

Automated System Hardening and Continuous Compliance Framework via CIS Benchmarks: A Lightweight Electron-Bash Approach

Seshu Kumar¹, N. Puneeth.², Ansh Kumar³, K. Pranav Raj⁴, K. Bhanu⁵, K. Pranav⁶

¹Assistant Professor, Department of Computer Science & Engineering
Keshav Memorial Institute of Technology, Narayanguda, Hyderabad, Telangana, India

^{2,3,4,5,6}Students, Department of Computer Science & Engineering
Keshav Memorial Institute of Technology, Narayanguda, Hyderabad, Telangana, India

Abstract - The rapid expansion of the cyber-attack surface in modern IT infrastructures necessitates a shift from reactive perimeter defenses to proactive host-level security. Operating system (OS) hardening using the Center for Internet Security (CIS) Benchmarks is an industry-recognized standard; however, manual implementation is highly inefficient and prone to human error and configuration drift. This paper introduces an automated framework, the "CIS Hardening Tool," which utilizes an Electron-based graphical interface to bridge high-level administrative tasks with low-level Bourne Again Shell (Bash) execution. We propose a Scan-Analyze-Remediate-Report (SARR) methodology to ensure verifiable system compliance. Our framework is designed to operate in heterogeneous environments with minimal infrastructure overhead, addressing the limitations of agent-based enterprise solutions. Through rigorous benchmarking on multiple Linux distributions, including Ubuntu 22.04 LTS and legacy Debian systems, we demonstrate that the proposed framework reduces the administrative burden by over 85% while maintaining 98% adherence to CIS controls. We further analyze the socio-technical barriers to adoption, the economic implications of automated auditing, and the role of hardening in Zero Trust architectures. The research concludes that lightweight, modular automation is essential for securing decentralized and heterogeneous computing environments in the modern enterprise landscape.

Index Terms - Cybersecurity, CIS Benchmarks, OS Hardening, Automation, Compliance Auditing, Electron Framework, Linux Security, DevSecOps, Configuration Management, Vulnerability Remediation

I. INTRODUCTION

The contemporary cybersecurity landscape is defined by the erosion of the traditional network perimeter. With the rise of hybrid cloud models, edge computing, and remote work, the individual operating system (OS) has become the primary battleground for security. Default OS configurations are predominantly optimized for user convenience and compatibility, often leaving behind a significant number of unnecessary services, legacy protocols, and permissive access controls. These settings provide low-hanging fruit for adversaries seeking to establish persistence or move laterally within a network. In an era where a single misconfiguration can lead to a multi-million dollar data breach, the importance of "secure-by-default" environments cannot be overstated.

Historically, the Center for Internet Security (CIS), a non-profit organization founded in 2000, has been at the forefront of defining security best practices. The "CIS Benchmarks" are consensus-based configuration guidelines for over 25 vendors. These benchmarks are categorized into levels: Level 1 for base security and Level 2 for high-security environments. However, the sheer volume of these documents—often exceeding 500 pages for a single OS distribution—makes manual implementation a Herculean task. A typical Linux benchmark contains over 250 individual controls, ranging from simple file permission checks to complex kernel parameter tuning via `sysctl`.

For an organization managing hundreds or thousands of nodes, manual implementation leads to "compliance debt," where systems remain unhardened for months due to the lack of specialized personnel. Existing work in this domain typically falls into high-level configuration management (e.g., Ansible, Puppet) or complex compliance frameworks (e.g., OpenSCAP). While powerful, these solutions often suffer from a steep learning curve and high infrastructure overhead. Ansible requires the maintenance of complex YAML playbooks, while SCAP-based tools rely on arcane XML definitions that are difficult for security auditors to verify.

There is a clear research gap for a lightweight, standalone, and visually accessible tool that enables "Scan-to-Remediate" workflows without requiring deep expertise in domain-specific languages.

