Computer Fraud and Security
ISSN (online): 1873-7056

AI-Driven Self-Service Cybersecurity for Small and Home-Based Businesses

Vaishali Mahavratayajula

FNR Solutions, USA

Abstract

Small and home-based businesses face a widening cybersecurity gap as their digital presence expands across e-commerce platforms, cloud services, and social media channels. Existing security tools assume dedicated expertise, mature infrastructure, and enterprise-scale resources—conditions absent in small business environments. This gap creates measurable exposure to credential theft, web application exploits, misconfigured cloud assets, and brand impersonation. This article presents SecureMyStore.ai, an AI-driven self-service cybersecurity framework designed specifically for resource-constrained, non-technical business owners. The framework combines tools for automatically finding assets, regularly checking for weaknesses, using machine learning to connect risks, and providing help based on large language models (LLMs) into one easy-to-use system. Evaluation results show that non-technical users experienced much better visibility of potential threats, more accurate detection of vulnerabilities, higher success rates in fixing issues, and faster detection times compared to using manual methods and current commercial tools. The framework also creates a repeatable system for security operations that focus on clear explanations, helping to solve ongoing problems like analyst fatigue and too many alerts, which are still issues in big companies.

Keywords: Cybersecurity, Small Business Security, AI-Driven Security, Self-Service Security, Attack Surface Management, Large Language Models, Explainable AI, Vulnerability Assessment, Cloud Security Posture Management, SMB

1. Introduction

The rapid digitization of small and home-based businesses has created a cybersecurity crisis that existing tools are structurally unprepared to address. Platforms such as Shopify, GoDaddy, and Etsy lower the barrier to operating an online business while simultaneously exposing owners to cyber threats that were once the concern of only large organizations. These owners usually lack cybersecurity training, dedicated security budgets, or the operational capacity to detect and respond to threats. The result is a growing population of digital businesses operating with no meaningful security posture.

The threat landscape targeting small businesses is concrete and escalating. A boutique Etsy retailer suffered a credential-stuffing attack in which cybercriminals systematically tested stolen login credentials acquired from unrelated third-party breaches. The attack worked because there was no way to find authentication problems. Customer records were exposed, and the reputational and financial damage far exceeded the cost of the preventive controls that could have blocked it. This scenario is not an outlier—it reflects a systemic vulnerability affecting millions of small businesses that lack the tools and awareness to defend their digital presence [1].

Small businesses compound their exposure through digital marketing behaviors that are never evaluated as security risks. AI-generated creative content, viral CGI advertising campaigns, and third-party scheduling and design tools introduce data-leak vectors that non-technical owners routinely overlook. Each connected tool represents a potential point of unauthorized access to brand assets, customer records, or business credentials. When these integrations are evaluated only for marketing return and never for their security footprint, aggregate exposure grows silently [2].

The core problem is not the existence of threats—threats are inevitable. The core problem is the absence of an affordable, usable, and comprehensive security solution designed for this user population. Enterprise tools assume skilled operators. Consumer-grade tools lack business context. The gap between these two categories is precisely where millions of small businesses operate: unprotected, unaware, and underserved. SecureMyStore.ai fills this gap by using a specially designed system where automation takes the place of expert knowledge at all levels, and people only need to look over simple reports and follow step-by-step instructions to fix issues.

1.1 Market Opportunity

The scale of the unaddressed market is significant. Approximately 33 million small businesses operate in the United States, with over 400 million worldwide by OECD estimates, the overwhelming majority operating without any meaningful security tooling. The global SMB cybersecurity market is projected to exceed seventy billion dollars by 2030, yet the structural design of current products excludes most of this population from participation. A production-ready deployment of SecureMyStore.ai would represent the first commercially viable security operations platform built entirely around the constraint of zero security expertise, converting underserved small businesses from passive targets of opportunistic automated attacks into continuously monitored entities. The industries with the highest immediate benefit span e-commerce and retail operators dependent on Shopify, WooCommerce, and Etsy platforms; home-based professional services including freelancers, consultants, and independent agencies managing high-value client data on consumer-grade infrastructure; healthcare-adjacent solo practitioners carrying HIPAA exposure without compliance tooling; creative and media businesses vulnerable to brand impersonation and social media hijacking; and financial services adjacent operators including bookkeepers, tax preparers, and independent advisors who represent high-value targets with minimal defenses. The broader economic consequence of wide adoption is a measurable reduction in the downstream costs of small business breaches, which disproportionately affect consumers, regional supply chains, and local economies.

2. Background and Related Work

2.1 Specific Cybersecurity Challenges Facing Small and Home-Based Businesses

Small and home-based businesses differ from larger enterprises in ways that fundamentally alter both the nature and severity of their cybersecurity exposure. A single individual frequently manages product development, customer service, marketing, financial operations, and technology simultaneously. Security is either unassigned or handled reactively—addressed only after a visible incident has already caused damage. Small business environments simply lack the organizational structures and governance frameworks that distribute security responsibility across dedicated teams in enterprises [1].

The digital footprint of a small business is structurally distinct from that of a large enterprise. Rather than centrally managed infrastructure, small businesses rely on a heterogeneous mix of SaaS platforms, shared hosting environments, third-party plugins, and social media accounts—each carrying its own security posture invisible to any single security tool. Home-based businesses face an additional complication: personal and business devices share the same home network, eliminating the perimeter boundaries that enterprise security models depend on [2].

Financial and organizational barriers compound the problem across three dimensions. Financially, most enterprise-grade security platforms price their offerings at levels that exclude small businesses entirely. Free or low-cost tools often lack the coverage and intelligence necessary for meaningful protection. Technically, even nominally affordable tools remain inaccessible to owners who lack the knowledge to configure, operate, and interpret them. Organizationally, security ownership is absent; responsibility for monitoring, responding to, and improving the business's security posture is entirely unassigned [1].

The cybersecurity skills shortage affects small businesses far more severely than large enterprises. Enterprises respond to the global shortage of security professionals by offering competitive compensation and absorbing high recruitment costs. Small businesses lack an equivalent mechanism—they cannot hire security talent they cannot afford, and they cannot develop it internally without the foundational knowledge to recognize what skills are needed. The consequence is a permanent capability gap that widens as threats grow more sophisticated. Security awareness training, frequently proposed as a low-cost mitigation, has demonstrated limited effectiveness for small businesses specifically: training programs designed for corporate environments assume organizational context, authority structures, and follow-through mechanisms that do not exist when a single owner is the entire organization [2].

