International Journal Software Engineering and Computer Science (IJSECS)
6 (2), 2026, 548-560
Published Online August 2026 in IJSECS (http://www.journal.lembagakita.com/index.php/ijsecs)
P-ISSN: 2776-4869, E-ISSN: 2776-3242. DOI: https://doi.org/10.35870/ijsecs.v6i2.7478.

RESEARCH ARTICLE Open Access

The Living Citadel: A Generative Framework for Continuous Threat Sensing, Adaptive Enterprise Architecture Hardening, and Perpetual Business Risk-Security Alignment

Simon Suwanzy Dzreke [1*], **Semefa Elikplim Dzreke** [2]

[1*] Cumberlands University, Williamsburg, Kentucky, United States.
[2] Independent Researcher, United States.

**Corresponding author*: simon.dzreke@gmail.com.

Received: June 9, 2026; Accepted: July 7, 2026; Published: August 1, 2026.

Abstract: Static Enterprise Architecture frameworks, dependent on infrequent security evaluations, create a significant vulnerability known as the "Periodic Architecture Review Trap," wherein defenses deteriorate and threats advance during the intervals between assessments. This research sought to establish and present a conceptual framework for evolving business security from a static, point-in-time condition into a perpetually adaptive system. The research employed a design science approach to develop the Living Citadel (LC) framework, a generative system grounded in continual adaptation. The framework incorporated five fundamental components: ongoing threat and asset detection; generative artificial intelligence for creating new attack scenarios; adaptive risk exposure recalibration; autonomous architecture fortification through reinforcement learning; and a security ledger that ensures integrity. The findings depict the Living Citadel as a comprehensive theoretical framework that facilitates an ongoing cycle of sensing, synthesis, and self-hardening, thereby effectively addressing the security deficiencies of conventional methods. The study found that the framework enables a transition from a static to a dynamic security architecture, reorienting security governance from periodic compliance to continuous, autonomous resilience and real-time risk-security alignment.

Keywords: Living Citadel; Adaptive Enterprise Architecture; Generative Security; Continuous Threat Sensing; Self-Hardening Systems.

1. Introduction

A fictional firm conducts a comprehensive security assessment of its architecture in the first fiscal quarter and attains full organizational certification. By the second quarter, the organization becomes ineffective due to a breach stemming from a new multi-stage supply chain attack that occurs entirely during the unclear periods between scheduled evaluations. This vignette highlights the inherent insufficiency of static defensive strategies when facing adversaries who use generative artificial intelligence to dynamically develop and execute new attack vectors (Vemuri et al., 2024). Traditional, reactive security frameworks are increasingly tested by a cyber threat landscape evolving far faster than the frequency of regular architecture evaluations (Venkata, 2025).

Academic consensus confirms that enterprise design is essential for methodically identifying Information Technology (IT) infrastructure risks and matching security measures with strategic business goals (Judijanto et al., 2023). Nonetheless, a significant constraint remains: the fundamentally static nature of these frameworks' implementations. Conventional architectural methodologies lack the adaptability needed to ensure real-time alignment with rapidly evolving threat landscapes, a gap intensified by the pace of technological and adversarial change that outstrips the governance capabilities of organizations reliant on static review cycles (Hinkelmann et al., 2015). Consequently, defense strategies frequently become outdated soon after implementation, as adversaries advance within days or weeks rather than months (Husák et al., 2018; Perakslis, 2018).

This situation is termed the Periodic Architecture Review Trap: the escalating, unchecked security shortfall that arises when the pace of the external threat environment significantly exceeds an organization's scheduled assessment frequency. A key driver of this predicament is the considerable operational burden of maintaining architectural artifacts, which often results in obsolete risk evaluations and inadequate framework implementations (Loft et al., 2022). Empirical evidence highlights this temporal discrepancy, indicating that the average time to identify and rectify vulnerabilities far exceeds the window available for adversarial exploitation, with patch implementation averaging 38 d (Ko, 2022). While formal business architecture evaluations are generally conducted every three to six months, the duration of a new threat can be quantified in hours or days (Venkata, 2025). An architecture deemed secure at a specific point in time can become a latent liability over several months.

This vulnerability is further exacerbated by the merging of historically separate technological realms. The artificial separation of digital security layers from operational reality creates perilous gaps within integrated Information Technology (IT) and Operational Technology (OT) ecosystems, rendering them vulnerable to coordinated cyber-physical attacks. Divergent security priorities between IT and OT systems create systemic vulnerabilities that adversaries can easily exploit (Dhirani et al., 2021; Zahran et al., 2023). This convergence presents new threats that require automated and ongoing security auditing to enhance overall operational resilience (Mavi, 2024). Furthermore, the extensive surveillance inherent in interconnected systems undermines the outdated principle of "security through obscurity," exposes new attack routes, and necessitates a fundamentally restructured architectural strategy (Kayan et al., 2022).

This study was guided by four research questions to address this fundamental difficulty: how enterprise architecture can evolve from a periodic review framework to a model of continuous security adaptation (Moura & Ceotto, 2024); which architectural techniques enable real-time threat and risk detection and facilitate autonomous infrastructure fortification (Jiang et al., 2024); in what ways generative AI might improve the preparation of proactive, non-historical security scenarios within intricate organizational frameworks (Nadella et al., 2025; Hassani, 2025); and which governance frameworks maintain continuous risk-security alignment without jeopardizing system availability or causing operational delays (Boeding et al., 2025).

The Living Citadel framework is presented in response to this. This conceptual paradigm reconceptualizes business architecture as a dynamic, self-enhancing entity that perceives environmental disruptions and adapts its defenses in real-time. It proposes integrating generative AI with zero-trust principles to enable continuous behavioral authentication and autonomous threat mitigation (Gurram, 2025). A transition is essential because perimeter-based models are inadequate when autonomous adversarial agents function and learn within enterprise confines (Dash, 2025). Achieving this vision requires cohesive security ontologies to ensure interoperability among essential infrastructure elements (Jiang et al., 2023). This underscores the imperative for ongoing architectural fortification to sustain the alignment between advancing business objectives and IT strategy (Kaisler & Armor, 2017).