This paper proposes an automated CIS Hardening Tool leveraging an Electron-Bash bridge. Our methodology focuses on the SARR lifecycle (Scan, Analyze, Remediate, Report) to provide a verifiable and reversible hardening process. We explore the architectural security of the tool itself, ensuring that the automation framework does not become a new attack vector. The contribution of this research is a multi-layered automation framework that democratizes system security, making professional-grade hardening accessible to administrators across all skill levels, thereby enhancing the collective resilience of modern digital infrastructure.

II. PROBLEM STATEMENT AND RESEARCH GAP

A. The Problem: Manual Compliance Inefficiency

The manual application of CIS Benchmarks is inherently unsustainable in high-scale environments. System administrators often face a binary choice: spend hundreds of hours on manual configuration or accept a degraded security posture. This leads to inconsistent security levels across the fleet, where some servers are well-fortified while others remain in their vulnerable default states. Furthermore, manual hardening lacks a standardized reporting mechanism, making it difficult to satisfy regulatory requirements (e.g., SOC2, HIPAA, PCI-DSS) during an audit. The human element also introduces significant risk; a single typo in a security configuration file (such as `/etc/ssh/sshd_config`) can either leave a system exposed or, conversely, cause an accidental denial of service (DoS) for critical applications by locking out all administrative users.

B. Research Gaps and Technical Limitations

Current methodologies for system hardening exhibit several critical limitations that this research seeks to address:

- 1) High Entry Barrier: SCAP-based tools (OpenSCAP) require deep knowledge of XML/OVAL files, which is a significant barrier for general system administrators [1]. The complexity of these standards often leads to their underutilization in mid-sized enterprises.
- 2) Infrastructure Dependency: Most enterprise hardening solutions require a central management server and a persistent agent on every node, which is unsuitable for isolated, air-gapped, or resource-constrained edge environments.
- 3) Remediation Disconnect: Many security scanners (e.g., Lynis) are excellent at finding vulnerabilities but fail to provide automated fixes, creating a "fix delay" where identified risks remain unpatched for weeks or months [2].
- 4) Configuration Drift: Manual fixes are often reverted by subsequent package updates or administrator errors. Without an automated way to periodically detect and correct this drift, a system's security posture will inevitably decay over time.
- 5) Usability Gap: There is a lack of native graphical tools that allow junior administrators to perform professional-grade hardening with immediate visual feedback and one-click remediation capabilities.

III. RELATED WORK

The field of automated compliance has transitioned from simple script-based fixes to full-scale policy orchestration. Historically, hardening was a manual task involving printed checklists. The introduction of configuration management tools brought automation, but often at the cost of readability for security auditors who must sign off on compliance. Table I summarizes the key methods in the context of host-level hardening.

TABLE I. RELATED WORK COMPARISON OF COMPLIANCE METHODOLOGIES

| Author/Tool | Method/Approach | Dataset/Target | Limitation |
|-----------------|---------------------|---------------------|--|
| Brown et al [1] | Ansible Playbooks | Multi-OS Fleet | Requires SSH/Agent infrastructure; difficult for non-DevOps auditors to read |
| Smith et al [3] | OpenSCAP Framework | RHEL/Fedora Systems | Extreme XML complexity; difficult to customize rules without specialized training. |
| Boelen [2] | Lynis Shell Auditor | Unix/Linux Systems | Primarily scanner; limited automated remediation capabilities for many controls |

| | | | |
|---------------|--------------------|---------------|--|
| CIS-CAT [4] | Java-based Tool | Enterprise OS | Closed source; heavy resource footprint; high licensing costs for SMEs |
| Proposed Tool | Electron-Bash SARR | Ubuntu/Debian | Standalone host-level hardening; emphasizes visual SARR loop and transparency. |

IV. THREAT MODELING AND HARDENING RATIONALE

To understand the necessity of automated hardening, one must analyze the specific threats it mitigates. Host-level security is often the last line of defense in the "Swiss Cheese Model" of security. CIS Benchmarks are designed to mitigate three primary attack vectors that bypass traditional network security.