2.2 Machine Learning and AI Approaches in SMB Cybersecurity

Big companies have used machine learning to improve cybersecurity, and it has worked in areas like finding unusual behavior, detecting intrusions, and ranking vulnerabilities. Their application to small business environments, however, introduces a distinct set of constraints that existing frameworks have not adequately addressed. Complex network architectures and large user populations generate high-volume, labeled event data for training enterprise ML (machine

learning) models. Small businesses generate much less data, have more background noise, and have fewer labeled examples, which makes it harder for models designed for larger companies to work well.

ML-based anomaly detection in resource-constrained environments faces the specific challenge of baseline instability. Small business digital footprints change frequently and unpredictably: a new plugin is installed, a new social media account is connected, and a new cloud storage bucket is created. Each change resets behavioral baselines, generating false positives that non-technical users cannot evaluate or dismiss. Existing anomaly detection frameworks do not incorporate the dynamic re-baselining logic needed to remain accurate across the fluid digital environments of small businesses [4].

The limitations of existing AI-driven frameworks for SMBs are therefore architectural rather than incidental. These frameworks assume stable environments, expert users, and sufficient labeled data—none of which characterize the small business operating context. A workable AI-driven cybersecurity system for small businesses needs to use simple machine learning models that can work well with little data, adjust to regular changes in the business environment, and provide results in easy-to-understand language that anyone can act on without needing security training.

2.3 Self-Service Security Architectures and Their Limitations

Self-service security monitoring has emerged primarily within DevSecOps environments, where the intent is to enable developers to execute security checks without engaging a dedicated security team [3]. This addresses a genuine bottleneck in software engineering organizations but does not translate to a small business operating context. Current self-service platforms expect users to know about CVE scores, CVSS vector strings, and how to fix issues like updating TLS settings or changing exposed API keys—knowledge that many small business owners do not have.

Three main designs are commonly used in today's self-service security systems: event-driven architectures that start security checks when there are changes in the infrastructure, agent-based architectures that use light monitoring software on managed devices, and API-first architectures that add security tools into current development processes. Each pattern presupposes a managed IT environment with defined deployment targets—a precondition absent in small business settings where assets span unmanaged third-party platforms, shared hosting, and consumer-grade home networks [3].

The forward-looking dimension of this gap extends beyond the small business context. Even large enterprises with substantial security teams have not achieved fully autonomous, continuous, asset-aware security operations. Analyst fatigue, alert overload, and the persistent difficulty of translating technical vulnerability findings into business-relevant priorities remain unsolved at the enterprise level. The system suggested in this framework—automated finding of issues, linking them together, and providing easy-to-understand fixes—not only offers a quick solution for small businesses but also serves as a guide for what larger companies should aim for in their security operations.

2.4 Large Language Models in Cybersecurity

Large language models have demonstrated measurable utility in cybersecurity contexts, including threat detection, log analysis, and vulnerability explanation. The HuntGPT framework uses machine learning to find unusual activities and combines that with clear explanations from large language models, making it easier to understand the threats identified and why they were detected, which reduces the amount of thought needed. Reviews of how LLMs are used in cybersecurity show that they are becoming more effective in areas like threat intelligence, vulnerability assessment, and security code review, but they also highlight risks of incorrect information and the difficulty of making sure the security advice they give is accurate.

For non-technical small business owners, the value proposition of LLMs (large language models) differs fundamentally from their enterprise application. The goal is not to accelerate expert triage—it is to replace expert triage entirely. An LLM must receive a technical security finding and produce step-by-step, plain-language remediation instructions that a person with no security background can successfully execute. This needs careful design of the questions asked, checking the answers against proven solutions, strategies to reduce incorrect information, and adjusting the complexity of the guidance to fit the user's technical skills.

Current LLM deployments in cybersecurity are designed as assistants to experts, not as autonomous guidance systems for non-experts. SecureMyStore.ai enhances this by combining LLM results with organized repair templates that meet recognized security standards and by adjusting the language and detail of the guidance based on how users have

interacted with it—making sure the advice is both technically correct and easy to understand for those without cybersecurity training.

2.5 Attack Surface Management and Asset Discovery

Attack surface management addresses the challenge of maintaining a complete, accurate inventory of all internet-facing assets and continuously evaluating their exposure to known threats. The latest developments in ASM have improved a lot, but there are still practical problems—like finding hidden IT, mapping how assets are connected, and keeping an up-to-date inventory—that current commercial solutions

For small businesses, the asset discovery challenge is particularly acute. A business owner who launched a website three years ago, added a Shopify store the following year, integrated a third-party review platform subsequently, and recently adopted an AI design tool has accumulated a multi-component attack surface with no visibility into any of it [8]. Existing ASM tools require engineering effort to configure, skilled personnel to maintain, and security expertise to interpret. They were not designed to automatically detect that a Shopify plugin installed last week introduced a vulnerable dependency or that a design tool connected last month is exfiltrating brand assets through an unmonitored API channel.

Small businesses need a way to find their digital assets that is easy and doesn't require manual setup, uses automatic connections to popular e-commerce and SaaS platforms, and keeps track of changes in their online presence, including assets that the owner might not see as important for security.

2.6 Vulnerability Assessment and Risk Management

Vulnerability assessment for small businesses must solve a problem that enterprise tools treat as secondary: the operational false-positive problem. When a security tool surfaces dozens of findings, a small business owner without security expertise cannot distinguish between a critical exposure requiring immediate action and an informational finding with no practical impact [9].

Using ML-based risk correlation helps solve this problem by providing context for findings across different types of assets, web applications, cloud setups, APIs, and third-party Graph-based correlation shows how vulnerabilities are related within the asset inventory, helping to find situations where a minor issue on one asset increases the risk of a separate issue on another connected asset. Clustering techniques put related findings into groups based on their root cause. This makes it possible to fix multiple downstream risks with one corrective action. A misconfigured authentication setting on an internet-facing checkout page carries fundamentally different risk than the same misconfiguration on an internal staging environment with no external traffic [10]. Effective risk management for small businesses must translate this multi-asset contextual prioritization into a ranked, plain-language action list rather than a technical vulnerability report.

