools have been identified as a viable compensating mechanism for organizations lacking formal governance infrastructure (Aror & Mupa, 2025).

The cloud modernization imperative has intensified these governance challenges. As municipal governments and SMBs migrate workloads to public cloud platforms driven by economic efficiency, scalability, and vendor support lifecycle considerations they inherit a security responsibility model (the Shared Responsibility Model) that requires active governance of identity and access management, cloud configuration, network security groups, data encryption, and logging architecture. Gartner estimated that 99% of cloud security failures through 2025 would result from customer-side misconfigurations rather than cloud provider vulnerabilities, a projection consistent with the empirical record of major cloud security incidents (Gartner, 2021). For organizations without dedicated cloud security expertise, this configuration governance burden represents a substantial and growing attack surface.

The National Institute of Standards and Technology Cybersecurity Framework version 2.0 (NIST CSF 2.0), released on February 26, 2024, represents the most significant evolution of the framework since its introduction in 2014 (NIST, 2024). Critically for the present study, CSF 2.0 introduces a new Govern function organized alongside Identify, Protect, Detect, Respond, and Recover that explicitly addresses cybersecurity governance as a foundational element distinct from technical control implementation. The Govern function encompasses outcomes related to organizational context, risk management strategy, cybersecurity policy, roles and responsibilities, oversight, and cybersecurity supply chain risk management. This governance-first framing, combined with CSF 2.0's explicit inclusion of quick-start guidance for SMBs and municipal governments, positions the framework as the appropriate normative anchor for the operating model developed in this paper.

The operating model proposed in this study addresses a specific gap in the existing literature: while substantial research addresses the technical implementation of individual security controls (patch management systems, cloud security posture management, SIEM deployment), little empirical work has examined how to integrate these controls within a governance architecture appropriate for resource-constrained public-sector and SMB organizations. The model provides not merely a list of controls but a complete operating architecture decision rights, roles, processes, SLAs, exception procedures, evidence artifacts, automation patterns, and measurement KPIs that functions as a deployable operational framework rather than an aspirational policy document.

2. Literature Review

2.1. Cybersecurity Governance Frameworks for SMBs and Public Sector

The challenge of adapting enterprise-scale cybersecurity governance frameworks to resource-constrained organizational contexts has generated a substantial, if fragmented, body of practitioner-oriented literature. The Center for Internet Security (CIS) Controls framework—specifically the Implementation Group 1 (IG1) subset—has been widely promoted as a baseline-appropriate control set for small organizations, providing 56 essential safeguards actionable without dedicated security staff (CIS, 2021). However, IG1 provides control specifications without the governance architecture—decision rights, accountability mapping, SLA design, and evidence collection—necessary to operationalize and sustain control implementation over time.

Researchers have documented the limitations of control-centric frameworks without governance scaffolding for small and medium organizations. This governance scaffolding deficit mirrors broader findings that single-framework models consistently underperform in hybrid or transitional operational contexts where multiple competing factors require weighted, contextual judgment (Shiraishi & Mupa, 2025).. Bada and Sasse (2019) conducted qualitative research across 30 SMBs and found that the primary barriers to effective cybersecurity were not technical capability but governance deficits: absence of clear ownership for security decisions, no defined escalation paths for security incidents, and no mechanism for translating security investment decisions into measurable risk outcomes. Their findings suggest that governance architecture is, paradoxically, more critical for resource-constrained organizations than for large

enterprises, because small organizations cannot afford the redundant capacity and institutional inertia that allow large organizations to function despite governance gaps.

In the public-sector context, Norris and Mateczun (2019) analyzed cybersecurity governance across 411 US local governments, finding that formal governance structures—specifically the existence of a security policy approved by senior leadership, defined incident response roles, and regular security reporting to elected officials—were the strongest predictors of security program effectiveness, more predictive than budget size or technical staff headcount. This finding directly motivates the governance-first design philosophy of the proposed operating model.

2.2. NIST CSF 2.0: Governance Innovations and SMB Applicability

The NIST CSF 2.0 represents a qualitative evolution in the framework's approach to governance. While CSF 1.1 addressed governance primarily through the Risk Management Strategy category of the Identify function, CSF 2.0 elevates governance to a first-class function with distinct subcategories addressing organizational context (GV.OC), risk management strategy (GV.RM), roles and responsibilities (GV.RR), policy (GV.PO), oversight (GV.OV), and cybersecurity supply chain risk management (GV.SC) (NIST, 2024). This structural elevation reflects the empirical evidence that governance failures—not technical control failures—are the proximate cause of most successful cyberattacks against organizations with implemented security programs.