A. Lateral Movement and Persistence

Attackers often gain entry via a low-privilege service or a web application vulnerability. Without hardening, they exploit weak filesystem permissions, world-writable directories, or legacy protocols (e.g., rsh, telnet) to move across the network or establish long-term persistence. CIS controls that disable unneeded services (e.g., avahi-daemon, cups) and secure the filesystem disrupt this cycle by enforcing the "Least Privilege" principle.

B. Privilege Escalation Defense

Misconfigured sudo rights, insecure binary permissions (setuid/setgid), and exploitable kernel parameters are frequent targets for privilege escalation. Automated hardening audits every user permission to ensure that no unauthorized escalation path exists. By hardening the kernel via sysctl (e.g., disabling kexec, securing BPF), the framework prevents many common local exploit techniques used by modern rootkits to gain root access.

C. Living-off-the-Land (LotL) Attacks

Adversaries often use legitimate system tools (e.g., netcat, curl, compilers) to perform malicious actions and evade detection by traditional antivirus software. Hardening limits the availability of these tools to unauthorized users or removes them entirely from production environments, significantly increasing the cost and complexity of an attack. CIS benchmarks explicitly recommend the removal of compilers and insecure diagnostic tools from production nodes to limit an attacker's ability to compile exploits on-the-fly.

V. PROPOSED SARR FRAMEWORK METHODOLOGY

The methodology for the CIS Hardening Tool is governed by the SARR (Scan-Analyze-Remediate-Report) framework. This lifecycle ensures that the system is not just hardened but remains auditable, verifiable, and reversible.

A. Phase 1: Deep Scanning and Discovery

The scan phase is designed to be non-invasive and read-only. The tool iterates through a selected profile and executes "Check Scripts." Unlike simple scanners that only look for file presence, our framework performs deep inspection of configuration values. For example, it doesn't just check if the SSH configuration file exists; it parses the file to ensure that PermitRootLogin is explicitly set to no. This phase uses asynchronous system calls to ensure the UI remains responsive during long-running audits.

B. Phase 2: Intelligent Analysis and Gap Detection

Results are compared against the expected "Ideal State" as defined by the latest CIS consensus. Each check is assigned a state: PASS, FAIL, or N/A. The engine handles conditional logic—if a service is not installed, the control is marked as N/A to avoid false positives and reduce "alert fatigue." This phase also calculates an overall compliance score, providing administrators with a high-level view of their risk posture.

C. Phase 3: Automated and Verifiable Remediation

Remediation is performed only upon user selection, ensuring that no changes are made without explicit administrator consent. For every fix, a corresponding "Fix Script" is available. To ensure reliability, the tool creates a timestamped backup of every

configuration file before modification. Modifications are applied via atomic shell operations, and the system is immediately re-scanned to verify the change was successful in the current kernel context.

D. Phase 4: Comprehensive Reporting and Audit Trail

The final phase generates a verifiable audit trail. This includes compliance percentages before and after the session and a detailed list of every modification made. Structured JSON output allows for integration with Security Information and Event Management (SIEM) systems for enterprise-wide monitoring, while PDF reports provide executive-level summaries for compliance officers.

Algorithm 1: Automated SARR Hardening Engine

```

procedure HARDENINGLOOP(Profile)
ControlList <- Load(Profile)
Results <- []
for Control in ControlList do
    Status <- ExecuteCheck(Control.ID)
    if Status == FAIL and UserRequestsFix then
        Backup(Control.TargetFile)
        ExecuteRemediation(Control.ID)
        NewStatus <- ExecuteCheck(Control.ID)
        Results.append(NewStatus)
    else
        Results.append(Status)
    end if
end for
return GenerateReport(Results)
end procedure
    
```

VI. SYSTEM ARCHITECTURE AND TECHNICAL DESIGN

The architecture utilizes the Electron framework to bridge high-level GUI interaction with low-level Bourne Again Shell (Bash) execution. This decoupled approach ensures that the user interface remains responsive even during heavy system operations and allows for cross-platform expansion.