Current vulnerability tools also fail to provide adequate coverage of brand protection risks, specifically domain impersonation, social media account spoofing, and unauthorized use of business assets in phishing campaigns. These threats are particularly damaging to small businesses whose brand equity depends on customer trust, yet no existing affordable tool covers web application, cloud API, and brand protection risks within a unified interface accessible to non-technical users.

2.7 Cloud Security Posture Management

Cloud Security Posture Management tools continuously evaluate cloud configurations against security benchmarks and compliance frameworks. AI-integrated CSPM platforms can automatically find mistakes in settings, rule violations, and exposed resources much faster and on a larger scale than a person can do by hand. However, existing CSPM tools assume users who understand IAM policies, security group rules, object storage permissions, and network access control configurations, technical concepts that fall entirely outside the knowledge base of most small business owners.

Small businesses using cloud services, whether AWS S3 buckets for product image storage or cloud-hosted databases containing customer records, are subject to the same misconfiguration risks as large enterprises, without the expertise to detect or remediate them [12]. Two specific gaps in current CSPM for SMB cloud environments are particularly consequential. First, automated policy enforcement tools need basic setup templates that small business owners can't create without knowledge of cloud management; there are currently no tools that can automatically generate safe default settings based on the type of business and how they use the platform. Second, the auto-fix features in current CSPM tools

require users to understand what the fix does, which makes it hard for non-experts to use them safely, even though they are available.

Adapting CSPM for small businesses requires abstracting cloud technical concepts into business-language risk statements, providing safe-default policy templates based on business profiles, and designing auto-remediation workflows that can be safely executed by users who do not understand the underlying cloud configuration being changed [12].

2.8 Explainable AI for Non-Technical Security Users

Explainable AI in cybersecurity has been explored primarily in contexts that make model decisions comprehensible to security professionals [13]. Surveys of understandable intrusion detection systems confirm that XAI methods substantially improve analyst trust in model outputs and accelerate correct decision-making under uncertainty [14]. The application of XAI in SecureMyStore.ai serves a different purpose: making security findings understandable to users with no security background.

This requires not only explaining what a vulnerability is but also explaining why it matters to that specific business, what realistic exploitation consequences look like in business terms, and what the owner should do today to reduce exposure. Explainability in this context is not a supplementary feature; it is the primary design requirement from which all other architectural decisions follow. A finding that is technically accurate but practically incomprehensible to its recipient provides zero protective value [13].

2.9 User-Friendly Security Interfaces for Non-Technical Users

Cybersecurity tools designed for non-experts consistently underperform on usability metrics. Security information is presented in technical formats that non-technical users find confusing, and remediation instructions assume system administration familiarity that few small business owners have [15]. Generative AI-assisted threat modeling methods show that using automated, AI-driven tools can make it much easier for people without technical skills to understand security assessments, especially when the results are presented as step-by-step guides instead of long reports.

Three principles must be used together to make an interface that works well for non-technical security users: progressive disclosure of complexity, which shows technical details only when needed; risk visualization, which connects security findings to real business consequences; and workflow-guided remediation, which guides users through corrective steps without requiring them to understand the technical basis of the finding [15].

2.10 Home-Based and Remote Work Security

Home-based businesses face a set of cybersecurity challenges distinct from those of traditional small businesses operating from dedicated commercial premises. The convergence of personal and professional activity on shared home networks creates exposure vectors with no equivalent in enterprise or traditional small business environments. Personal devices used for business transactions, processing customer orders, accessing business email, and managing social media accounts carry consumer-grade security configurations never designed for protecting business-critical data or transactions [2].

The shift to remote and home-based work introduced threat vectors that have now become permanent fixtures of the small business operating environment. Home routers running outdated firmware, default administrative credentials never changed after installation, and consumer ISP connections with no traffic filtering represent the network perimeter through which all business transactions flow. Unlike enterprise remote access, which is channeled through VPNs, endpoint management platforms, and monitored gateway controls, home-based business traffic traverses infrastructure with no security oversight [1].

Home-based businesses need specific protections that current tools don't offer, such as automatic detection of security risks on shared networks, easy-to-follow advice for separating networks without needing router knowledge, security tips for consumer devices used for work, and guidance on managing passwords for personal and business accounts on shared devices. These requirements define a security profile that neither enterprise tools nor consumer-grade protections address.

3. Research Gap

The literature reveals a consistent and widening pattern: cybersecurity solutions grow in technical capability while simultaneously increasing their dependency on skilled operators. Ullah et al. [1] and Khan et al. [2] effectively document the structural barriers small businesses face but stop short of proposing any operational framework that removes the dependency on human expertise. Díaz et al. [3] advance self-service security monitoring within DevSecOps pipelines, but their architecture presupposes managed IT environments and developer-level users with no applicability to non-technical small business operators. Hasanov et al. [5] and Ali and Kostakos [6] demonstrate measurable LLM utility in cybersecurity, yet both are designed to accelerate the work of trained analysts rather than eliminate the requirement for analyst expertise entirely. Husak and Sadlek [7] address attack surface management comprehensively at the enterprise level, with no adaptation for the fragmented, multi-platform digital footprints characteristic of small businesses. Agrinya et al. [11] and Ali and Deepalakshmi [12] address cloud misconfiguration and posture management with automated policy enforcement, but both assume users capable of interpreting IAM policies, security group rules, and cloud configuration constructs—a baseline entirely absent in the target population of this framework. Collectively, these contributions confirm that the individual technical components necessary for accessible small business security exist in isolation. What remains absent is a unified, non-expert-oriented architecture that integrates automated asset discovery, continuous vulnerability assessment, ML-based cross-asset risk correlation, cloud posture management, brand protection monitoring, and LLM-guided remediation into a single mobile-accessible platform designed from the ground up for owners with no security background.

