WHITE PAPER

ARGUSS

Defensible Autonomous Remediation for Software Supply Chain Security

---

A framework for context-aware vulnerability prioritization and confidence-based automated remediation of open-source software dependencies.

Arguss Research Team

2026

Table of Contents

Abstract 8

Executive Summary 8

1. Introduction 10

1.1 Software Supply Chain Security Landscape 10

1.2 Problem Statement 10

1.3 Research Contribution 11

2. Background and Related Work 11

2.1 Software Composition Analysis and Vulnerability Management 11

2.2 Existing Approaches to Dependency Security 12

2.3 Limitations of Current Vulnerability Remediation Approaches 13

2.3.1 Vulnerability Volume and Alert Fatigue 13

2.3.2 Lack of Dependency Context 14

2.3.3 Trust and Reliability of Open-Source Packages 14

2.4 Need for Defensible Autonomous Remediation 14

3. Threat Model 15

3.1 Threat Model Overview 15

3.2 Threat Actors 15

3.3 Security Risks Addressed by Arguss 16

4. Arguss System Overview 17

4.1 Design Principles 17

4.1.1 Risk-Aware Automation 17

4.1.2 Explainability 17

4.1.3 Human-in-the-Loop Control 18

4.1.4 Multi-Signal Decision Making 18

4.1.5 Developer-Centric Workflow Integration 19

4.2 System Architecture 19

4.3 Arguss Workflow 19

Stage 1: Repository Scanning 20

Stage 2: Dependency Mapping and SBOM Generation 20

Stage 3: Security Intelligence Collection 20

Stage 4: Three-Lens Risk Analysis 21

Stage 5: Fix Confidence Evaluation 21

Stage 6: Remediation and Reporting 22

4.4 System Differentiation 22
5. Multi-Lens Risk Assessment Framework 22
5.1 Overview 22
5.2 Lens 1: Vulnerability Analysis 23
5.2.1 Purpose 23
5.2.2 CVE and CVSS Analysis 23
5.2.3 Exploit Prediction Scoring System (EPSS) 23
5.2.4 Known Exploited Vulnerabilities (KEV) 24
5.2.5 Vulnerability Risk Output 24
5.3 Lens 2: Package Trust Analysis 24
5.3.1 Purpose 24
5.3.2 OpenSSF Scorecard Integration 24
5.3.3 Repository Health Indicators 25
5.3.4 Package Trust Output 25
5.4 Lens 3: Dependency Context Analysis 25
5.4.1 Purpose 25
5.4.2 Software Bill of Materials (SBOM) 25
5.4.3 Direct vs. Transitive Dependencies 26
5.4.4 Upgrade Impact Assessment 26
5.5 Combined Risk Assessment Model 26
5.6 Advantages of the Three-Lens Approach 27
6. Fix Confidence Engine 27
6.1 Overview 27
6.2 Motivation for Confidence-Based Automation 28
6.2.1 Limitations of Fully Automated Remediation 28
6.2.2 Limitations of Fully Manual Remediation 28
6.2.3 Arguss Approach 28
6.3 Fix Confidence Inputs 29
6.3.1 Vulnerability Risk Signals 29
6.3.2 Upgrade Complexity Signals 29
6.3.3 Package Trust Signals 29
6.3.4 Validation Signals 29
6.4 Confidence Classification Model 29
6.4.1 High Confidence Remediation 30
6.4.2 Medium Confidence Remediation 30

6.4.3 Low Confidence Remediation 30
6.5 Fix Confidence Decision Workflow 30
6.6 Explainability and Decision Transparency 31
6.7 Advantages of the Fix Confidence Engine 31
6.8 Role of Human Oversight 31
6.9 Summary 32
7. System Implementation 32
7.1 Overview 32
7.2 System Architecture Components 32
7.2.1 Repository Scanner 32
7.2.2 Dependency Analysis Engine 33
7.3 Software Bill of Materials Generation 33
7.3.1 Purpose of SBOM Integration 33
7.3.2 SBOM Data Model 34
7.4 Security Intelligence Integration 34
7.4.1 Vulnerability Intelligence Sources 34
7.4.2 Package Trust Intelligence 34
7.5 Backend Architecture 35
7.5.1 Backend Framework 35
7.5.2 Data Processing Pipeline 35
7.6 Frontend Dashboard 35
7.6.1 Purpose 35
7.6.2 Dashboard Capabilities 35
7.7 AI-Assisted Security Analysis 36
7.7.1 Role of AI Integration 36
7.7.2 Human-Centered AI Design 36
7.8 GitHub Workflow Integration 36
7.8.1 Development Workflow Integration 36
7.8.2 Automated Remediation Pull Requests 36
7.9 Technology Stack 36
7.10 Deployment Architecture 37
7.11 Implementation Challenges 37
7.12 Summary 37
8. Evaluation Framework 38
8.1 Evaluation Overview 38