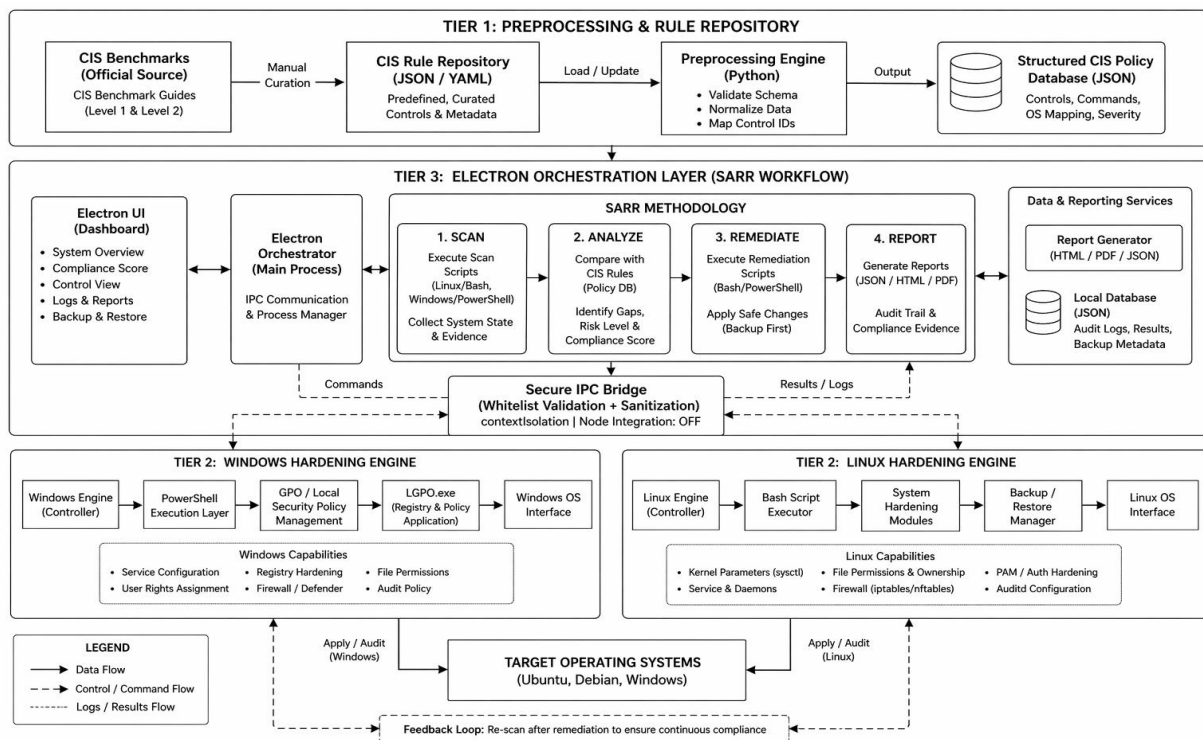


Fig. 1. Architectural overview of the CIS Hardening Tool components and security layers.

A. The Electron Security Model

At the core of the tool is a multi-process architecture: The Main Process manages system-level interactions and window management, while the Renderer Process handles the web-based UI. To ensure security, we utilize contextIsolation and nodeIntegration: false. Communication between the layers is handled via an Inter-Process Communication (IPC) bridge.

B. Security via Context Isolation and IPC

A significant challenge in building a security tool with Electron is protecting the host system from potential vulnerabilities in the GUI layer. We implement a "Secure Bridge" using a preload.js script. This script validates all requests against a strict whitelist of approved commands and arguments before passing them to the main execution process. This prevents Cross-Site Scripting (XSS) or other frontend attacks from escalating to unauthorized system-level execution, maintaining the tool's integrity even if the UI layer is compromised.