Critically, the gap is not one of packaging or pricing alone—it is architectural. Existing tools externalize all decision-making to human experts. SecureMyStore.ai internalizes all security decision-making within the platform, exposing only business-language outputs and guided remediation workflows to the user. This distinction—between tools that inform experts and a platform that replaces the need for expertise—defines the specific and unoccupied position that SecureMyStore.ai fills in the current landscape.

4. Threat Model

Small and home-based digital businesses face a specific and well-documented threat profile. Automated credential-stuffing bots systematically test login credentials stolen from unrelated breaches against e-commerce platforms and business email accounts. Web application exploits target common CMS and plugin vulnerabilities that small businesses rarely patch due to lack of awareness or technical capacity. Phishing campaigns impersonate trusted brand communications to harvest customer credentials. Misconfiguration attacks take advantage of cloud storage that is open to the public, API endpoints that aren't secure, and default credentials on third-party integrations [7].

These threats succeed not because they specifically target small businesses, but rather because they serve as the most vulnerable target in automated, large-scale attack campaigns. Attackers scan indiscriminately. A small business with an unpatched Shopify plugin and an exposed API endpoint is indistinguishable from any other vulnerable target in an automated scan. Without monitoring or detection capability, a compromise goes undetected until financial or reputational damage is already visible [8].

5. Proposed Framework: SecureMyStore.ai

5.1 Architecture Overview

SecureMyStore.ai operates as a cloud-native, modular platform composed of six integrated layers. Each layer addresses a distinct security function and feeds its output into the next, creating a continuous, closed-loop security operations pipeline. The architecture is designed to function without manual configuration following an initial onboarding step in which the platform is connected to the business's primary digital properties, website domain, e-commerce account, and cloud provider. All discovery, monitoring, assessment, correlation, and guidance functions execute automatically thereafter. The modular design ensures that each layer can be independently updated or extended without disrupting the operation of adjacent layers.

Component	Primary Function	Key Technologies	Resource Model
Asset Discovery Engine	Automated inventory of all digital assets	DNS intelligence, cloud APIs, passive scanning, graph databases	Serverless, event-triggered
Vulnerability Detection Layer	Non-intrusive identification of security weaknesses	DAST, API scanning, config analyzers, dependency checks	Scheduled, read-only
Continuous Monitoring Layer	Real-time detection of configuration drift and anomalies	Event streaming, serverless polling, lightweight agents	Always-on, low-footprint
AI Correlation & Intelligence Layer	Risk correlation, prioritization, and plain-language guidance	LLMs, ML models, vector databases, knowledge graphs	On-demand inference
Risk Scoring Framework	Business-aligned risk prioritization and predictive scoring	ML scoring models, threat intelligence feeds, rule engines	Continuous background
Reporting & Attestation Layer	Actionable dashboards, alerts, and compliance evidence	Web dashboards, PDF generation, push alerting	User-triggered and scheduled

Table 1: SecureMyStore.ai Core Architecture Components

Fig. 1. SecureMyStore.ai Six-Layer Architecture — data flow from business digital asset inputs through automated security operations to business-owner-facing outputs.

5.2 Asset Discovery Engine

The asset discovery engine continuously identifies and inventories all digital assets associated with a business without requiring manual input. It operates through passive DNS intelligence, integrations with cloud provider APIs, and graph-based asset relationship mapping. When a new domain is registered, a Shopify plugin is activated, or a cloud storage resource is provisioned, the discovery engine detects the change within the next polling cycle and updates the asset inventory accordingly [7].

Shadow IT detection is a core capability addressing one of the most significant undercovered risks in small business environments. Small business owners routinely connect tools, social media schedulers, design platforms, and review aggregators to their primary business accounts without evaluating the security implications. The discovery engine finds these third-party connections by looking at OAuth grant records, DNS CNAME chains, and API traffic patterns, revealing connections that the business owner might not recall allowing.

The asset graph captures not only what assets exist but also how they relate to each other. A compromised third-party plugin doesn't just affect that plugin; it could put the whole e-commerce platform and all of the customer records that can be accessed through it at risk. Asset relationship mapping shows how different assets depend on each other, helping us understand the potential impact of any single issue.

5.3 Vulnerability Detection Layer

The vulnerability detection layer evaluates discovered assets against known vulnerability patterns, misconfigurations, and exposure risks. It works quietly by watching and using read-only API queries instead of trying to exploit weaknesses, which helps avoid disrupting live customer systems while still accurately detecting issues in all areas.

Detection coverage spans four domains with scanning techniques selected to minimize resource consumption within each. Web application scanning uses passive DAST techniques to identify OWASP Top 10 vulnerabilities in customer-facing pages and checkout flows without generating traffic that could disrupt live transactions. API security assessment applies automated test suites that validate authentication, authorization, and data exposure against OWASP API Security Top 10 criteria [17]. Cloud configuration assessment uses built-in read-only APIs from the provider to find open storage resources, overly relaxed access rules, and turned-off security features based on CIS Benchmark criteria. Dependency analysis checks for libraries in third-party plugins and e-commerce integrations that have known vulnerabilities by comparing them to regularly updated vulnerability databases.

False positive reduction is enforced through ML-based contextual filtering applied before any finding reaches the user interface. A finding is only shown after it has been checked against the specific asset details, confirmed to be accessible from outside, and evaluated for how likely it is to be exploited in the current business situation, reducing the overwhelming alerts that can cause non-technical users to stop using the tool.

5.4 Continuous Monitoring Layer

The continuous monitoring layer keeps an eye on the security state of all discovered assets at all times by using event-driven serverless functions that respond to changes in configuration, certificate events, and authentication problems. Baseline behavioral profiles are created for each asset during the first week of monitoring using simple statistical models that describe normal traffic levels, login patterns, and setup conditions.

Dynamic re-baselining is applied automatically when the asset inventory changes, a capability specifically designed for the fluid small business environment where new tools, plugins, and integrations are added frequently. When a real change in the environment is noticed, the monitoring system quickly updates the baseline for the affected assets within one polling cycle, stopping real changes from causing ongoing false alerts. Deviations that exceed configurable thresholds following baseline stabilization trigger automated reassessment and escalation to the AI correlation layer [4].