8.2 Research Questions 38
8.3 Evaluation Metrics 39
8.3.1 Vulnerability Detection Metrics 39
8.3.2 Remediation Recommendation Metrics 39
8.4 Evaluation Dataset and Test Scenarios 40
8.4.1 Vulnerable Repository Testing 40
8.4.2 Remediation Scenario Testing 40
8.5 Comparative Evaluation 41
8.6 Developer Efficiency Evaluation 41
8.7 Security Validation Approach 42
8.8 Evaluation Limitations 42
8.9 Summary 42
9. Security and Ethical Considerations 43
9.1 Overview 43
9.2 Human-in-the-Loop Security 43
9.2.1 Importance of Human Oversight 43
9.2.2 Confidence-Based Human Involvement 43
9.3 Explainability and Transparency 44
9.3.1 Importance of Explainable Security Decisions 44
9.3.2 Explainable Remediation Example 44
9.4 Responsible Use of Artificial Intelligence 44
9.4.1 Role of AI in Arguss 44
9.4.2 Preventing AI Over-Reliance 44
9.5 Security of the Arguss Platform 45
9.5.1 Protecting Repository Information 45
9.5.2 Credential Management 45
9.6 Privacy Considerations 45
9.6.1 Minimizing Data Collection 45
9.6.2 Transparency with Users 45
9.7 Equity and Accessibility Considerations 46
9.7.1 Making Security Automation Available 46
9.7.2 Avoiding Unequal Security Outcomes 46
9.8 Ethical Alignment with the E3 Framework 46
9.9 Potential Risks and Mitigations 46
9.10 Summary 47

10. Limitations 47
10.1 Overview 47
10.2 Dependency Ecosystem Coverage 47
10.2.1 Current Coverage Limitations 47
10.2.2 Future Expansion Needs 48
10.3 Dependency Metadata Limitations 48
10.3.1 Reliance on Available Security Information 48
10.3.2 Incomplete Package Context 48
10.4 Limited Runtime Exploitability Analysis 48
10.4.1 Static Analysis Focus 48
10.4.2 Need for Runtime Context 49
10.5 Dependence on CI/CD Validation 49
10.5.1 Importance of Automated Testing 49
10.5.2 Impact on Automation Confidence 49
10.6 AI-Generated Explanation Limitations 49
10.6.1 AI Reasoning Boundaries 49
10.6.2 Mitigation Approach 49
10.7 Limited Historical Remediation Data 50
10.7.1 Challenge 50
10.7.2 Future Improvement 50
10.8 Human Decision Dependency 50
10.8.1 Remaining Need for Expertise 50
10.8.2 Balancing Automation and Control 50
10.9 Scalability Considerations 50
10.9.1 Large-Scale Repository Management 50
10.9.2 Enterprise Integration 50
10.10 Summary 51
11. Future Work 51
11.1 Overview 51
11.2 Expanded Programming Language and Ecosystem Support 51
11.2.1 Broader Dependency Analysis 51
11.2.2 Cross-Ecosystem Risk Analysis 52
11.3 Runtime-Aware Security Analysis 52
11.3.1 Moving Beyond Static Dependency Analysis 52
11.3.2 Reachability Analysis 52