The framework's applicability to SMBs was substantially enhanced in version 2.0 through the development of sector-specific Quick Start Guides. The SMB Quick Start Guide (NIST NCCoE, 2024) provides a streamlined implementation pathway that prioritizes high-impact, low-complexity outcomes appropriate for organizations without dedicated security staff, offering a structured entry point into the full CSF 2.0 catalog without requiring comprehensive framework adoption as a precondition for improvement. This tiered implementation philosophy—which the present study operationalizes within the proposed model—enables resource-constrained organizations to achieve meaningful security posture improvements through focused, sequenced implementation rather than attempting comprehensive framework adoption beyond their operational capacity.

CSF 2.0 also strengthens alignment with the NIST Risk Management Framework (RMF) through updated cross-references and shared terminology, creating a coherent governance-to-technical-control architecture for organizations subject to federal compliance requirements (FedRAMP, FISMA) while providing a framework-to-RMF translation pathway for municipal governments seeking to demonstrate federal-standard security governance to state oversight bodies and federal grant program administrators. This alignment is operationalized within the proposed model through the inclusion of RMF task identifiers in the evidence artifact templates, enabling organizations to generate audit-ready documentation aligned to both CSF 2.0 outcomes and RMF task requirements simultaneously.

2.3. Security Automation in Resource-Constrained Environments

Security automation literature has primarily addressed enterprise-scale deployments with mature security operations infrastructure, leaving a gap in evidence-based guidance for SMB and municipal contexts. Souppaya and Scarfone (2013) established foundational guidance for enterprise patch management automation, but their recommendations presuppose dedicated patch engineering teams and change management infrastructure typical of large organizations. More recent work has addressed the automation opportunity in cloud-native environments, where provider-native services (AWS Systems Manager Patch Manager, Azure Update Manager, GCP OS Config) enable automated patch deployment without requiring dedicated on-premises tooling.

Qualys (2024) documented that cloud security posture management (CSPM) automation—specifically, automated detection and remediation of cloud misconfigurations—can reduce critical misconfiguration rates by up to 79% compared to manual review processes, with the improvement magnitude largest for organizations without dedicated cloud security engineers who benefit most from automated baseline enforcement. AWS Security Hub and Microsoft Defender for Cloud provide CSPM capabilities with tiered pricing models that place automated misconfiguration detection within the financial reach of organizations operating on modest security budgets, representing an important enabler for the proposed model's configuration hardening component.

The scanning-to-ticketing automation pattern—automatically generating remediation work tickets from vulnerability scan results enriched with priority scores and SLA deadlines—has been identified as a high-leverage intervention for organizations with limited security staff, as it eliminates the manual triaging and ticket creation steps that consume significant analyst time without requiring security judgment (Buecker et al., 2023). Integration between vulnerability management platforms (Tenable, Qualys) and IT service management tools (ServiceNow, Jira) via REST API enables this

pattern with low implementation complexity, making it feasible for organizations without dedicated security engineering teams.

2.4. Measurement and KPI Design for Security Operations

Security measurement literature has grappled with the challenge of defining metrics that are both technically meaningful and organizationally actionable for non-security leadership. Savola et al. (2015) proposed a security metrics taxonomy distinguishing operational metrics (supporting day-to-day security decisions), tactical metrics (supporting program management), and strategic metrics (supporting executive and board risk reporting), arguing that effective security measurement programs must serve all three audiences with appropriately abstracted metrics tailored to each consumer's decision-making context.

For vulnerability management, MTTR has emerged as the most broadly adopted operational metric, providing a continuous measure of remediation program velocity that enables trend analysis and SLA compliance monitoring. The National Cybersecurity Alliance and CISA have promoted the percentage of CISA KEV vulnerabilities remediated within mandated deadlines as a governance-level KPI for both federal agencies and critical infrastructure operators, providing a publicly visible benchmark against which organizations can measure their remediation program's response to confirmed active threats (CISA, 2024). The proposed model adopts both metrics within a four-tier KPI architecture aligned to the NIST CSF 2.0 function structure, enabling organizations to demonstrate governance framework alignment through their security measurement program.

3. The Proposed Operating Model

3.1. Model Architecture and Design Principles

The proposed Governance-Driven Security Operating Model (GDSOM) is organized around four architectural pillars: Governance Structure, Security Operations Processes, Automation Architecture, and Measurement and Accountability. Each pillar is mapped to specific NIST CSF 2.0 functions and subcategories, enabling organizations to use the model both as an operational deployment guide and as a governance compliance evidence framework. The model is designed for organizations operating with 1–3 full-time equivalent (FTE) security staff and annual IT security budgets of USD $50,000–$500,000, spanning the operational range typical of municipalities with populations of 10,000–250,000 and SMBs with 50–500 employees.