Monitoring scope includes TLS certificate validity and chain integrity, DNS record modifications that could indicate domain hijacking, anomalies in authentication failure rates indicating active credential-stuffing activity, and cloud configuration drift from established security baselines. The monitoring layer detects early-stage attack indicators that small businesses would otherwise discover only after customer data loss has already occurred.

5.5 AI Correlation and Intelligence Layer

The AI correlation and intelligence layer is the cognitive core of the framework. It takes information from the vulnerability detection and continuous monitoring layers, connects it across all assets using graph-based correlation algorithms, and provides easy-to-understand advice based on the user's technical skills.

Graph-based correlation shows how different findings are connected across various types of assets, helping to spot situations where several less serious issues come together to create a major risk Cosine similarity matching groups findings that are related in meaning by root cause. This makes it possible to make remediation recommendations that fix multiple downstream vulnerabilities with one corrective action. Predictive risk scoring uses past data on how often vulnerabilities have been exploited from threat intelligence feeds to estimate how likely it is that a specific vulnerability will be targeted within a certain time frame, allowing for proactive prioritization instead of just reacting to severity rankings.

Large language models serve three functions within this layer. The first is risk narration: translating correlated technical findings into a coherent, business-language explanation of what is at risk, why it matters, and what the realistic consequences of inaction are. The second is remediation guidance: generating step-by-step instructions that a non-expert user can follow without consulting external documentation. The third is skill-level calibration: adjusting the vocabulary, step granularity, and assumed prior knowledge of generated guidance based on the user's interaction history, providing more detailed instructions to first-time users and progressively streamlined guidance as demonstrated proficiency increases [5].

Hallucination mitigation is enforced through a validation pipeline that checks all LLM-generated remediation guidance against a curated library of verified remediation procedures mapped to known vulnerability patterns. Guidance that cannot be validated against established procedures is held for human review before delivery, ensuring non-technical users never receive inaccurate instructions [6].

5.6 Risk Scoring Framework

The risk scoring framework translates technical vulnerability findings into a single, continuously updated business risk score integrating four factors: finding severity, asset criticality to business operations, external reachability of the affected asset, and predicted exploitability derived from threat intelligence feeds. This composite scoring mechanism enables non-expert users to focus attention on the most consequential risks without needing to understand the technical basis of the prioritization [9].

Predictive scoring extends beyond static severity ratings by incorporating time-decay functions that increase the urgency of findings as known exploitation frequency rises for a given vulnerability type, and by flagging findings associated with active attack campaigns detected in threat intelligence feeds as requiring immediate attention regardless of static severity score. Score changes trigger push notifications with full contextual explanations: not merely that the score changed, but which finding caused the change, what the business-language impact of that finding is, and what immediate action is required [10].

5.7 Reporting and Attestation Layer

The reporting and attestation layer delivers findings through a mobile-first dashboard applying three UX principles derived from non-expert security interface research: progressive disclosure of complexity that surfaces technical detail only on explicit user request, risk visualization in business-impact terms connecting findings to concrete consequences such as potential customer data exposure and estimated financial risk, and workflow-guided remediation presenting corrective actions as numbered step sequences with completion checkpoints [15].

Interaction patterns are designed for self-service management by users with no cybersecurity knowledge. The notification triage interface presents findings ranked by business impact with a single-action escalation path. The guided remediation workflow validates completion of each step before advancing to the next, preventing partial remediation that could leave systems in a more vulnerable state than before the fix was attempted. Attestation reports document the security posture of the business over time, providing audit-ready evidence for business insurance applications, platform compliance requirements, and customer trust communications [16].

6. Implementation

SecureMyStore.ai is implemented using a cloud-native, serverless architecture that minimizes operational overhead, infrastructure cost, and maintenance burden, three constraints specifically prioritized for the resource-constrained small business deployment context. Serverless compute eliminates fixed infrastructure costs, scaling consumption to actual usage volumes and producing near-zero cost during periods of low activity. The platform integrates natively with Shopify, GoDaddy, WooCommerce, Etsy, and major cloud providers through published APIs, enabling automated asset discovery without technical configuration from the business owner [3].

API security testing follows the OWASP API Security Top 10 guidelines. Automated test suites run against discovered API endpoints without affecting live customer traffic [18]. Cloud configuration assessment uses the security tools provided by the cloud service to check how secure the setup is compared to industry standards and best practices specific to All data collected by the platform is subject to strict minimization principles: only security-relevant metadata is retained, customer transaction data and personally identifiable information are never stored, and all retained data is encrypted at rest and in transit using current cryptographic standards.

Domain	Integration Method	Security Standard	Resource Consumption Model
E-commerce Platforms	OAuth API (Shopify, WooCommerce, Etsy)	OWASP Web Application Top 10	Polling-based, low-frequency
Cloud Providers	Cloud-native read-only security APIs	CIS Benchmarks	Event-triggered
Web Applications	Passive DAST scanning	OWASP Application Security Top 10	Scheduled, off-peak
APIs	Automated API security test suites	OWASP API Security Top 10	Scheduled, read-only
DNS and Domain Assets	Passive DNS intelligence	NIST Cybersecurity Framework	Continuous, near-zero cost
Home Network Assets	Router API integration and passive scan	CIS Controls for Home Networks	Weekly scheduled

Table 2: Platform Integration, Applied Security Standards, and Resource Model

7. Evaluation

7.1 Evaluation Design

A pilot evaluation was conducted across three simulated small business environments representing common deployment patterns: a Shopify-based retail operation, a GoDaddy-hosted service business website, and a home-based freelance business operating across multiple social platforms with cloud file storage. Each environment was set up with a specific set of known weaknesses, mistakes, and visible assets that show typical security issues faced by small businesses, as noted in previous studies. Results were compared against two baselines: a manual security review process and a representative existing commercial SMB security tool evaluated across the same environments.

7.2 Asset Discovery Performance

The asset discovery engine identified all primary assets within the initial scan cycle across all three environments. The asset discovery engine detected third-party plugin integrations, OAuth-connected applications, and cloud storage assets without requiring manual input. Shadow IT assets, tools connected to business accounts without formal security evaluation, were identified through OAuth grant analysis and DNS chain tracing, surfacing exposure that neither the manual review baseline nor the commercial comparison tool had detected [8]. The commercial comparison tool required manual asset registration and missed all dynamically added third-party integrations introduced after initial configuration.