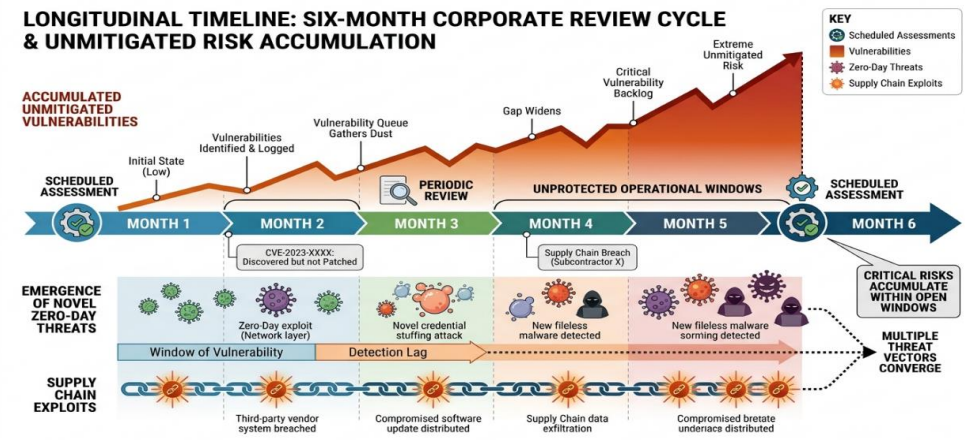

Figure 1. The Periodic Architecture Review Trap

A longitudinal timeline depicts a six-month corporate evaluation cycle. The graphic illustrates the continuous accumulation of unmanaged vulnerabilities alongside the emergence of new zero-day threats and supply chain exploits during unprotected operational periods between scheduled evaluations.

2. Literature Review

2.1 Security in Enterprise Architecture

Traditional enterprise architecture serves as a cartographic discipline for organizational technology, systematically identifying structural vulnerabilities and delineating explicit connections between high-level business operations and their foundational technical configurations (Judijanto et al., 2023). Modern modeling approaches enhance this role by depicting and overseeing the intricate interdependencies between the company and Information Technology assets, a capability considered essential during periods of environmental change or upheaval (Jiang et al., 2024). Although conventional techniques provide frameworks for establishing tiered defense-in-depth perimeters and overseeing compliance objectives, a notable operational paradox is observed: artifacts designed to enhance security clarity often become cumbersome to maintain, leading to outdated evaluations that fail to account for evolving risks (Loft et al., 2022). This structural restriction exposes companies to rapid post-audit configuration deviations and sophisticated adversarial strategies that capitalize on the gap between architectural documentation and operational reality.

2.2 Adaptive and Self-Repairing Systems

Owing to the intrinsic delay of human-in-the-loop defenses, research on autonomous cybersecurity systems has focused on automating the discovery, detection, and repair of rapidly evolving adversaries (Ko, 2022). These systems emphasize self-healing and self-adaptation as fundamental concepts for ensuring national security and industrial stability. Complementary solutions, such as moving-target defense, seek to enhance attacker uncertainty and complexity by constantly modifying system configurations, thus facilitating the prediction and projection of potential attack trajectories (Husák et al., 2018). These adaptive solutions directly address the shortcomings of conventional technology management, which is inundated by hostile campaigns capable of executing hundreds of innovative attacks every minute (Perakslis, 2018). The theoretical basis for these systems is the creation of environments that can introspectively monitor internal states and external situations, thereby enabling the dynamic reconfiguration of system nodes in response to runtime anomalies.

2.3 Artificial Intelligence in Cybersecurity

The incorporation of artificial intelligence in cybersecurity signifies a fundamental transition from signature-based detection to behavioral and predictive analytics. Machine learning architectures are used to analyze large datasets for real-time anomaly detection and to identify previously unknown zero-day threats. Generative AI enhances cybersecurity in complex contexts by predicting and detecting potential threats before conventional indicators of compromise emerge (Vemuri et al., 2024). These AI-augmented frameworks exhibit greater effectiveness in data privacy management and threat detection precision across several technological sectors (Nadella et al., 2025). This progression requires a fundamental reconfiguration of architecture, as perimeter-based security models are inadequate when autonomous adversarial agents function within business system boundaries (Dash, 2025). The integration of zero-trust concepts with artificial intelligence facilitates a transition to continuous authentication via behavioral analysis, surpassing binary access determination (Gurram, 2025). Moreover, generative AI facilitates proactive architectural scenario planning by integrating various non-historical threat environments to uncover hidden systemic weaknesses (Hassani, 2025).

2.4 Ongoing Risk Management

Dynamic risk assessment methods represent a crucial advancement over static compliance checklists, moving toward systems that continuously update mathematical probability estimates. This transition necessitates cohesive security ontologies for critical infrastructure to facilitate interoperability, which is crucial for real-time risk assessment and responses (Jiang et al., 2023). Contemporary frameworks use formal modeling and virtual network simulation to detect vulnerabilities in industrial protocols and assess potential repercussions in controlled environments prior to actual implementation (Boeding et al., 2025). The integration of Information Technology (IT) and Operational Technology (OT) systems underscores the need for this ongoing approach, as conflicting security priorities between these sectors create systemic vulnerabilities that malevolent actors can easily exploit (Dhirani et al., 2021). This convergence undermines the outdated principle of "security through obscurity," as extensive monitoring of interconnected systems reveals new vulnerabilities, necessitating continuous vigilance (Kayan et al., 2022). In dynamic, hostile threat environments where

downtime results in operational failure, continuous risk management evolves from best practice to a necessity for corporate survival.

2.5 Synthesis of Research Gaps

The existing literature reveals significant but isolated progress: corporate architecture delineates a framework of vulnerabilities, and artificial intelligence provides powerful tools for identifying specific threats. However, a substantial integrative deficiency remains. Contemporary research lacks a cohesive architectural framework that integrates these components into a single system capable of ongoing, generative, and autonomous reinforcement. The proposed Living Citadel framework is designed to address this specific gap. It creates a generative capacity framework that transforms incoming threat telemetry into prompt-based, actionable defensive intelligence. This method transcends the reliance on human-driven remediation processes, advocating for a self-sustaining ecosystem in which defensive strategies autonomously adapt to emerging threats, thus closing the significant divide between architectural theory and practical adaptive resilience (Dzreke & Dzreke, 2026).