11.4 Advanced Fix Confidence Modeling 52
11.4.1 Machine Learning-Based Confidence Prediction 52
11.4.2 Organization-Specific Learning 53
11.5 Enhanced AI Security Assistant 53
11.5.1 More Advanced Security Reasoning 53
11.5.2 Security Policy-Aware Recommendations 53
11.6 Continuous Software Supply Chain Monitoring 53
11.6.1 From Periodic Scanning to Continuous Protection 53
11.6.2 Real-Time Risk Updates 53
11.7 Enterprise Security Integration 53
11.7.1 Security Operations Integration 53
11.7.2 Identity and Access Integration 54
11.8 Improved Validation and Testing Capabilities 54
11.8.1 Automated Remediation Verification 54
11.8.2 Pre-Merge Security Validation 54
11.9 Community-Driven Security Intelligence 54
11.9.1 Shared Remediation Knowledge 54
11.9.2 Balancing Collaboration and Privacy 54
11.10 Long-Term Vision 54
11.11 Summary 55
12. Conclusion 55
12.1 Overview 55
12.2 Key Contributions 55
12.2.1 Moving Beyond Vulnerability Detection 55
12.2.2 Multi-Lens Security Analysis 55
12.2.3 Confidence-Based Autonomous Remediation 56
12.3 Impact on Cybersecurity Operations 56
12.3.1 Reducing Security Team Workload 56
12.3.2 Improving Developer Security Experience 56
12.3.3 Strengthening Software Supply Chain Security 56
12.4 Responsible Automation as a Security Principle 56
12.5 Broader Implications for AI in Cybersecurity 57
12.6 Final Perspective 57

Abstract

Modern software development increasingly depends on open-source software components, creating a complex and rapidly expanding software supply chain. While open-source dependencies accelerate innovation and reduce development costs, they also introduce significant security risks. Vulnerable, outdated, or compromised dependencies can provide attackers with opportunities to infiltrate applications, escalate privileges, and compromise downstream organizations.

Existing Software Composition Analysis (SCA) tools have improved vulnerability visibility by identifying known vulnerabilities in third-party dependencies. However, these tools often produce large volumes of findings without providing sufficient context about exploitability, package trustworthiness, or remediation safety. As a result, developers and security teams frequently experience alert fatigue, spending significant time manually investigating vulnerabilities rather than resolving the most impactful risks.

Arguss is a defensible autonomous remediation platform designed to address this challenge. Instead of functioning solely as a vulnerability scanner, Arguss evaluates whether a vulnerability requires action and whether a proposed remediation can be safely automated. The platform combines vulnerability intelligence, package trust analysis, dependency context, Software Bill of Materials (SBOM) generation, and artificial intelligence-assisted reasoning to provide explainable remediation recommendations.

The primary contribution of Arguss is the Fix Confidence Engine, which determines the appropriate level of automation for each remediation opportunity. High-confidence fixes can be automatically proposed, medium-confidence fixes require developer review, and low-confidence cases are escalated for manual investigation. This human-centered approach enables organizations to benefit from automation while maintaining accountability and operational reliability.

This paper presents the motivation, architecture, implementation, evaluation framework, and ethical considerations behind Arguss as a next-generation approach to software supply chain security.

Executive Summary

Software supply chain security has become one of the most important challenges facing modern cybersecurity teams. Applications today are rarely developed entirely from internally written code. Instead, they depend on thousands of external libraries, packages, and frameworks maintained by global open-source communities.

This dependency ecosystem introduces significant risk. A vulnerability discovered in a commonly used package can affect thousands of organizations simultaneously. Recent software supply chain incidents have demonstrated that attackers increasingly target dependencies because compromising a trusted component can provide access to many downstream applications.

Organizations currently rely on vulnerability scanners and Software Composition Analysis tools to identify vulnerable dependencies. While these tools provide valuable visibility, they primarily answer one question:

> *"Which dependencies have known vulnerabilities?"*

However, developers need answers to additional questions:

- How dangerous is this vulnerability in my environment?
- Is the package itself trustworthy?
- Is the dependency actually being used?
- Will upgrading break my application?
- Can this fix be safely automated?

Without these answers, security teams face thousands of findings with limited ability to prioritize. This creates remediation delays and increases organizational exposure.

Arguss addresses this problem by shifting from vulnerability detection toward defensible autonomous remediation.

The platform combines several security intelligence sources:

- Vulnerability severity information
- Exploit likelihood
- Known exploited vulnerabilities
- Package reputation and maintenance signals
- Dependency relationships
- Software Bill of Materials analysis
- Repository security indicators

These signals are evaluated through Arguss's Three-Lens Risk Analysis Framework:

1. Vulnerability Analysis Lens — Determines the technical severity and exploitability of vulnerabilities.
2. Package Trust Lens — Evaluates whether a dependency is maintained, trustworthy, and actively supported.
3. Dependency Context Lens — Determines how the vulnerable dependency affects the application environment.