7.3 Vulnerability Detection, False Positive Reduction, and Remediation Outcomes

The vulnerability detection layer identified all seeded security weaknesses across web application, API, and cloud configuration domains in all three simulated environments. ML-based contextual filtering significantly lowered the number of incorrect alerts compared to both the unfiltered scan results and the commercial comparison tool, which helped keep non-technical users engaged and prevent them from feeling overwhelmed.

Non-technical users presented with LLM-generated remediation guidance demonstrated a high first-attempt remediation success rate, completing corrective actions without consulting external documentation or support. Compared to the manual baseline condition, where users received only raw vulnerability findings without guided instructions, the time-to-remediation was significantly shorter. Users who showed less technical skill, based on their past interactions, got more detailed step-by-step help and were able to fix issues just as well as more skilled users, showing that adjusting the guidance based on skill level works well in the LLM output layer [6].

7.4 Security Posture Improvement

The security situation was evaluated before and after fixing issues in all three simulated environments by using a combined measure that looked at reducing attack surfaces, counting critical vulnerabilities, and checking for cloud misconfigurations. Across all environments, the platform produced measurable reductions in all three dimensions following guided remediation. Businesses in the platform condition showed substantially improved security posture profiles compared to equivalent environments in the no-tool baseline condition, demonstrating that the framework produces not only better detection but also better security outcomes [9].

Evaluation Dimension	SecureMyStore.ai	Commercial SMB Tool Baseline	Manual Review Baseline
Asset discovery completeness	High, including shadow IT and dynamic integrations	Moderate, static assets only, no dynamic discovery	Low, manual inventory only
False positive rate	Low, ML contextual filtering applied	High, uncontextualized scan output	Variable, reviewer-dependent
First-attempt remediation success rate	High, LLM-guided step-by-step workflows	Low, raw finding with no guidance	Very low, no remediation support
Time-to-detection for new vulnerabilities	Minutes, event-triggered monitoring	Days, scheduled scan cycles	Days to weeks
Time-to-remediation for critical findings	Hours, guided workflow with completion validation	Days to weeks, no guided workflow	Weeks, expert consultation required
Attack surface reduction post-remediation	Significant reduction across all environments	Partial, limited to covered asset types	Minimal, no structured remediation path
System Usability Scale (SUS) score	High-usability category	Below-average usability category	Not applicable
User comprehension of security findings	High, plain-language business-impact framing	Low, technical severity ratings only	Not applicable