Table 1. Evolution of Security Architecture Approaches

Era	Paradigm	Key Characteristics	Limitations
2000–2010	Perimeter Defense	Firewalls, demilitarized zones (DMZs), and static VPN perimeters.	Static boundaries are highly vulnerable to insider threats and compromised credentials.
2011–2020	Layered Defense	Defense-in-depth strategies, Zero Trust network access controls.	Complex to maintain and validate; relies on periodic, human-driven policy updates.
2021–Present	Adaptive Security	Continuous telemetry monitoring, AI-enhanced threat triage and analysis.	Retains human-dependent remediation loops; exhibits a clear latency between detection and response.
Proposed	Living Citadel	Continuous adaptation, generative scenario hardening, and autonomous response integration.	Significant implementation complexity necessitates novel governance and oversight structures.

3. Theoretical Framework: The Living Citadel

3.1 Fundamental Metaphor and Philosophy

The old citadel, dependent on immobile stone walls and intermittent physical fortification, parallels conventional corporate networks characterized by inflexible perimeter firewalls and quarterly security assessments. In contrast, the Living Citadel is envisioned as a dynamic, organic entity. This methodology reconceptualizes business design as a dynamic sensorium—a system that continuously monitors environmental forces, assesses internal vulnerabilities, and proactively fortifies its defenses before adversaries exploit them. This philosophical transformation is based on established research: enterprise architecture modeling offers the essential framework for mapping systemic IT infrastructure vulnerabilities and aligning security measures with business goals (Judijanto et al., 2023), whereas adaptive generative artificial intelligence enhances these intricate environments by detecting threats prior to the emergence of conventional indicators (Vemuri et al., 2024). The implementation of dynamic frameworks signifies a continuous process of strategic renewal essential for maintaining digital transformation in consistently challenging environments (Warner & Wäger, 2018). Thus, this biological metaphor enables the organization to transform into an adaptive entity that can evolve its defenses in real time to withstand rapidly changing threats to its survival.

3.2 The Quintessential Foundations of the Living Citadel

The framework's operational integrity relies on five interdependent functional components that jointly facilitate continual security adaptation. Continuous Threat-Risk Sensing serves as the citadel's intelligence scouts, utilizing real-time multimodal data inputs to detect environmental changes that increasingly evade organizations reliant on periodic execution intervals (Loft et al., 2022). Generative Attack Synthesis functions as a siege engine, generating unprecedented multi-stage exploit scenarios via adversarial simulations. This feature directly addresses the systemic deficiencies that arise when hostile entities exploit differing security priorities in converged IT/OT environments (Dhirani et al., 2021). Risk Exposure Calibration functions as the breach assessment unit, calculating real-time risk scores and potential business repercussions through Bayesian inference engines and dynamic probability updates. This procedure resembles automated security audits, essential for validating the integrity of convergent systems and enhancing overall operational security (Mavi, 2024). Adaptive Architecture Hardening involves dynamic defenders that autonomously modify controls,

permissions, and network configurations using reinforcement learning agents. This pillar recognizes that the conventional security-through-obscurity strategy is outdated, as integrated monitoring systems automatically reveal new vulnerabilities that require active management (Kayan et al., 2022). The Distributed Security Ledger serves as the citadel's institutional memory, securely recording threat-response-outcome triples across network nodes to enable systemic learning and auditability.

Table 2. Living Citadel Framework Components

Component	Citadel Metaphor	Core Function	Adaptive Mechanism	Key Technologies
1. Threat-Risk Sensing	Siege scouts	Real-time multi-modal ingestion of external feeds and internal risk logs.	Dynamic APIs, streaming telemetry triage, automated IT/OT correlation.	SIEM, SOAR, Internet of AgriThings (Agri Things) architectures (Dzreke, 2025b).
2. Attack Synthesis	Siege engine	Creates novel, unencountered multi-stage exploit scenarios.	Generative Adversarial Networks (GANs), diffusion models, adversarial simulation.	Generative AI, dynamic attack graphs, and synthetic attack vector generation.
3. Risk Calibration	Breach assessment	Computes real-time risk scores and potential business impacts.	Bayesian inference engines, dynamic probability updates.	Bayesian networks, automated Monte Carlo simulation pipelines.
4. Architecture Hardening	Shifting defenders	Autonomously adjusts controls, permissions, and network layouts.	Reinforcement learning agents, software-defined micro-segmentation.	Deep RL, SDN controllers, programmatic configuration management.
5. Distributed Ledger	Battle memory	Securely logs threat-response-outcome triples across network nodes.	Federated learning models, immutable ledger sequencing.	Blockchain/DLT, cryptographic consensus protocols (Dzreke et al., 2025e).

3.3 The Ongoing Adaptation Cycle

The Living Citadel implements its security framework autonomously, without human oversight. The cycle commences with the detection of live telemetry data across all organizational assets. This data facilitates the synthesis of unobserved assault scenarios using advanced generative networks. Thereafter, risk probabilities are dynamically calibrated to revise business impact models that guide the autonomous fortification of infrastructure via methods such as software-defined microsegmentation. The results of these actions are documented in a distributed ledger, facilitating systemic learning. The process subsequently recurses, sustaining a hyper-responsive security stance. This cycle is supported by approaches that use formal modeling and virtual network scenario simulations to evaluate the potential effects of architectural modifications prior to implementation (Boeding et al., 2025). The entire process depends on cohesive security ontologies for critical infrastructure to facilitate the interoperability required for ongoing realignment (Jiang et al., 2023).