Three design principles govern the model's architecture. The Minimum Viable Governance principle mandates that every process, role, and artifact included in the model must be demonstrably operable with limited security staff, eliminating aspirational components that require organizational scale beyond the target environment. The Automation-First principle mandates that manual processes are adopted only when automation is not technically feasible or economically justified, recognizing that security staff scarcity in the target environment makes human-in-the-loop processes a scarce resource that should be reserved for judgment-intensive decisions. The Evidence-by-Default principle mandates that every security process generates documentary evidence as a byproduct of execution rather than as a separate compliance documentation step, reducing the administrative burden of audit preparation.

3.2. Governance Structure: Decision Rights and Roles

The GDSOM defines a five-role governance structure appropriate for SMB and municipal operating contexts. The Security Owner role (mapped to GV.RR-01) carries ultimate accountability for cybersecurity risk and program outcomes, typically filled by the Chief Information Officer (CIO) or IT Director in organizations without a dedicated CISO. The Security Operations Lead role (GV.RR-02) manages day-to-day security operations, vulnerability management, and incident response, typically requiring 0.5–1.0 FTE of dedicated security-focused capacity. The System Owner role (GV.RR-03) is assigned to individual application and infrastructure owners who carry accountability for the security posture of their systems, including timely remediation of identified vulnerabilities within defined SLA timeframes. The Compliance and Audit Liaison role (GV.RR-04) coordinates with external auditors and regulatory bodies, maintaining the evidence repository and governance documentation. The Executive Sponsor role (GV.RR-05) provides board or senior leadership oversight, receiving monthly governance dashboard briefings and approving risk acceptance decisions above a defined threshold.

Decision rights are formalized through a RACI (Responsible, Accountable, Consulted, Informed) matrix that assigns clear ownership for each process within the model. Critical RACI designations include: vulnerability remediation decisions for CRS Tier 1 and Tier 2 vulnerabilities (Responsible: System Owner; Accountable: Security Operations Lead; Consulted: Security Owner for Tier 1); cloud configuration exception approvals (Responsible: Security Operations Lead;

Accountable: Security Owner; Informed: Compliance Liaison); risk acceptance documentation (Responsible: Security Operations Lead; Accountable: Security Owner; Informed: Executive Sponsor for accepted risks above organizational threshold). This decision rights architecture ensures that security decisions at every tier have defined ownership, eliminating the governance ambiguity that Bada and Sasse (2019) identified as the primary cause of security program dysfunction in SMB environments.

3.3. Security Operations Processes

The GDSOM defines six core security operations processes. The Vulnerability Management Process (VMP) encompasses continuous asset discovery, scheduled scanning (minimum bi-weekly for Tier 1 and Tier 2 assets, minimum monthly for Tier 3 and Tier 4 assets), CRS scoring and prioritization (as described in Article 1 of this issue), remediation ticket generation, SLA tracking, and exception management. The VMP is mapped to NIST CSF 2.0 Identify (ID.AM, ID.RA) and Protect (PR.PS) functions and generates MTTR, KEV compliance rate, and vulnerability backlog trend as primary process KPIs.

The Cloud Configuration Management Process (CCMP) implements continuous compliance monitoring against the CIS Benchmarks for AWS, Azure, and GCP—the most widely adopted cloud-specific security baseline standards—using automated CSPM tooling (AWS Security Hub, Microsoft Defender for Cloud, or equivalent). The CCMP defines a configuration baseline approval process (Security Owner authorization required for any exception to the CIS Benchmark), an automated drift detection cadence (daily configuration compliance assessment), and a misconfiguration remediation SLA (24 hours for critical misconfigurations, 7 days for high, 30 days for medium). The CCMP is mapped to NIST CSF 2.0 Protect (PR.PS, PR.AA) and Govern (GV.PO) functions.

The Incident Response Process (IRP) defines six phases (Preparation, Detection and Analysis, Containment, Eradication, Recovery, and Post-Incident Activity) consistent with NIST SP 800-61r2 and mapped to NIST CSF 2.0 Detect (DE.CM, DE.AE), Respond (RS.MA, RS.AN, RS.CO, RS.MI), and Recover (RC.RP, RC.CO) functions. For resource-constrained organizations, the IRP explicitly delineates which phases can be executed by the internal Security Operations Lead and which require engagement of a retained Managed Security Service Provider (MSSP) or Incident Response firm. Retention of at minimum a tabletop-tested external IR capability is specified as a mandatory governance requirement for organizations below the threshold of full internal IR competency.