The results are processed by the Fix Confidence Engine, which determines whether remediation should be automated, reviewed, or deferred.

Unlike traditional automated update systems, Arguss does not assume that every available patch is safe. Instead, it follows a principle of responsible automation:

> *Automate when confidence is high. Escalate when uncertainty remains.*

This approach reduces unnecessary developer workload while preserving security, transparency, and human control.

1. Introduction

1.1 Software Supply Chain Security Landscape

Modern software development relies heavily on open-source ecosystems. Package managers such as npm, PyPI, Maven, and others allow developers to rapidly integrate external functionality into applications. This approach has significantly accelerated software innovation, but it has also expanded the attack surface organizations must manage.

A modern application may contain hundreds or thousands of dependencies, including direct dependencies explicitly selected by developers and transitive dependencies introduced indirectly through other packages.

While dependency reuse improves efficiency, it creates several security challenges:

- Vulnerabilities may exist in outdated packages.
- Organizations may unknowingly rely on abandoned libraries.
- Attackers may introduce malicious packages into public repositories.
- Security teams may lack visibility into complete dependency chains.
- Developers may delay updates due to concerns about breaking functionality.

The software supply chain has therefore become a critical cybersecurity concern. Protecting applications requires not only identifying vulnerable components but also making intelligent decisions about remediation.

1.2 Problem Statement

Current vulnerability management approaches focus primarily on discovery. Security tools generate findings, assign severity scores, and recommend available updates. However, remediation decisions often remain manual.

This creates three major challenges:

Alert Overload

Large organizations may receive thousands of vulnerability alerts. Many are low-risk, difficult to exploit, or unrelated to the application's actual usage.

Lack of Context

Severity scores alone do not explain whether a vulnerability represents a meaningful business risk.

A critical vulnerability in an unused development dependency may require less urgency than a moderate vulnerability affecting an internet-facing production service.

Unsafe Automation

Automatically applying updates can introduce operational failures, compatibility issues, or unexpected application behavior.

As a result, organizations face a difficult tradeoff:

- Manual review is safer but slow.
- Full automation is faster but risky.

Arguss addresses this gap by introducing confidence-based autonomous remediation.

1.3 Research Contribution

Arguss contributes a new approach to software supply chain security through three primary innovations:

1. Context-Aware Vulnerability Prioritization

Arguss combines multiple security signals rather than relying on vulnerability severity alone.

2. Confidence-Based Remediation

The Fix Confidence Engine determines whether automated action is appropriate.

3. Defensible Automation

Every recommendation is accompanied by reasoning, allowing developers and security teams to understand why an action was suggested.

Together, these capabilities enable safer automation while maintaining human oversight.

2. Background and Related Work

2.1 Software Composition Analysis and Vulnerability Management

Software Composition Analysis (SCA) has become a foundational capability in modern application security programs. SCA tools analyze application dependencies and compare identified components against vulnerability databases to determine whether known security issues exist.

Traditional SCA solutions typically perform three primary functions:

4. Dependency Identification — Detecting third-party libraries and package versions used within an application.
5. Vulnerability Mapping — Matching dependencies against vulnerability intelligence sources such as Common Vulnerabilities and Exposures (CVE) databases.
6. Risk Reporting — Generating findings that describe affected packages, severity ratings, and recommended upgrades.

These capabilities provide organizations with essential visibility into their software supply chain. However, the rapid growth of open-source dependencies has created a scalability challenge. Modern applications may contain thousands of packages, producing a continuous stream of security findings.

The primary limitation of traditional SCA tools is that they typically stop at detection. They identify vulnerable dependencies but provide limited assistance in determining:

- Which vulnerabilities should be addressed first.
- Whether a vulnerability is exploitable in the current environment.
- Whether the dependency is trustworthy.
- Whether a recommended upgrade is safe.
- Whether remediation can be automated.

As a result, security teams must manually investigate large numbers of alerts before taking action.

2.2 Existing Approaches to Dependency Security

Several existing approaches attempt to improve software dependency security. While each provides valuable capabilities, significant gaps remain.