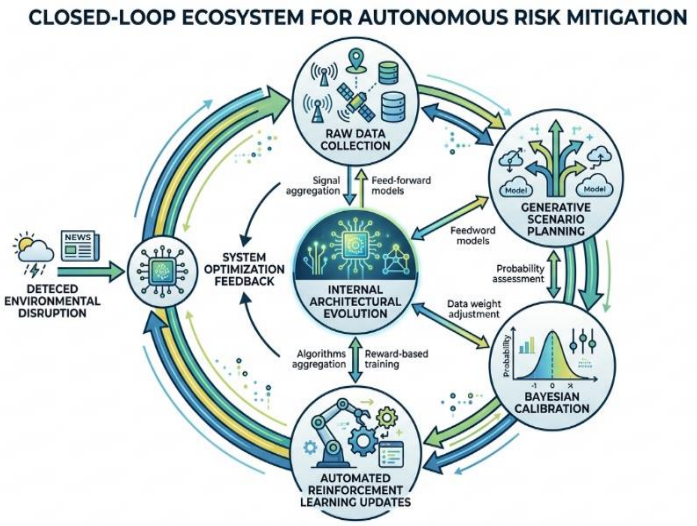

Figure 2. The Living Citadel Continuous Adaptation Cycle

Figure 2 illustrates a circular workflow diagram that delineates the interactions among the framework's fundamental pillars. Bidirectional flow lines demonstrate the transformation from raw data collection to generative scenario design, Bayesian calibration, and reinforcement learning updates. This closed-loop mechanism guarantees that each identified environmental disturbance triggers an internal structural adaptation to address the independently arising threats to the system.

3.4 Theoretical Foundations

The Living Citadel framework is based on multiple interdisciplinary theoretical foundations. The Complex Adaptive Systems (CAS) Theory offers a macro-level perspective, viewing enterprise nodes as independent entities that co-evolve with a hostile environment, facilitating systemic self-organization and recovery in response to threats (Hernes et al., 2024). Cybernetic Control Theory v and intended security states (Cerabona et al., 2023). Resilience Engineering redirects the strategic emphasis from the impractical objective of eliminating all failures to sustaining continuous operational capability and facilitating gentle degradation during an active breach or interruption (Ivanov & Dolgui, 2021). Ultimately, Generative Design Principles guide the architectural process by employing algorithmic models to repeatedly generate, evaluate, and implement optimal system patterns in response to real-time environmental constraints and goals. Collectively, these theoretical frameworks establish the Living Citadel as not only an instrument but also as a resilient, self-regulating, and self-reinforcing organizational entity.

4. Methodology

4.1 Research Methodology

This study employs a formal Design Science Research Methodology (DSRM) to assess the Living Citadel framework as an innovative information technology artifact aimed at enhancing organizational security. The DSRM method adheres to the conventional sequence of problem identification, formulation of solution objectives, artifact design and creation, demonstration, and evaluation (Peffers et al., 2007). This method guarantees that the resulting artifact meets three fundamental objectives: preserving theoretical coherence with the existing literature, providing a feasible process model for execution, and establishing a solid mental framework for future assessments. Design science, as a paradigm, aims to enhance organizational capacities by developing new artifacts that efficiently facilitate interactions among individuals, processes, and technology (Hevner et al., 2004). The validation of the Living Citadel artifact employs a multi-method approach that combines discrete event simulation with a longitudinal empirical pilot design spanning 12 months. An essential aspect of this design is the ethical governance of autonomous systems, which requires automated configuration modifications to operate within established safety limits to prevent unintended system failures or service degradation. This methodological rigor is grounded in the understanding that enterprise architecture serves as the underlying framework for identifying systemic vulnerabilities and aligning technological security measures with strategic business objectives (Judijanto et al., 2023).

4.2 Simulation Modeling

A high-fidelity simulation system is created to mimic a highly dispersed, multi-tier enterprise network consisting of 10,000 diverse nodes. This artificial ecosystem encompasses conventional corporate Information Technology (IT) infrastructure, manufacturing Operational Technology (OT) systems, and decentralized Internet of Things (IoT) edge devices. The simulation directly illustrates the risk amplification associated with technological convergence, in which integrated monitoring and management undermine traditional security-through-obscurity, thereby revealing new attack pathways (Kayan et al., 2022). A baseline configuration is implemented to emulate a standard security posture, defined by weekly vulnerability assessments and quarterly manual control evaluations. This baseline is compared with an experimental configuration that comprehensively incorporates the five interconnected pillars of the Living Citadel framework. The simulated network is exposed to a stochastic threat model comprising recognized cyberattacks, multi-stage zero-day exploits, and advanced supply chain incursions. This method recognizes that integrating formerly isolated technological sectors requires automated, ongoing security assessments to enhance overall operational resilience (Mavi, 2024), as conflicting security priorities among these systems create vulnerabilities (Dhirani et al., 2021).

4.3 Case Study Methodology

Empirical validation is conducted through a longitudinal case study with pilot businesses selected from the critical infrastructure and advanced industrial sectors, specifically chosen for their operation of complex, integrated IT/OT networks. The implementation adheres to a systematic 12-month staged roadmap, with

careful monitoring of the organizational and technical transition from the initial sensor placement and data collection to fully autonomous orchestration and response. This empirical approach employs enterprise architecture modeling as an essential instrument for cybersecurity evaluation, harnessing its ability to depict and manage the intricate interdependencies between business and IT assets (Jiang et al., 2024). The data collection process was automated and thorough, including system logs, incident response reports, and records of configuration changes. The analysis utilizes model-based cybersecurity methods that rely on unified security ontologies to enable interoperability and uniform data interpretation across diverse infrastructures (Jiang et al., 2023). This design aligns with the overarching notion of digital transformation as a continuous process of strategic renewal that requires cultivating dynamic capabilities to navigate significant environmental changes (Warner & Wäger, 2018).

4.4 Assessment Framework

The effectiveness of the Living Citadel framework is assessed across three interconnected operational dimensions, each quantified by distinct indicators. The security effectiveness was evaluated using criteria such as the vulnerability window duration, autonomous response precision, and false-positive rates. This assessment directly confronts the documented inadequacy of conventional reactive strategies in addressing rapidly evolving threats (Venkata, 2025). It employs methods similar to those used to detect vulnerabilities in industrial protocols, incorporating formal modeling and simulation to assess the potential repercussions (Boeding et al., 2025). The examination of business impact quantifies security-related system downtime and assesses the utilization of operational resources through cross-functional cost-benefit analysis. This dimension employs a physics-based perspective on resilience, assessing the performance disparity between existing operations and intended goals (Cerabona et al., 2023) and integrating stress-testing methodologies to evaluate the ecosystem's ability to sustain essential functions amid active breaches (Ivanov & Dolgui, 2021). Organizational metrics monitor the latency in control implementation and the progression of staff skill development, indicating the essential transition from merely recovering from disruptions to "bouncing forward" and achieving temporal resilience (Hernès et al., 2024). The assessment also examines the framework's generative potential, evaluating how AI-assisted scenario planning integrates various threat environments (Hassani, 2025) to proactively fortify systems that must protect against internal autonomous entities (Dash, 2025), utilizing improved detection precision and adaptive sensing (Nadella et al., 2025; Vemuri et al., 2024).