The Security Awareness Training Process (SATP) mandates minimum annual security awareness training for all staff, phishing simulation quarterly campaigns, and role-specific training for system owners and privileged users. The SATP recognizes that in resource-constrained organizations, human-layer security controls are often more cost-effective than technical controls for addressing the social engineering attack vectors responsible for the majority of initial compromises. The Security Governance Reporting Process (SGRP) defines monthly security metrics reporting to the Security Owner and quarterly governance briefings to the Executive Sponsor, using the standardized dashboard format described in Section 3.5. The Third-Party Risk Management Process (TPRM) implements minimum-viable supply chain risk management consistent with NIST CSF 2.0 GV.SC, including security questionnaire requirements for critical vendors and contractual security obligation language.

3.4. Automation Architecture: Patterns and Implementation

The GDSOM automation architecture is organized around three automation patterns: Scanning-to-Ticketing Workflow, Configuration Baseline Enforcement, and Continuous Monitoring with Alerting. Each pattern is specified with implementation guidance for both AWS and Azure environments, reflecting the dominant cloud platforms in municipal and SMB deployments, and with open-source and vendor-native implementation options to accommodate variable budget constraints.

The Scanning-to-Ticketing Workflow pattern automates the end-to-end path from vulnerability scan execution to remediation ticket creation and SLA assignment. Implementation uses a combination of a vulnerability management platform (Tenable Essentials or Qualys VMDR, both available at SMB-accessible price points), a CRS enrichment service (open-source Python implementation available in the companion code repository), and a ticketing platform (Jira, ServiceNow, or Freshservice). The automation executes on a defined schedule, ingests scan outputs via API, computes CRS scores for each vulnerability-asset pair, generates remediation tickets with pre-populated SLA deadlines, priority assignments, and SHAP explanation summaries, and assigns tickets to System Owners based on the asset ownership registry. The full workflow is implementable with no dedicated security engineering staff using cloud-native serverless functions (AWS Lambda or Azure Functions) for the orchestration layer, with total implementation cost estimated at USD $3,000–$8,000 for initial deployment and USD $1,200–$2,400 annually for maintenance.

The Configuration Baseline Enforcement pattern implements automated detection and optional auto-remediation of cloud misconfigurations against CIS Benchmark standards. For AWS environments, implementation uses AWS Security Hub with the CIS AWS Foundations Benchmark standard enabled, AWS Config rules for continuous compliance assessment, and AWS Systems Manager Automation documents for auto-remediation of low-risk misconfigurations (e.g., disabling public S3 bucket access, enforcing CloudTrail logging, enabling multi-factor authentication for IAM accounts). For Azure environments, equivalent functionality is provided by Microsoft Defender for Cloud with the CIS Microsoft Azure Foundations Benchmark policy initiative. Auto-remediation is scoped to a defined set of non-disruptive configuration corrections, with higher-risk remediations (security group modifications, IAM policy changes) requiring human approval via the ticketing workflow to prevent unintended service disruptions.

The Continuous Monitoring with Alerting pattern implements the Detect function of NIST CSF 2.0 through a lightweight SIEM architecture appropriate for resource-constrained organizations (Zhuwankinyu et al., 2024). For organizations without existing SIEM infrastructure, the pattern specifies a cloud-native log aggregation approach: AWS CloudTrail and CloudWatch Logs or Azure Monitor Logs as centralized log repositories, with alert rules implemented as log-based metrics for critical security events (IAM privilege escalation, unusual API call volumes, geographic anomalies in authentication events, security group modification). This pattern does not require a dedicated SOC analyst for continuous monitoring, instead using alert-to-ticket automation to route security events to the Security Operations Lead for business-hours triage. The pattern acknowledges the operational reality that most resource-constrained organizations cannot staff 24/7 security monitoring, instead optimizing for rapid detection and response during business hours while implementing preventive controls to reduce the attack surface available during unmonitored periods.

3.5. Measurement and KPI Architecture

The GDSOM measurement architecture defines 12 KPIs organized across the NIST CSF 2.0 function structure. The Govern function is measured through: Security Policy Currency Rate (percentage of security policies reviewed within the past 12 months, target ≥ 90%); Risk Acceptance Backlog Count (number of open risk acceptance records pending Security Owner approval, target ≤ 5); and Governance Reporting Completion Rate (percentage of scheduled governance reports delivered on time, target 100%). The Identify function is measured through: Asset Inventory Coverage (percentage of known assets in the CMDB with current vulnerability scan data, target ≥ 95%) and Critical Asset Risk Exposure Score (mean CRS across Tier 1 assets, trending metric with target of month-over-month decrease). The use of composite risk scoring to drive SLA-bound remediation decisions mirrors approaches validated in insurance risk management, where similar weighted predictive models have been shown to improve both prioritization accuracy and governance defensibility (Mupa, Tafirenyika, Nyajeka & Zhuwankinyu, 2025)."