Dependency Update Automation

Tools such as automated dependency update systems create pull requests when newer package versions become available. These systems reduce manual maintenance effort by automatically identifying available updates.

However, version availability does not necessarily indicate security relevance or operational safety.

Potential challenges include:

- Updates may introduce breaking changes.
- Security improvements may not address the actual risk.
- Developers may ignore automated pull requests due to volume.
- No contextual risk analysis may be performed.

Automation without sufficient reasoning can create additional workload rather than reducing it.

Vulnerability Scanning Platforms

Modern vulnerability management platforms provide deeper security visibility by aggregating vulnerability intelligence from multiple sources.

These platforms commonly incorporate:

- CVSS severity scoring.
- Vulnerability age.

- Exploit availability.
- Asset importance.

However, many systems still depend heavily on human analysts to determine remediation priority. A vulnerability score alone does not provide enough information to decide whether a fix should be automatically applied.

For example, two applications may contain the same vulnerable dependency:

- Application A uses the dependency in an isolated internal testing environment.
- Application B exposes the dependency through a public-facing production service.

Although the technical vulnerability is identical, the practical risk is significantly different.

Container and Infrastructure Scanning

Tools designed for container and infrastructure security have expanded vulnerability visibility beyond application dependencies. These solutions analyze:

- Container images.
- Operating system packages.
- Cloud configurations.
- Infrastructure-as-code templates.

While valuable, they primarily focus on identifying insecure conditions rather than determining remediation confidence.

Arguss extends these concepts by focusing specifically on the decision-making process required before automated dependency remediation.

2.3 Limitations of Current Vulnerability Remediation Approaches

The current software security ecosystem faces several unresolved challenges.

2.3.1 Vulnerability Volume and Alert Fatigue

The number of disclosed vulnerabilities continues to increase each year. Security teams often receive more findings than they can manually investigate.

This creates several problems:

- Critical issues may be overlooked among lower-priority findings.
- Developers may become desensitized to security notifications.
- Security teams spend significant time reviewing alerts instead of reducing risk.

The challenge is no longer simply discovering vulnerabilities. The challenge is determining which vulnerabilities require immediate action.

2.3.2 Lack of Dependency Context

Many vulnerability management systems evaluate dependencies individually rather than considering their application context.

Important questions often remain unanswered:

- Is the vulnerable code actually reachable?
- Is the package directly or indirectly imported?
- Is the affected functionality enabled?
- Does the dependency have active maintainers?
- Are safer upgrade paths available?

Without context, organizations may spend resources fixing vulnerabilities that represent minimal risk while delaying more important issues.

2.3.3 Trust and Reliability of Open-Source Packages

Security decisions increasingly require understanding the health and reliability of software projects themselves.

A package with no recent maintenance, unknown ownership, or suspicious development activity may represent a greater long-term risk than a package with an active and trustworthy community.

Important package trust indicators include:

- Repository activity.
- Release frequency.
- Contributor diversity.
- Security practices.
- Maintenance history.

Traditional vulnerability scanners often do not incorporate these signals into remediation decisions.

2.4 Need for Defensible Autonomous Remediation

The limitations of current approaches demonstrate the need for a new security model.

Organizations require automation, but automation must be:

- Context-aware — understanding the environment where vulnerabilities exist.
- Explainable — showing why a recommendation was made.

- Controlled — avoiding unsafe automatic changes.
- Adaptive — increasing automation when confidence improves.

Arguss introduces the concept of defensible autonomous remediation, where automation is not based solely on the existence of a vulnerability or available patch.

Instead, remediation decisions are based on evidence. The system asks:

7. How severe is the vulnerability?
8. How likely is exploitation?
9. How trustworthy is the affected dependency?
10. What impact could an upgrade introduce?
11. How confident are we that remediation is safe?

By answering these questions, Arguss transforms vulnerability management from a detection-focused process into a decision-support system.

3. Threat Model

3.1 Threat Model Overview

The threat model defines the security risks Arguss is designed to address. The platform focuses primarily on software supply chain threats where attackers exploit weaknesses in dependencies, package ecosystems, or automated development workflows.

Arguss assumes that attackers may attempt to compromise software through:

- Vulnerable third-party libraries.
- Malicious dependency releases.
- Abandoned packages.
- Dependency confusion attacks.
- Compromised developer workflows.