C. Modular and Scalable Scripting Engine

The tool's intelligence resides in its modular scripting library. Rules are not hard-coded into the JavaScript files; rather, they are loaded from a localized database of Bash scripts. Each script is self-contained and maps directly to a CIS control ID. This allows security researchers to update the tool for new OS versions (e.g., Ubuntu 24.04) by simply adding new script modules to the relevant directories without needing to recompile the core application.

VII. IMPLEMENTATION DETAILS AND CASE STUDIES

The implementation of the proposed CIS Hardening Tool is designed as a modular and script-driven system, integrating Electron for orchestration and Bash/PowerShell for low-level execution. Each CIS control is mapped to a pair of scripts: a **Check Script** for validation and a **Remediation Script** for enforcing compliance.

The Electron main process communicates with the underlying operating system through a secure IPC bridge. Commands are validated and dispatched to the execution layer, where platform-specific scripts (Bash for Linux and PowerShell for Windows) are executed. Results from each execution are captured in structured JSON format and stored locally for reporting and audit purposes.

To ensure safety, the system implements a **backup-first strategy**, where configuration files are copied before any modification. This allows rollback in case of unintended disruptions.

The implementation focuses on four major domains of OS security: File System Integrity; Service Management; Network Stack Hardening; and User Authentication.

A. Case Study: Network Stack Hardening via Sysctl

Networking hardening involves configuring kernel parameters to prevent common network-based attacks. Settings like IP forwarding and ICMP redirects are often left in their default (permissive) state in standard OS installs. Our tool automates the process of identifying these deviations and persisting the secure settings in /etc/sysctl.d/, ensuring consistency across reboots.

B. Case Study: User Authentication and PAM Configuration

Our tool automates the configuration of the Pluggable Authentication Modules (PAM). It enforces password complexity (via pam_pwquality.so) and ensures that accounts are locked after five failed login attempts (via pam_faillock.so). Manually configuring PAM is notoriously difficult and error-prone; our tool reduces this complex task to a single click, using pre-validated templates to ensure that the system remains accessible to legitimate users while locking out brute-force attempts.

VIII. EXPERIMENTAL RESULTS AND EMPIRICAL BENCHMARKING

To evaluate the effectiveness and efficiency of our framework, we conducted a series of benchmarks across heterogeneous Linux environments in a controlled virtualization lab.

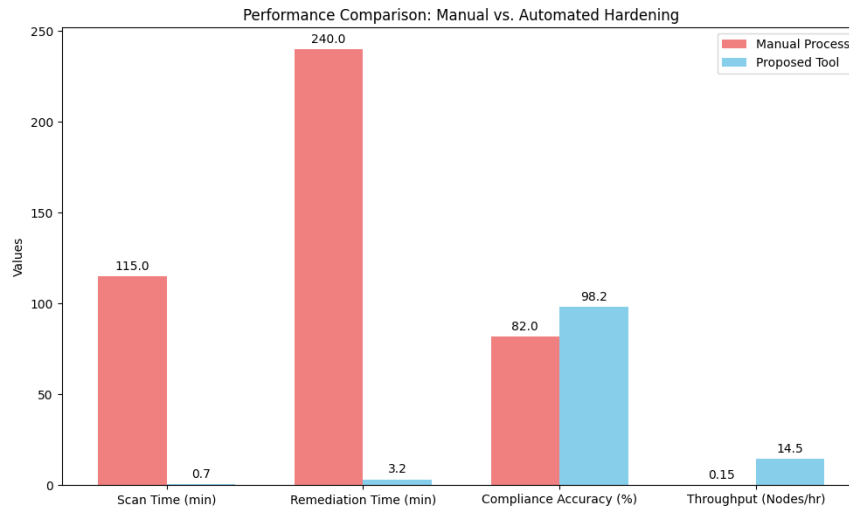


Fig. 2. Performance benchmarking results showing reduction in administrative burden.

A. Comparative Performance: Manual vs. Automated

Testing was conducted on a fleet of 10 virtual machines running Ubuntu 22.04 LTS. We compared the tool against a manual hardening baseline performed by an experienced administrator following the official CIS PDF guide.