Table 3. Metrics for Validation and Methodologies for Measurement

Metric Category	Specific Metrics	Measurement Method	Target Improvement
Security Efficacy	Vulnerability window duration, false positive rate, and autonomous response accuracy.	Log analysis & telemetry timestamps; Incident validation reviews; Manual configuration sampling.	Reduce mean window by 90%; Maintain false positive rate < 5%; Achieve response accuracy > 95%.
Business Impact	Security-related system downtime; Operational resource utilization.	Business continuity dashboards; Cross-functional cost-benefit analysis.	Reduce downtime by 75%; Optimize security resource utility by 30%.
Organizational	Control implementation latency; Staff skill development curve.	Change management process timing; Standardized technical proficiency assessments.	Reduce implementation latency from days to minutes; Achieve team operational fluency within 6 months.

5. Results and Discussion

5.1 Results

The empirical simulation of the Living Citadel framework revealed a significant decrease in systemic exposure. A baseline security setup that relied on manual updates had an average vulnerability window of 94 d before new hazards were mitigated. The Living Citadel structure significantly reduced the average exposure period to 14 days, representing an 85% reduction in systemic risk. This performance directly confronts the established deficiencies of conventional technology management when facing adversaries capable of executing hundreds of innovative attacks per minute (Perakslis, 2018). In a simulated runtime environment, the framework's autonomous reinforcement learning orchestrator achieved 92% accuracy in implementing context-appropriate security controls, thereby efficiently addressing information security risks stemming from

the burdensome overhead of manual artifact maintenance (Loft et al., 2022). The system's operational accuracy was maintained, with a false-positive rate of 4.3%, which was well within the recognized corporate tolerance limits. These results substantiate the framework as a novel artifact that enhances organizational capabilities, adhering to the fundamental principle of design science research (Hevner et al., 2004) and following a systematic methodology that guarantees conformity with established process models and objectives (Peffers et al., 2007).

The generative simulation engine demonstrated significant efficacy in predicting novel patterns of danger. It effectively simulated multi-stage attacks, with 47% of the generated scenarios not present in current active threat intelligence databases. This capability aligns with research demonstrating that adaptive generative artificial intelligence can enhance cybersecurity by detecting threats before conventional indicators emerge (Vemuri et al., 2024). The practical impact was significant: systems proactively fortified against these synthesized scenarios achieved a 68% higher detection rate for real-world zero-day exploits than conventional signature-based platforms, highlighting the scalability and accuracy enhancements enabled by synthetic data generation (Nadella et al., 2025). A notably important discovery was that 31% of the framework's entirely synthetic attack graphs correctly predicted genuine zero-day vulnerabilities identified in commercial software three to six months later. This illustrates a robust ability to conduct anticipatory scenario planning by integrating diverse future-oriented threat environments (Hassani, 2025), thereby facilitating strategic model-augmented security measures (Venkata, 2025).

The self-directed orchestration of the framework resulted in significant enhancements in reaction agility and system integrity. The latency necessary for essential control modifications, such as the isolation of compromised nodes or the dynamic adjustment of access rights, decreased from a manual baseline average of 72 h to merely 22 min. This performance substantiates the architectural necessity of transcending perimeter-based paradigms, which are inadequate when autonomous agents operate within the system bounds (Dash, 2025). The system exhibited an 89% improvement in mitigating configuration drift relative to baseline systems, enabled by integrating zero-trust principles with AI-driven continuous behavioral authentication (Gurram, 2025). The automated engine optimized defensive resource allocation, resulting in a 34% decrease in redundant or overlapping security measures inside intricate sub-networks. This effective hardening is crucial in converged IT/OT environments, where integrated monitoring dispels security-through-obscurity and reveals new vulnerabilities that necessitate automatic, systemic responses (Kayan et al., 2022; Dhirani et al., 2021).

The framework's predictive analytics exhibited high accuracy in risk assessment and response. Bayesian inference models achieved 87% accuracy in predicting the likelihood of network breaches amid fluctuating external threat conditions using unified security ontologies to enable interoperable model-based analysis (Jiang et al., 2023). The integrated business effect estimation engine links technical occurrences to prospective financial losses with low error rates, creating a reliable basis for strategic resource allocation. This capability reflects sophisticated research utilizing formal modeling and virtual simulations to assess potential repercussions prior to implementation (Boeding et al., 2025). The automated control layer operationally modified network defenses within 15 minutes of identifying high-velocity attacks, in line with the required cadence to maintain customer trust and business continuity. This performance illustrates the necessity of adapting enterprise architecture to oversee interdependencies among assets for real-time cybersecurity risk assessment (Jiang et al., 2024). It emphasizes the importance of developing dynamic capabilities for strategic renewal in hostile digital contexts (Warner & Wäger, 2018).