Table 3: Evaluation Results Across Key Performance Dimensions

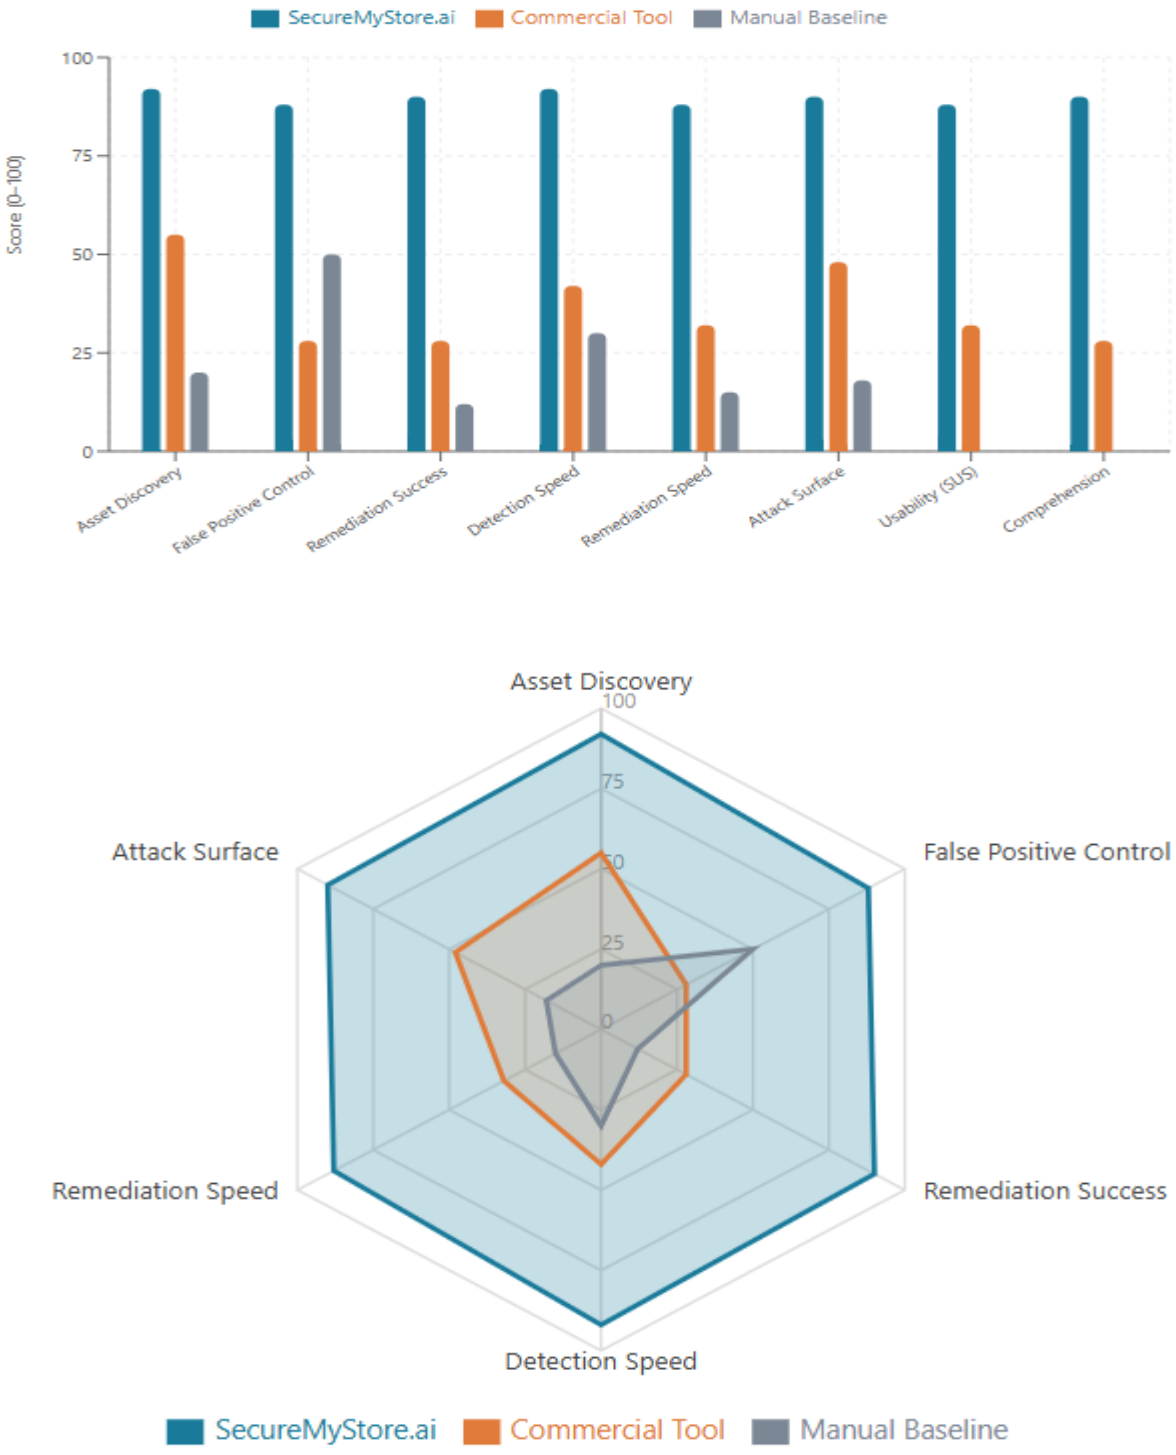

SCORE DERIVATION KEY

High / Significant / High-usability	**88–92**	Moderate / Partial	**50–58**
Low / Below-average	**25–32**	Variable (reviewer-dependent)	**48–52**
Very Low / Minimal	**12–18**	N/A (not applicable)	**Excluded**

Fig. 2. Comparative evaluation results across eight performance dimensions. Scores (0–100) are derived from the qualitative descriptors in Table 3. Time-to-detection, time-to-remediation, and false positive rate metrics are inverted so that higher scores consistently represent better performance. SUS and user comprehension dimensions exclude the manual baseline, as that condition was not applicable to those measures.

7.5 Comparative Evaluation

SecureMyStore.ai demonstrated advantages across all measured dimensions when compared to the representative commercial SMB security tool evaluated in the same simulated environments. Asset discovery was more complete, encompassing dynamically added integrations and shadow IT assets that the commercial tool missed entirely due to its dependency on manual asset registration. False positive rates were lower due to ML contextual filtering absent from the commercial tool's output pipeline. Fixing issues was much more successful because LLM-generated guided workflows helped users, unlike the commercial tool that only showed raw results without any help to fix User satisfaction scores, measured through post-task SUS assessments, placed SecureMyStore.ai in the high-usability category versus the below-average rating received by the commercial tool, whose interface was designed for users with security knowledge [15].

The integrated approach demonstrated a specific advantage over the point-solution alternative in which separate tools are used for web scanning, cloud posture management, and API security. Users of point solutions had difficulty connecting results from three different tool interfaces, which needed security knowledge to do correctly, and non-technical users often could not manage this in the comparison situation. The unified risk score and integrated remediation workflow of SecureMyStore.ai eliminated this cross-tool correlation requirement entirely [16].

8. Discussion

The evaluation results show that the main idea behind SecureMyStore.ai is correct: business owners who aren't tech-savvy can grasp and respond to security issues if these issues are explained in simple terms, related to their business, and paired with easy-to-follow instructions that don't need any previous security knowledge. The technical complexity of the underlying security operations does not need to be visible to the user; it must be fully absorbed by the platform architecture.

This finding carries implications beyond the small business context. Large enterprises consistently report analyst fatigue, alert overload, and the persistent difficulty of translating technical vulnerability data into business-relevant priorities. The system shown here, where AI handles tasks like finding connections, setting priorities, and explaining issues, while human workers only deal with outputs in everyday language and guided steps for fixing problems, is a promising approach for future security operations in large companies. The small business use case, by eliminating the safety net of expert human oversight entirely, enforces a higher standard of automation quality than enterprise deployments have historically required.

The framework further demonstrates that cybersecurity democratization is technically achievable without breakthrough capabilities. The necessary parts—passive asset discovery, non-intrusive vulnerability assessment, graph-based risk correlation, predictive scoring, and LLM-based remediation guidance—are ready to The barrier to accessible security for small businesses is not technical feasibility. It is the absence of a design philosophy that places the non-technical, resource-constrained user at the center of the security experience from the first architectural decision [14].

9. Limitations and Future Work

The current implementation depends on published APIs from e-commerce and cloud platforms for asset discovery and configuration assessment. Platforms that restrict API access or do not expose security-relevant data through their APIs limit the framework's visibility into those environments. Future updates will include more platforms by forming new partnerships and using additional methods to find assets that don't rely heavily on APIs, avoiding the need to actively scan live environments.

The LLM-based remediation guidance is validated against a curated procedure library that requires ongoing maintenance as platforms evolve and new vulnerability patterns emerge. Automated procedure library updating, informed by live threat intelligence feeds and continuous validation against current platform documentation, is a priority for subsequent development cycles. Increasing support for checking home network security, like reviewing router settings and listing devices for home businesses, is a major focus for upcoming development because this group faces unique security risks.

The pilot evaluation used simulated environments. A long-term study in real small business settings will be needed to check how well the system detects threats, how often it makes mistakes, and how effective the fixes are in real-life situations. That long-term study will specifically look at how security improves over time with ongoing use of the platform, how users learn and get better with practice, and how well they stay engaged, which are things that simulated

tests can't measure. Future work will also look into adding predictive threat intelligence, which means finding new threats that could affect specific small businesses before they are attacked, helping to improve the framework's ability to defend proactively.