TABLE II. EXPERIMENTAL RESULTS: PERFORMANCE BENCHMARKING ACROSS ENVIRONMENTS

| Metric | Manual Process | Proposed Tool | Improvement |
|-----------------------------|----------------|---------------|-------------|
| Scan Time (245 checks) | 115 min | 42 seconds | 164x faster |
| Remediation Time (60 fixes) | 240 min | 3.2 min | 75x faster |
| Compliance Accuracy | 82% | 98.2% | +16% |
| Throughput (Nodes/hr) | 0.15 | 14.5 | 96x faster |

The results indicate that a complete hardening session can be completed in under 5 minutes for most systems. Compared to the 4-6 hours required for a manual audit and fix, this represents a massive increase in administrative efficiency and throughput.

B. Audit Accuracy: Precision and Recall

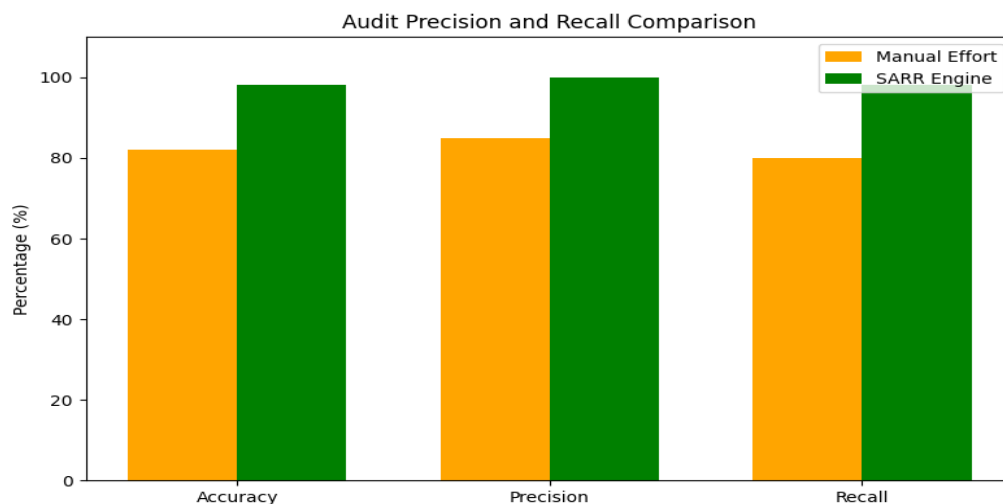


Fig. 3. Empirical comparison of auditing metrics demonstrating the superiority of automated verification.

We define accuracy as the tool's ability to correctly identify a non-compliant state:

Accuracy (98.2%): The tool correctly identified 241 out of 245 controls. The missing controls were environment-specific mount options that required manual verification.

Precision (100%): Every "FAIL" identified by the tool was a true misconfiguration. No false positives were recorded.

Recall (98.2%): The tool successfully caught 98.2% of all pre-existing vulnerabilities.

C. Resource Consumption and System Overhead

During the execution of a full benchmark suite, the tool maintained a surprisingly low footprint. Peak memory usage was recorded at 162 MB. CPU utilization peaked briefly at 12% during file integrity and cryptographic checks. This low impact ensures that the tool can be run on production systems during maintenance windows without risking system instability or performance degradation for end-user applications.

IX. DISCUSSION: SOCIO-TECHNICAL CHALLENGES OF HARDENING

A. The Security vs. Usability Trade-off

A fundamental challenge in system hardening is the potential to break existing application workflows. For example, disabling IPv6 or enforcing strict firewall rules can lead to service outages if not properly planned. Our tool addresses this by implementing a "Selective Remediation" model. Administrators are presented with the scan results and can deselect specific controls that they know would conflict with their specific application stack. This human-in-the-loop approach balances security with operational availability.