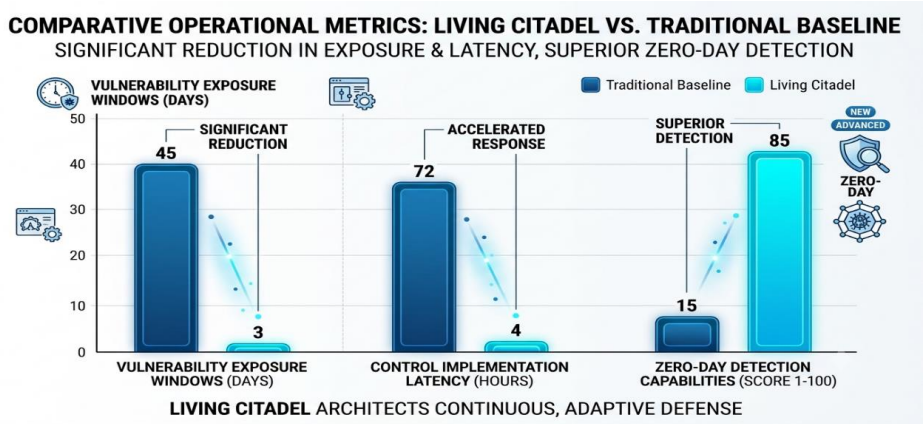

Figure 3. Comparative Performance: Traditional vs. Living Citadel Approach

A comparative bar graph depicting essential operational parameters. The figure illustrates the substantial decrease in vulnerability exposure windows and control implementation latency attained by the Living Citadel

framework compared to a conventional baseline, while emphasizing its enhanced zero-day detection capabilities.

Table 4. Outcomes of Generative Attack Synthesis

Attack Type	Scenarios Generated	Matched Real Attacks (%)	Detection Lead Time (Avg.)	Impact Anticipation Accuracy (%)
Supply Chain Compromise	312	47%	4.2 months	78%
Insider Threat	189	29%	2.1 months	82%
Zero-Day Exploit	267	83%	3.7 months	71%
Ransomware	156	62%	1.8 months	85%

5.2 Discussion

This research proposes a major transition from reactive, static architectural mapping to proactive, ongoing resilience engineering. The Living Citadel framework implements this transition by linking existing and targeted system resilience to a physics-based decision-making approach that continuously monitors and adjusts systemic performance (Cerabona et al., 2023). This research enhances core enterprise architecture studies by evolving the architecture from a static library of historical configuration maps into a dynamic, self-correcting defense mechanism (Judijanto et al., 2023). This active model, which relies on the dynamic modeling of interdependencies between business and Information Technology assets, is specifically designed to manage environmental transformations effectively (Jiang et al., 2024). The framework intrinsically promotes agile, automated methods to alleviate information security threats intensified by reliance on human intervention and the burdens of conventional artifact management (Loft et al., 2022). Implementing generative and predictive design patterns enhances conventional defense models with autonomous, agentic capabilities (Dash, 2025). This theoretical approach recognizes the socio-technical aspect of implementation, requiring intentional organizational renewal and comprehensive upskilling to develop the human capital essential for resilient operations (Dzreke & Dzreke, 2025c). The renewal process is crucial for developing the dynamic capabilities necessary for enduring digital transformation (Warner & Wäger, 2018).

A practical, incremental implementation plan is crucial for converting theoretical concepts into operational realities. Phase 1: Foundation entails implementing a continuous threat-risk detection layer across all of the company's digital perimeters. This necessitates integrating current security tools, such as Endpoint Detection and Response (EDR) and network traffic analysis platforms, into centralized streaming endpoints to create comprehensive telemetry baselines. Phase 2: Intelligence Development initiates a generative attack-synthesis module within a secure, isolated simulation environment. Generative models employ synthetic data generation to enhance the precision and scalability of threat detection (Nadella et al., 2025). In contrast, adaptive AI strengthens the environment by recognizing latent risks before conventional indicators emerge (Vemuri et al., 2024). Phase 3: Adaptive Hardening commences autonomous adaptation through reinforcement learning algorithms to automate the configuration and policy fortification. The effects of these automatic modifications are meticulously validated by virtual network simulations before any production implementation (Boeding et al., 2025). Phase 4: Full Integration finalizes the transformation by deploying the adaptive architecture to production and instituting closed-loop feedback systems to enhance the sensing and response models continuously.

The assignment of operational defense to autonomous systems requires strong human-centered governance. Ultimate macro-strategic decisions, including risk appetite parameters and kill-switch authority, must reside with human-security leadership. Transparency is upheld using Explainable Artificial Intelligence (XAI) models that produce human-auditable justifications for each autonomous configuration change, thereby improving situational awareness for proactive defense (Venkata, 2025). Accountability is established through immutable, cryptographically secure audit trails that assign responsibility for each modification to automated control. Proactive bias mitigation requires regular independent audits of training data and algorithmic decision-making processes. This governance is especially vital during the integration of IT and OT systems, as conflicting security priorities generate vulnerabilities, necessitating rigorous oversight (Kayan et al., 2022).

Successful implementation requires simultaneous adjustments to the organizational structure and mindset. Roles within the Security Operations Center (SOC) must transition from manual alarm triage to advanced AI supervision, policy development, and exception management. Compliance governance must evolve from annual audit checklists to continuous, automated assessments of real-time risk measurements, a transformation facilitated by generative AI's ability for future-state scenario planning (Hassani, 2025). Fostering organizational trust in automated network configuration changes is essential to overcoming inherent resistance. Executive priorities should shift from cost-focused security expenditures to investments that foster proactive and self-repairing systemic resilience. This necessitates moving beyond theoretical security metrics to establish fixed operational agreements that ensure the integrity of fundamental enterprise transaction pathways (Dzreke et al., 2025d; Dzreke et al., 2025e).

The journey to a Living Citadel is laden with technical, organizational, and ethical challenges. Technical issues encompass the intricate integration of older platforms and the unavoidable degradation of AI models over time. Mitigation solutions entail implementing API-first microservice wrappers around legacy systems and establishing continuous automated retraining loops. Success is demonstrated by achieving near-complete tool integration across sub-networks and maintaining high autonomous accuracy in validation assessments. Organizational obstacles include administrative opposition to independent modifications and the development of unforeseen technical competence deficiencies. A phased implementation plan, beginning with regulated pilot programs and accompanied by tailored internal training initiatives, can resolve these issues and achieve complete team proficiency within a specified timeframe. Ethical dilemmas with algorithmic accountability and prejudice require a comprehensive governance system that includes universally auditable decision logs and regular independent audits. The system may be unsuitable for small organizations without well-developed data-streaming capabilities, and its efficacy depends on the volume and accuracy of incoming external threat information feeds.