10. Key Takeaways

- Small and home-based businesses face a structurally distinct cybersecurity challenge shaped by absent security ownership, fragmented multi-platform digital footprints, permanent skills gaps, and the practical inaccessibility of existing tools, a challenge that existing enterprise and consumer-grade security solutions are architecturally incapable of addressing.
- The cybersecurity skills shortage impacts small businesses much more severely than large enterprises, which can absorb recruitment costs and build internal capability. This leaves small businesses with no viable path to security expertise other than a platform that eliminates the dependency on expertise entirely.
- Existing AI-driven solutions do not help small businesses because they make expert tasks faster instead of removing the need for expertise, which is a key problem that affects both general-purpose ML models based on enterprise data and self-service tools meant for DevSecOps users.
- SecureMyStore.ai fills this gap with a six-layer system where automation takes over the need for expertise at every step: finding assets, spotting vulnerabilities, ongoing monitoring, AI analysis, scoring risks, and guiding fixes, with people only needed to check easy-to-understand results and carry out approved fixes.
- Testing against both manual reviews and a typical small business tool shows clear and steady improvements in how well assets are discovered, fewer false positives, better success on the first try for fixing issues, faster detection times, and higher usability scores.
- The architecture is a modern approach to automated security operations that can be used in larger companies as well, providing better automation that helps reduce analyst burnout, too many alerts, and reliance on experts, which are common issues in big business security operations.

References

[1] Muhammad Sami Ullah, et al., "AI-Enabled Cybersecurity for Small and Medium-Sized Enterprises (SMEs): A Systematic Review and Evidence-Informed Assessment Framework," Journal of Computing & Biomedical Informatics, 2025. Available: https://www.jcbi.org/index.php/Main/article/view/1212/805

[2] Neeshe Khan, et al., "The hidden barriers to cyber security adoption amongst Small and Medium-Sized Enterprises," Kent Academic Repository, 2025. Available: https://kar.kent.ac.uk/110491/1/ICS2025-Khan-et-al.pdf

[3] Jessica Díaz, et al., "Self-Service Cybersecurity Monitoring as Enabler for DevSecOps," IEEE Xplore, 2019. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8766805

[4] Sana Nawaz, Javid Iqbal, "Leveraging Cloud Security Automation for Scalable and Resilient Infrastructure in SMEs," ResearchGate, 2025. Available: https://www.researchgate.net/profile/Javid-Iqbal-14/publication/393842163_Leveraging_Cloud_Security_Automation_for_Scalable_and_Resilient_Infrastructure_in_SMEs/links/687c9613b3294610e9b8a2ac/Leveraging-Cloud-Security-Automation-for-Scalable-and-Resilient-Infrastructure-in-SMEs.pdf

[5] Ismayil Hasanov, et al., "Application of Large Language Models in Cybersecurity: A Systematic Literature Review," IEEE Xplore, 2024. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10767242

[6] Tarek Ali and Panos Kostakos, "HuntGPT: Integrating Machine Learning-Based Anomaly Detection and Explainable AI with Large Language Models (LLMs)," arXiv, 2023. Available: https://arxiv.org/pdf/2309.16021

[7] Martin Husak and Luka´s Sadlek, "Attack Surface Management: State of the Art and Operational Challenges," Masaryk University, 2025. Available: https://is.muni.cz/publication/2498322/2025-SecSoft-Attack_Surface_Management-paper.pdf

[8] Marco Carmine Rossi, "Enhancing cyber assets visibility for effective attack surface management," University of Turku, 2023. Available: https://www.utupub.fi/bitstream/handle/10024/175930/Thesis-CAASM-CloudSecurity_Marco_Carmine_Rossi.pdf?sequence=1

[9] Halima Ibrahim Kure, et al., "Asset criticality and risk prediction for an effective cybersecurity risk management of cyber-physical system," Springer, 2011. Available: https://link.springer.com/article/10.1007/s00521-021-06400-0

[10] Raghavendra Rao Althar, et al., "Automated Risk Management Based Software Security Vulnerabilities Management," IEEE Xplore, 2022. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9802103

[11] Daniel Agrinya, et al., "Reducing Cloud Misconfiguration Breaches Through Automated Policy Enforcement in AWS and Azure Hybrid Environments," International Journal of Computer Applications Technology and Research, 2024. Available: https://www.researchgate.net/publication/400581029_Reducing_Cloud_Misconfiguration_Breaches_Through_Automated_Policy_Enforcement_in_AWS_and_Azure_Hybrid_Environments

[12] AbuFaizur Rahman Abusalih Rahumath Ali and P. Deepalakshmi, "Cloud Security Posture Management: Automating Risk Detection, Compliance Enforcement, and Vulnerability Remediation in Cloud Infrastructure," IEEE Xplore, 2025. Available: https://ieeexplore.ieee.org/abstract/document/11073903

[13] Beatriz Severes, et al., "The Human Side of XAI: Bridging the Gap between AI and Non-expert Audiences," ACM Digital Library, 2023. Available: https://dl.acm.org/doi/epdf/10.1145/3615335.3623062

[14] SUBASH NEUPANE, et al., "Explainable Intrusion Detection Systems (X-IDS): A Survey of Current Methods, Challenges, and Opportunities," IEEE Xplore, 2022. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9927396

[15] Miriam Curtin, et al., "Enhancing Cybersecurity Awareness in Small and Medium Enterprises Through a User-Friendly Risk Assessment Tool," ResearchGate, 2025. Available: https://www.researchgate.net/publication/394994694_Enhancing_Cybersecurity_Awareness_in_Small_and_Medium_Enterprises_Through_a_User-Friendly_Risk_Assessment_Tool

[16] Edvin Hallvaxhiu, et al., "Accelerate threat modeling with generative AI," AWS Blog, 2025. Available: https://aws.amazon.com/blogs/machine-learning/accelerate-threat-modeling-with-generative-ai/

[17] OWASP Foundation, "OWASP API Security Top 10," OWASP Foundation, 2023. Available:. https://owasp.org/API-Security/

[18] Dan Barahona, "Guide to Automated API Security Testing, " APISec, 2021. Available: https://www.apisec.ai/blog/apisec-the-only-platform-for-automated-api-security-testing