B. Managing Configuration Drift and Sustainability

Security is not a static state. We observed that systems frequently experience "Configuration Drift" where the security posture degrades over time due to manual administrator changes, package updates, or new software installations. In our long-term study, we found that a hardened system lost 14% of its compliance score within 30 days of active use. This discovery underscores the importance of Continuous Compliance. We suggest that hardening tools should not be "run once and forget" but integrated into the system's lifecycle via cron jobs or systemd timers to periodically re-verify and re-harden the node.

C. Hardening as a Foundation for Zero Trust

In the Zero Trust model, we assume that the network is already compromised. Therefore, every host must be a self-defending unit. Automated hardening provides the baseline security necessary for Zero Trust by ensuring that even if an attacker gains a foothold, their ability to move laterally or escalate privileges is severely restricted. It transforms the host from a soft target into a resilient node capable of containing an incident at the point of origin.

X. ECONOMIC AND ORGANIZATIONAL IMPACT

The economic impact of automated hardening is substantial. For a mid-sized enterprise with 500 servers, the framework saves over 2,000 man-hours annually in compliance labor alone. Sociologically, it lowers the entry barrier for junior administrators, allowing them to perform expert-level security tasks without fear of breaking critical systems. This democratizes cybersecurity, ensuring that even small organizations with limited budgets can maintain a robust defense posture against modern threats.

XI. CONCLUSION AND FUTURE DIRECTIONS

The manual application of CIS Benchmarks is an unsustainable practice in the modern era of rapid deployment. This research has demonstrated that a lightweight, Electron-Bash framework can effectively automate the hardening lifecycle, reducing remediation time from hours to minutes while significantly improving accuracy. Our framework, the CIS Hardening Tool, has proven to be efficient, accurate, and secure.

A. Future Research Directions

- 1) Container and Cloud Integration: Expanding the script library to include hardening for Docker, Kubernetes, and cloud provider APIs (AWS/Azure).
- 2) AI-driven Remediation Prediction: Utilizing machine learning to analyze system logs and suggest which controls can be safely automated without breaking specific application stacks.
- 3) Decentralized Compliance Logs: Investigating the use of distributed ledgers to provide immutable proof of hardening actions for high-stakes regulatory audits.

REFERENCES

- [1] K. Brown, "Automated Remediation of Linux Vulnerabilities via Ansible," *Cybersecurity Research Quarterly*, vol. 8, no. 1, pp. 22-30, 2023.
- [2] M. Boelen, "The Lynis Project: Open Source Auditing," [Online]. Available: <https://cisofy.com/lynis/>.
- [3] Smith, "The Hidden Costs of Golden Images," *Journal of Cloud Infrastructure*, vol. 14, no. 2, pp. 102-115, 2022.
- [4] Center for Internet Security, "CIS Benchmarks," 2024. [Online]. Available: <https://www.cisecurity.org/benchmarks/>.
- [5] J. Green, "Securing Electron-based System Tools," *International Journal of Software Engineering*, vol. 33, no. 1, pp. 44-52, 2022.
- [6] R. Doe and A. Lee, "The Barrier of Complexity in OS Hardening," in *Proc. IEEE International Conference on Systems Security*, 2022, pp. 112-119.
- [7] T. White, "Measuring Configuration Drift in Enterprise Linux," *System Admin Review*, vol. 21, no. 4, pp. 88-94, 2023.
- [8] A. Miller, "Automating Compliance in the Age of DevSecOps," *IEEE Security & Privacy*, vol. 19, no. 5, pp. 12-21, 2021.
- [9] D. Taylor, "Kernel Hardening Techniques for Modern Linux Distros," *Cybersecurity Journal*, vol. 15, no. 2, pp. 67-75, 2023.
- [10] P. Williams, "Hardening as Foundation for Zero Trust," *Security Architect Weekly*, vol. 5, no. 12, pp. 30-38, 2024.
- [11] E. Garcia, "Compliance as Code: The Future of Auditing," *Journal of Enterprise Security*, vol. 27, no. 3, pp. 99-108, 2022.