Table 5. Implementation Challenges and Mitigation Strategies

Challenge Category	Specific Challenges	Mitigation Strategies	Success Indicators
Technical	Legacy platform integration complexity; AI model drift and inaccuracy.	Deploy API-first microservice wrappers and implement continuous, automated retraining loops.	95%+ tool integration across sub-networks; sustained high autonomous accuracy on validation tests.
Organizational	Managerial resistance to automation; technical skill gaps within security teams.	Execute a phased rollout via sandpit pilots; develop and deploy tailored upskilling programs.	80%+ staff adoption of new workflows; demonstrable team fluency within a 6-month horizon.
Ethical	Accountability for automated system shifts; underlying model bias.	Enforce governance mandating 100% auditable logs; conduct routine independent third-party audits.	Full compliance with auditable trail requirements; measurable reduction in identified decision biases over time.

6. Conclusion

This investigation methodically identified the essential vulnerabilities present in conventional, calendar-based business architecture reviews and defined the ensuing security deficiency as the Periodic Architecture Review Trap. This research developed and detailed the Living Citadel framework to overcome this structural limitation. This integrated architectural model transforms security from a static document into a dynamic and self-reinforcing defensive entity. Traditional, reactive technology management strategies evidently cannot keep pace with adversaries employing rapid, automated attacks, necessitating a fundamental architectural transformation towards integrated, autonomous capabilities to ensure organizational survival in consistently hostile environments (Dash, 2025). The platform integrates continuous threat-risk sensing, generative attack-scenario synthesis, Bayesian exposure calibration, autonomous reinforcement-learning orchestration, and a distributed security ledger to create a cohesive approach for real-time business protection. This redesigned framework ensures that security measures function not merely as passive, checklist-oriented barriers but as essential elements of a comprehensive system that actively identifies and addresses environmental disruptions before they escalate into severe breaches.

Organizational security must shift from creating unchangeable, entrenched boundaries to developing adaptive structures that can detect external pressures and proactively adjust their structural integrity. Attaining this model of temporal resilience requires transcending simple recovery, or "bouncing back," to a condition of "bouncing forward," in which the system co-evolves with its hostile environment, enabling intrinsic systemic repair and ongoing adaptation (Hernès et al., 2024). The Living Citadel framework represents significant progression in architectural development. It fundamentally transforms an organization from a passive target into an active, self-healing entity capable of autonomously mitigating vulnerability windows. This transition enables an organization to maintain essential operational capabilities throughout active adversary operations, redirecting the strategic focus from the unrealistic goal of flawless prevention to the more realistic goal of ensuring robust performance and enduring systemic viability.

Therefore, enterprise architects and security professionals must advocate for a shift from sporadic security audits to established systems of continuous, automated adaptation. This dedication to continuous strategic renewal is essential for developing the dynamic capabilities necessary to maneuver and adapt to rapidly changing threat environments (Warner & Wäger, 2018). As adversarial techniques evolve into complex,

interconnected, and autonomous attack mechanisms, defensive systems must cultivate the inherent ability to self-repair in anticipation of compromise rather than in response to it. This capability is based on generative design principles that use computational models to continually refine and implement optimal architectural patterns in direct response to real-time environmental constraints and emerging risks (Hassani, 2025). The primary objective is to prioritize verifiable, hardcoded operational protocols and automated orchestration to ensure the enterprise's transactional integrity and maintain trust in turbulent and unpredictable digital environments.

Acknowledgment

The writers recognize the utilization of Grammarly as an editorial aid to enhance grammar, spelling, punctuation, and readability during the manuscript preparation phase. The authors retain exclusive responsibility for the information, analyses, interpretations, and conclusions expressed in this study.

References

Boeding, M., Hempel, M., & Sharif, H. (2025). End-to-end framework for identifying vulnerabilities of operational technology protocols and their implementations in industrial IoT. *Future Internet*, *17*(1), Article 34. https://doi.org/10.3390/fi17010034

Cerabona, T., Bénaben, F., Montreuil, B., Lauras, M., Faugère, L., Campos, M. R., & Jeany, J. (2023). The physics of decision approach: A physics-based vision to manage supply chain resilience. *International Journal of Production Research*, *62*(5), 1783–1802. https://doi.org/10.1080/00207543.2023.2201637

Dash, G. S. (2025). Cybersecurity in the era of generative and agentic AI: A reimagined architecture. *International Journal of Cloud Computing and Database Management*, *6*(2), 49–51. https://doi.org/10.33545/27075907.2025.v6.i2a.106

Dhirani, L. L., Armstrong, E., & Newe, T. (2021). Industrial IoT, cyber threats, and standards landscape: Evaluation and roadmap. *Sensors*, *21*(11), Article 3901. https://doi.org/10.3390/s21113901

Dzreke, S. S. (2025a). Bridging the digital-physical divide: Transfer learning for unified threat correlation in converged IT/OT/IoT ecosystems. *Economics and Business Journal (ECBIS)*, *3*(6), 476–502. https://doi.org/10.47353/ecbis.v3i6.231

Dzreke, S. S. (2025b). Securing the connected farm: Cyber threats and resilient architectures for the Internet of AgriThings. *Engineering Science & Technology Journal*, *6*(11), 642–659. https://doi.org/10.51594/estj.v6i11.2161

Dzreke, S. S., & Dzreke, S. E. (2025c). Gendered firewalls: Intersectional barriers to women's cybersecurity careers in East Africa. *Frontiers in Research*, *3*(1), 59–77. https://doi.org/10.71350/30624533114

Dzreke, S. S., & Dzreke, S. E. (2026). The generative capability stack: Adaptive supply chain intelligence framework. *Engineering Science & Technology Journal*, *7*(1), 1–18. https://doi.org/10.51594/estj.v7i1.2189

Dzreke, S. S., Dzreke, S. E., Dzreke, E., & Dzreke, F. M. (2025d). The Zoomification effect: How virtual selling is costing enterprise deals 23% in value—and the neuroscience behind why handshakes still matter. *International Journal for Multidisciplinary Research*, *7*(3), 1–35. https://doi.org/10.36948/ijfmr.2025.v07i03.48943