The objective of Arguss is not only to identify these risks but also to help organizations respond with secure and appropriate remediation actions.

3.2 Threat Actors

Arguss considers several categories of adversaries.

External Attackers

External attackers may exploit known vulnerabilities in publicly available dependencies to compromise applications.

Common objectives include:

- Unauthorized access.
- Data theft.
- Remote code execution.
- Privilege escalation.

Malicious Package Publishers

Attackers may publish malicious packages or compromise existing packages within public repositories.

Potential techniques include:

- Typosquatting.
- Package takeover.
- Malicious updates.
- Hidden backdoors.

Supply Chain Compromise Actors

Advanced attackers may target trusted software components because compromising a widely used dependency can impact many organizations simultaneously.

These attacks highlight the importance of evaluating not only vulnerabilities but also package trustworthiness.

3.3 Security Risks Addressed by Arguss

Arguss focuses on reducing several key software supply chain risks:

Threat	Arguss Response
Vulnerable dependencies	Vulnerability analysis and prioritization
Exploitable packages	EPSS and KEV integration
Unsafe updates	Fix Confidence Engine
Untrusted packages	Package trust analysis
Dependency visibility gaps	SBOM generation

Through these capabilities, Arguss provides organizations with a more complete understanding of software dependency risk.

4. Arguss System Overview

4.1 Design Principles

Arguss was designed around the principle that cybersecurity automation should be defensible, explainable, and risk-aware. Traditional automation systems often prioritize speed by applying changes whenever a vulnerability or update is detected. While this approach can reduce manual effort, it introduces operational risks when automated decisions are made without sufficient context.

Arguss takes a different approach by treating automation as a decision-making problem. Instead of asking:

> *"Can this vulnerability be automatically fixed?"*

Arguss asks:

> *"Do we have enough evidence to safely automate this remediation?"*

This philosophy guides the platform's architecture and implementation. The system follows five core design principles:

4.1.1 Risk-Aware Automation

Arguss does not automatically remediate every vulnerability. The platform evaluates multiple risk factors before recommending action.

A vulnerability's remediation priority is determined by combining:

- Technical severity
- Exploit likelihood
- Real-world exploitation status
- Dependency importance
- Package reliability
- Upgrade impact

This prevents organizations from spending time on low-impact issues while ensuring critical risks receive appropriate attention.

4.1.2 Explainability

Security automation requires trust. Developers and security teams must understand why a system recommends a specific action.

Arguss provides explanations for remediation decisions by identifying:

- Which security signals influenced the recommendation.
- Why a dependency was prioritized.

- Why automation was allowed or restricted.
- What risks remain after remediation.

Rather than producing a simple instruction such as:

> *"Upgrade package X."*

Arguss provides reasoning such as:

> *"Package X contains a known exploited vulnerability with high exploit probability. The recommended upgrade is a minor version change, maintains compatibility, and the package demonstrates strong maintenance activity. Automated remediation confidence is high."*

This transparency allows users to validate decisions rather than blindly trust automation.

4.1.3 Human-in-the-Loop Control

Although automation can improve security operations, fully autonomous decision-making introduces risks.

Arguss follows a human-in-the-loop security model:

- High-confidence actions may proceed automatically.
- Medium-confidence recommendations require developer approval.
- Low-confidence scenarios require manual investigation.

This approach balances efficiency with accountability.

The goal of Arguss is not to replace security professionals or developers. Instead, it reduces repetitive analysis while allowing humans to focus on complex decisions.

4.1.4 Multi-Signal Decision Making

Security decisions should not rely on a single metric.

For example, CVSS severity alone does not determine whether a vulnerability represents immediate risk. A critical vulnerability in an unused package may be less urgent than a moderate vulnerability actively exposed in production.

Arguss combines multiple sources of intelligence:

- CVSS severity scoring
- EPSS exploitation probability
- CISA Known Exploited Vulnerabilities (KEV)
- OpenSSF Scorecard metrics
- Dependency relationships
- SBOM information

- Package maintenance signals

This multi-signal approach provides a more accurate representation of software risk.

4.1.5 Developer-Centric Workflow Integration

Security tools often fail when they disrupt developer workflows.