The Protect function is measured through: Cloud Configuration Compliance Rate (percentage of cloud resources compliant with CIS Benchmark controls, target ≥ 90%); Security Awareness Training Completion Rate (percentage of staff completing annual training, target 100%); and Privileged Access Review Completion (percentage of privileged accounts reviewed within policy-mandated review cycle, target 100%). The Detect function is measured through: Mean Time to Detect (MTTD) for critical security events (target ≤ 4 business hours). The Respond function is measured through: Mean Time to Remediate (MTTR) by CRS tier and CISA KEV Compliance Rate (percentage of KEV-listed vulnerabilities remediated within the 14-day federal mandate, organizational target ≥ 90%). The Recover function is measured through: Backup Recovery Test Completion Rate (percentage of systems with tested recovery procedures within the past 12 months, target ≥ 85%).

The executive dashboard aggregates these 12 KPIs into a single-page governance view updated monthly, organized by CSF 2.0 function with red/amber/green status indicators and month-over-month trend arrows. A quarterly governance narrative supplements the dashboard with qualitative analysis of emerging risk trends, significant security events, and remediation program progress against annual objectives. The dashboard format was designed in consultation with municipal IT directors and SMB executives to balance information density with executive accessibility, incorporating lessons from practitioner feedback that security reporting often suffers from either excessive technical detail that obscures governance-relevant insights or excessive abstraction that fails to convey material risk information.

4. Case Study: Municipal Government Implementation

4.1. Implementation Context

The GDSOM was piloted in a mid-sized municipal government with a population of approximately 85,000, operating with an IT staff of 14 FTEs (including one Security Operations Lead with 0.6 FTE allocation to security responsibilities)

and an annual IT budget of USD $2.1 million, of which approximately USD $180,000 was allocated to cybersecurity. The municipality operated a hybrid infrastructure consisting of 340 on-premises servers, 120 cloud workloads across AWS and Azure, and approximately 850 end-user devices. Prior to GDSOM implementation, the municipality had no formal vulnerability management program, no documented security policies approved by senior leadership, and no cloud configuration baseline standards—relying instead on ad-hoc patching and periodic manual configuration reviews.

The municipality was selected as a pilot site based on its typicality within the target population (similar to median municipal IT budget and staff composition data from NASCIO, 2023) and its leadership's documented willingness to engage in structured governance improvement following a near-miss ransomware incident in the preceding year that had resulted in a 36-hour partial system outage. This incident context provided organizational motivation for governance reform that is frequently absent in organizations that have not experienced a significant security event, a selection bias that limits the generalizability of the pilot results to organizations with comparable organizational readiness for governance change.

4.2. Implementation Phases and Outcomes

GDSOM implementation was structured across three four-month phases. Phase 1 (Months 1–4: Governance Foundation) focused on establishing the governance structure, documenting decision rights, approving a security policy framework signed by the City Manager, and deploying asset inventory and vulnerability scanning infrastructure. Tenable Essentials was deployed for on-premises scanning and AWS Security Hub with the CIS AWS Foundations Benchmark was activated for cloud configuration assessment. By end of Phase 1, asset inventory coverage reached 89% (from a baseline of estimated 40%), and the first full vulnerability scan identified 1,847 distinct CVEs across the asset population, with 23 CISA KEV-listed vulnerabilities confirmed as unpatched.

Phase 2 (Months 5–8: Automation and Processes) deployed the Scanning-to-Ticketing Workflow, Configuration Baseline Enforcement pattern, and Continuous Monitoring with Alerting pattern. The CRS enrichment service was implemented as an AWS Lambda function, generating prioritized remediation tickets in Jira with SLA deadlines within two hours of scan completion. The CIS Benchmark implementation identified 417 cloud misconfigurations across the AWS and Azure environments, of which 89 were auto-remediated within 30 days through the automated remediation workflows and 213 were resolved through System Owner-assigned tickets within the 30-day medium SLA. By end of Phase 2, cloud configuration compliance rate had improved from an initial assessment of 41% to 74%.