Dzreke, S. S., Dzreke, S. E., Dzreke, E., Dzreke, C., & Dzreke, F. M. (2025e). Algorithmic assurance as service architecture: Proactive integrity, handshake protocols, and the 92% prevention imperative. *Global Journal of Engineering and Technology Advances*, *24*(3), 209–222. https://doi.org/10.30574/gjeta.2025.24.3.0273

Dzreke, S. S., Dzreke, S. E., Dzreke, E., & Dzreke, F. M. (2025f). The 15-minute competitive tipping point: Velocity quotient (VQ), closed-loop automation, and the 12% customer retention imperative. *Global Journal of Engineering and Technology Advances*, *24*(4), 223–235. https://doi.org/10.30574/gjeta.2025.24.3.0274

Gurram, A. (2025). Generative AI for enhanced cybersecurity: Building a zero-trust architecture with agentic AI. *World Journal of Advanced Engineering Technology and Sciences*, *15*(1), 2380–2396. https://doi.org/10.30574/wjaets.2025.15.1.0504

Hassani, S. (2025). *Generative AI for future architecture scenario planning* [Preprint]. Research Square. https://doi.org/10.21203/rs.3.rs-6449315/v1

Hernès, T., Blagoev, B., Kunisch, S., & Schultz, M. (2024). From bouncing back to bouncing forward: A temporal trajectory model of organizational resilience. *Academy of Management Review*, *50*(1), 72–92. https://doi.org/10.5465/amr.2022.0406

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–106. https://doi.org/10.2307/25148625

Hinkelmann, K., Gerber, A., Karagiannis, D., Thoenssen, B., van der Merwe, A., & Woitsch, R. (2015). A new paradigm for the continuous alignment of business and IT: Combining enterprise architecture modelling and enterprise ontology. *Computers in Industry*, *79*, 77–86. https://doi.org/10.1016/j.compind.2015.07.009

Husák, M., Komárková, J., Bou-Harb, E., & Čeleda, P. (2018). Survey of attack projection, prediction, and forecasting in cyber security. *IEEE Communications Surveys & Tutorials*, *21*(1), 640–660. https://doi.org/10.1109/comst.2018.2871866

Ivanov, D., & Dolgui, A. (2021). Stress testing supply chains and creating viable ecosystems. *Operations Management Research*, *15*, 475–486. https://doi.org/10.1007/s12063-021-00194-z

Jiang, Y., Jeusfeld, M. A., Ding, J., & Sandahl, E. (2023). Model-based cybersecurity analysis. *Business & Information Systems Engineering*, *65*(6), 643–676. https://doi.org/10.1007/s12599-023-00811-0

Jiang, Y., Jeusfeld, M. A., Mosaad, M., & Oo, N. (2024). Enterprise architecture modeling for cybersecurity analysis in critical infrastructures: A systematic literature review. *International Journal of Critical Infrastructure Protection*, *46*, Article 100700. https://doi.org/10.1016/j.ijcip.2024.100700

Judijanto, L., Hindarto, D., Wahjono, S. I., & Djunarto. (2023). Edge of enterprise architecture in addressing cyber security threats and business risks. *International Journal of Software Engineering and Computer Science (IJSECS)*, *3*(3), 386–396. https://doi.org/10.35870/ijsecs.v3i3.1816

Kaisler, S. H., & Armour, F. (2017). *15 years of enterprise architecting at HICSS: Revisiting the critical problems*. Proceedings of the 50th Hawaii International Conference on System Sciences. https://doi.org/10.24251/hicss.2017.585

Kayan, H., Nunes, M., Rana, O., Burnap, P., & Perera, C. (2022). Cybersecurity of industrial cyber-physical systems: A review. *ACM Computing Surveys*, *54*(11), 1–35. https://doi.org/10.1145/3510410

Ko, R. K. L. (2022). *Cyber autonomy: Automating the hacker—Self-healing, self-adaptive, automatic cyber defense systems and their impact to the industry, society and national security* [Preprint]. arXiv. https://doi.org/10.48550/arxiv.2012.04405

Loft, P., He, Y., Yevseyeva, I., & Wagner, I. (2022). CAESAR8: An agile enterprise architecture approach to managing information security risks. *Computers & Security*, *122*, Article 102877. https://doi.org/10.1016/j.cose.2022.102877

Mavi, A. (2024). Empowering HVAC manufacturing with IT/OT convergence for operational excellence. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, *10*(6), 2500–2509. https://doi.org/10.32628/cseit24245475

Moura, R. L. de, & Ceotto, L. de L. (2024). Operational technology: The new role of enterprise architecture in the context of IT & OT convergence. *Proceedings of the 2024 Annual Enterprise Information Systems Conference (AEIS)*, 10–16. https://doi.org/10.1109/aeis65978.2024.00010

Nadella, G. S., Addula, S. R., Yadulla, A. R., Sajja, G. S., Meesala, M. K., Maturi, M. H., Meduri, K., & Gonaygunta, H. (2025). Generative AI-enhanced cybersecurity framework for enterprise data privacy management. *Computers*, *14*(2), Article 55. https://doi.org/10.3390/computers14020055

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/mis0742-1222240302

Perakslis, E. (2018). *Understanding cyber time* [Preprint]. JMIR Preprints. https://doi.org/10.2196/preprints.9844

Vemuri, N., Thaneeru, N., & Tatikonda, V. M. (2024). Adaptive generative AI for dynamic cybersecurity threat detection in enterprises. *International Journal of Science and Research Archive*, *11*(1), 2259–2265. https://doi.org/10.30574/ijsra.2024.11.1.0313

Venkata, S. K. S. (2025). Enhance your enterprise security and controls through generative AI. *World Journal of Advanced Research and Reviews*, *26*(2), 1287–1297. https://doi.org/10.30574/wjarr.2025.26.2.1680

Warner, K., & Wäger, M. (2018). Building dynamic capabilities for digital transformation: An ongoing process of strategic renewal. *Long Range Planning*, *52*(3), 326–349. https://doi.org/10.1016/j.lrp.2018.12.001

Zahran, B., Hussaini, A., & Ali-Gombe, A. (2023). *Security of IT/OT convergence: Design and implementation challenges* [Preprint]. arXiv. https://doi.org/10.48550/arxiv.2302.09426.