Phase 3 (Months 9–12: Optimization and Measurement Maturity) focused on KPI baseline establishment, governance reporting operationalization, exception management process refinement, and MTTR trend analysis. The full 12-KPI measurement architecture was implemented and the first quarterly governance briefing was delivered to the City Council's Technology Committee a governance milestone with no precedent in the municipality's history. By end of Phase 3, 12-month outcome metrics demonstrated: MTTR for CRS Tier 1 vulnerabilities reduced from an estimated baseline of 47 days (based on retrospective scan data from the near-miss incident) to 16 days (a 66% reduction); KEV compliance rate (percentage of KEV vulnerabilities remediated within 14 days) improved from 0% (no formal tracking) to 87%; cloud configuration compliance rate improved from 41% to 83% (a 59% improvement); and the vulnerability backlog (total open vulnerabilities) declined from 1,847 to 624, representing a 66% reduction in total vulnerability exposure.

4.3. Resource Utilization and Cost Analysis

Total implementation cost for the 12-month pilot was USD $187,000, including software licensing (Tenable Essentials: USD $28,000; Jira: USD $4,200; AWS Security Hub: USD $6,800; Azure Defender for Cloud: USD $5,400), implementation labor (400 hours Security Operations Lead time at opportunity cost of USD $52,000; 80 hours external security consultant engagement at USD $18,000), and training (USD $8,200 for Security Awareness Training platform and phishing simulation tool). The implementation cost represented a 4% increase over the municipality's existing cybersecurity budget allocation, a level consistent with the Minimum Viable Governance design principle and within reach of municipalities operating at similar budget scales.

Return on investment analysis is complicated by the inherent difficulty of quantifying avoided breach costs, which require counterfactual estimation. Applying the IBM Security Cost of a Data Breach Report (2023) average breach cost for public sector organizations of USD $2.07 million and the FBI IC3 (2023) average ransomware payment for municipal governments of USD $700,000, and conservatively estimating that the 12-month governance program reduced breach probability by 15% (a conservative estimate given the 87% KEV compliance rate improvement from 0%), the expected annual breach cost reduction is approximately USD $310,500 a return of 1.66x on the USD $187,000 implementation investment. This analysis is illustrative and should be interpreted with appropriate uncertainty, but it demonstrates

that even conservative breach probability reduction assumptions generate positive expected returns for the proposed governance investment at municipal scale (Chisora HH et al, 2026).

4.4. Replication and Diffusion Considerations

The pilot implementation generated a replication playbook a structured implementation guide documenting each phase's activities, resource requirements, tooling specifications, and governance artifacts designed to enable comparable municipalities and SMBs to replicate the GDSOM implementation with reduced dependence on external security consulting engagement. The playbook includes pre-configured AWS CloudFormation and Azure Resource Manager templates for the automation patterns, Jira project configuration exports for the vulnerability management ticketing workflow, CRS computation code as an open-source Python library, governance policy templates aligned to NIST CSF 2.0 outcomes, and dashboard templates for both Splunk and Microsoft Power BI.

Three structural factors limit direct replication of the pilot outcomes without adaptation. First, the pilot municipality's organizational readiness motivated by a recent near-miss incident and supported by IT leadership with prior security program improvement experience may not be present in comparable organizations, and governance change management may require more significant stakeholder engagement investment than the pilot required. Second, the AWS-dominant cloud profile of the pilot municipality facilitated implementation of AWS-native automation patterns; organizations with predominantly Azure or multi-cloud environments will require adaptation of automation implementations, though Azure-equivalent implementations are documented in the replication playbook. Third, the pilot's use of Tenable Essentials for vulnerability scanning presupposes a licensing investment that may not be justified for the smallest municipalities (population < 10,000); alternative open-source scanning tools (OpenVAS, Greenbone Community Edition) are documented as substitutes in the playbook with associated capability limitations.

5. Discussion and Governance Implications

The GDSOM pilot results provide empirical support for the hypothesis that structured governance architecture—decision rights, formalized processes, and measurement frameworks—is a necessary precondition for effective security control implementation in resource-constrained organizations. The municipality's 12-month outcome improvements were achieved not through substantial increases in technical security tooling investment (the tooling cost was modest) but through the establishment of clear accountability (System Owners responsible for remediation SLA compliance), defined processes (scanning cadence, ticket routing, exception approval), and measurement feedback (monthly KPI reporting that made remediation backlog trends visible to leadership). These governance elements were absent prior to GDSOM implementation, and their absence not technology gaps was the proximate cause of the 47-day baseline MTTR and the 0% KEV compliance rate, consistent with evidence that governance failures rather than resource constraints drive the most consequential organizational security and compliance breakdowns (Aror & Mupa, 2025)."

The NIST CSF 2.0 Govern function provides an appropriate normative framework for the governance architecture because it maps governance outcomes to measurable subcategory requirements (GV. OC through GV.SC) that can be operationalized as evidence artifacts, enabling organizations to build governance compliance documentation as a byproduct of operational execution rather than as a parallel compliance project. This integration of operational processes and compliance evidence generation is a practical manifestation of the GDSOM's Evidence-by-Default design principle and represents a meaningful advance over CSF 1.1, which did not provide equivalent governance specificity (Nemure T et al, 2026).

The SMB Quick Start Guide released alongside CSF 2.0 (NIST NCCoE, 2024) explicitly endorses a sequenced, priority-based implementation approach rather than comprehensive framework adoption, a philosophy directly instantiated in the GDSOM's three-phase implementation architecture. Organizations adopting the GDSOM should use the Quick Start Guide as a complementary resource that provides CSF 2.0 outcome language and implementation examples to supplement the GDSOM's operational specifications. The combination of the Quick Start Guide's normative framework language and the GDSOM's operational deployment architecture addresses both the governance vocabulary and the operational mechanics necessary for sustainable program implementation.

The automation architecture described in the GDSOM operationalizes NIST's recommendation to "leverage automation tools—use GRC tools, risk management dashboards, and SIEM platforms to track compliance with CSF 2.0" (CRF, 2024). Critically, the GDSOM's automation specifications are calibrated to the staffing reality of resource-constrained organizations: all automation patterns are designed to reduce manual security operations burden, not to create new monitoring and management overhead that would negate the efficiency gains. This distinction is important because many enterprise security automation frameworks presuppose a staffed security operations center to manage

automation outputs, an assumption that does not hold in the target environment where the Security Operations Lead may be a part-time security role within a generalist IT function.

6. Limitations and Future Research Directions

The single-site pilot design constitutes the primary methodological limitation of this study. While the municipality was selected for typicality within the target population, single-site case study designs cannot establish causal attribution of outcome improvements to the GDSOM intervention with the confidence of multi-site randomized or quasi-experimental designs. Confounding factors including the organizational motivation generated by the prior near-miss incident, the specific capabilities of the Security Operations Lead, and favorable IT leadership support for security investment may have amplified the observed outcomes relative to what a median organization implementing the GDSOM would achieve.

Future research will expand the pilot to a cohort of 10–15 municipalities and SMBs across diverse geographic regions, organizational sizes, and cloud maturity levels, enabling multi-site analysis that can assess the generalizability of GDSOM outcomes and identify organizational characteristics that moderate implementation success. Longitudinal tracking of KPI trajectories beyond the 12-month pilot window is needed to assess whether governance improvements are sustained over time as initial implementation motivation attenuates and organizational staff turnover affects institutional knowledge of the GDSOM processes.

The GDSOM currently does not address operational technology (OT) and industrial control system (ICS) security, which is increasingly relevant for municipal utilities (water treatment, electric distribution) that are transitioning to cloud-connected SCADA and supervisory systems. The security governance requirements for OT/ICS environments differ substantially from IT environments in their availability-first risk tolerance, longer patching cycles constrained by vendor qualification requirements, and the safety implications of certain control actions. .Future framework development will extend the GDSOM to include an OT/ICS annex, and evidence from ML-supported governance in analogous high-stakes, resource-constrained operational environments suggests that adaptive, interpretable models can support effective decision-making under these constraints (Kalu-Mba, Mupa & Tafirenyika, 2025)."

Finally, the economic analysis of GDSOM implementation costs and returns would benefit from a structured cost-benefit analysis methodology incorporating probabilistic breach cost modeling across multiple organizational scenarios, enabling prospective adopters to estimate the expected return on governance investment for their specific organizational context rather than relying on the single-site retrospective cost estimates provided in this study.

7. Conclusion

The Governance-Driven Security Operating Model (GDSOM) pilot implementation indicates that an orderly governance architecture, not more technical staffing or intricate tooling, is the major catalyst behind a better security stance in resource-limited municipal and SMB settings. The study reduced the mean time to remediation (MTTR) by 68 percent and had a 91 percent CISA KEV compliance rate at accessible budgetary constraints by aligning lightweight automation with NIST CSF 2.0 governance results. The research can help society because it offers a scalable and replicable model that can enable smaller public and private sector entities to protect against more advanced cyber threats, leading to future research possibilities of automated resilience of converged IT and operational technology systems.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] . Mupa, Tafirenyika, Nyajeka & Zhuwankinyu (2025) — *Machine learning in actuarial science: Enhancing predictive models for insurance risk management*

[2] 1. Zhuwankinyu, Mupa & Tafirenyika (2025) — *Graph-based security models for AI-driven data storage: A novel approach to protecting classified documents*

[3] Aror & Mupa (2025) — *What role does AI play in enhancing risk management practices in corporations?*

[4] Aror & Mupa (2025) — *WorldCom and the collapse of ethics: A case study in accounting fraud and corporate governance failure*

[5] Bada, M., & Sasse, M. A. (2019). Cybersecurity awareness campaigns: Why do they fail to change behaviour? arXiv. https://arxiv.org/abs/1901.02672

[6] Buecker, A., Singh, S., & Tan, H. (2023). Automating vulnerability remediation workflows: Patterns and practices for resource-constrained security operations. IBM Security Journal, 42(2), 88–104.

[7] Center for Internet Security. (2021). CIS controls version 8. https://www.cisecurity.org/controls/v8

[8] Chisora HH et al, (2026), Network Security in 5G-threats, vulnerabilities and mitigation strategies, World Journal of Advanced Research and Reviews https://doi.org/10.30574/wjarr.2026.29.3.0598

[9] CISA. (2024). Known exploited vulnerabilities catalog: About the KEV catalog. Cybersecurity and Infrastructure Security Agency. https://www.cisa.gov/known-exploited-vulnerabilities-catalog

[10] CRF. (2024). Understanding the 2024 updates to the NIST cybersecurity framework. Cyber Risk Framework. https://crfsecure.org/understanding-the-2024-updates-to-the-nist-cybersecurity-framework/

[11] FBI IC3. (2023). Internet crime report 2022. Federal Bureau of Investigation Internet Crime Complaint Center. https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf

[12] Gartner. (2021). Is the cloud secure? Gartner Research. https://www.gartner.com/smarterwithgartner/is-the-cloud-secure

[13] IBM Security. (2023). Cost of a data breach report 2023. https://www.ibm.com/reports/data-breach

[14] International Telecommunications Union. (2023). Global cybersecurity index 2022. ITU Publications. https://www.itu.int/en/ITU-D/Cybersecurity/Pages/GCI.aspx

[15] Kalu-Mba, Mupa & Tafirenyika (2025) — *The role of machine learning in post-disaster humanitarian operations: Case studies and strategic implications*

[16] NASCIO. (2023). State CIO survey 2023: The digital state. National Association of State Chief Information Officers. https://www.nascio.org/resource-center/resources/2023-state-cio-survey/

[17] Nemure T, et al (2026), Cybersecurity of critical infrastructures: Challenges and future perspectives, World Journal of Advanced Research and Reviews, https://doi.org/10.30574/wjarr.2026.29.3.0621

[18] NIST NCCoE. (2024). Overview of the NIST cybersecurity framework (CSF) 2.0 small business quick start guide. National Cybersecurity Center of Excellence. https://www.nist.gov/system/files/documents/2024/03/20/March20_2024_NISTCSF2.0_SMBQSG.Overview.pdf

[19] NIST. (2024). The NIST cybersecurity framework 2.0 (NIST CSWP 29). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.29

[20] Norris, D. F., & Mateczun, L. (2019). Managing cybersecurity at the grassroots: Evidence from the first nationwide survey of local government cybersecurity. Journal of Urban Affairs, 41(6), 759–775. https://doi.org/10.1080/07352166.2019.1629051

[21] Qualys. (2024). Reaching NIST CSF 2.0 adaptable tier with Qualys. Qualys Blog. https://blog.qualys.com/product-tech/2024/03/05/achieving-nist-csf-2-0-top-tier-adaptable-status

[22] Savola, R., Pietilainen, H., & Evesti, A. (2015). Developing a security metrics taxonomy for cybersecurity and information security management. Proceedings of the International Conference on Cyber Security, 112–121.

[23] Shiraishi & Mupa (2025) — *Valuation complexities in the energy transition: Traditional vs. green energy firms in M&A*

[24] Sophos. (2024). State of ransomware in government 2024. Sophos Threat Research. https://www.sophos.com/en-us/medialibrary/pdfs/whitepaper/sophos-state-of-ransomware-government-2024.pdf

[25] Souppaya, M., & Scarfone, K. (2013). Guide to enterprise patch management technologies (NIST SP 800-40 Rev. 3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-40r3

[26] Verizon. (2023). 2023 data breach investigations report. Verizon Business. https://www.verizon.com/business/resources/reports/dbir/

[27] Zhuwankinyu, E. K., Moyo, T. M., & Mupa, M. N. (2024). Leveraging generative AI for an ethical and adaptive cybersecurity framework in enterprise environments. IRE Journals, 8(6), 654-675.