Arguss integrates security analysis into existing development processes through:

- Repository scanning
- GitHub workflow integration
- Pull request recommendations
- Dashboard visualization
- AI-generated security summaries

By meeting developers where they already work, Arguss reduces friction between security and engineering teams.

4.2 System Architecture

Arguss consists of several interconnected components that work together to analyze dependencies, evaluate risk, and recommend remediation actions.

The architecture follows a pipeline-based design:

Figure 1. Arguss High-Level Architecture

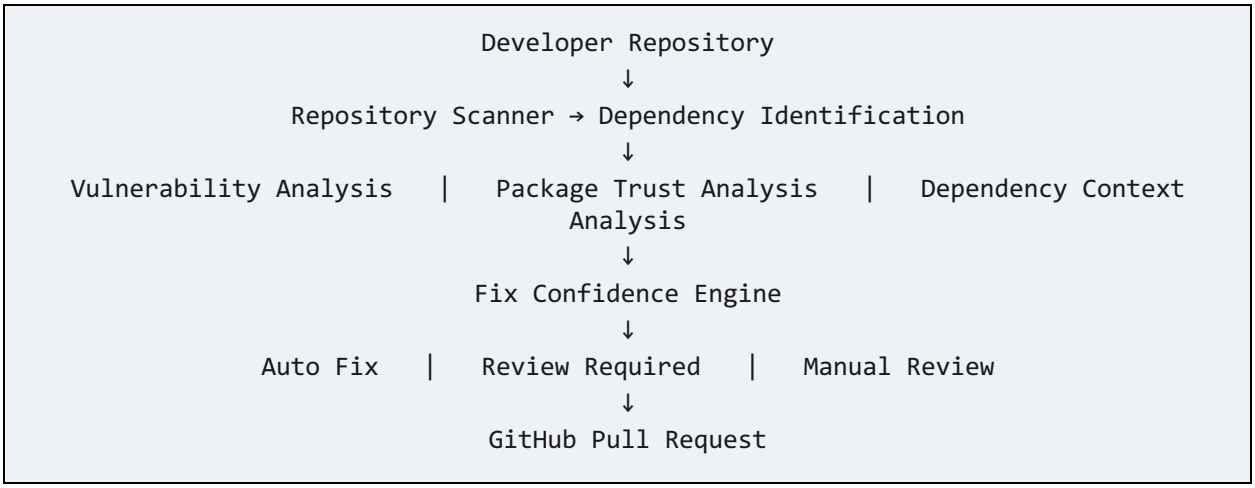

4.3 Arguss Workflow

The Arguss workflow consists of six major stages.

Stage 1: Repository Scanning

The process begins when a developer connects a repository to Arguss.

The scanner identifies:

- Programming language ecosystem.
- Dependency files.
- Package versions.
- Direct and indirect dependencies.

Examples of supported dependency files include:

- package.json
- requirements.txt
- pom.xml
- go.mod

The goal of this stage is to establish a complete understanding of the application dependency environment.

Stage 2: Dependency Mapping and SBOM Generation

After identifying dependencies, Arguss generates a Software Bill of Materials (SBOM).

An SBOM provides a structured inventory of software components, including:

- Package name.
- Package version.
- Dependency relationships.
- Component origin.
- Known security information.

SBOM generation improves visibility by answering:

> *"What software components exist inside this application?"*

This visibility is critical for managing software supply chain risks.

Stage 3: Security Intelligence Collection

Arguss collects information from multiple security sources. These sources provide different perspectives on dependency risk:

Vulnerability Intelligence

Includes: CVE information, CVSS severity, exploit availability, EPSS scores, and KEV status.

Package Intelligence

Includes: repository activity, maintenance history, community adoption, and security practices.

Dependency Intelligence

Includes: dependency depth, application usage, upgrade paths, and potential impact.

Combining these signals creates a more complete security assessment.

Stage 4: Three-Lens Risk Analysis

The collected information is processed through Arguss's Three-Lens Risk Analysis Framework. The framework evaluates each dependency from three perspectives:

Lens 1: Vulnerability Risk

Determines: how severe is the vulnerability, is exploitation occurring, and how likely is exploitation.

Lens 2: Package Trust

Determines: is this dependency actively maintained, does the project demonstrate healthy development practices, and are there signs of supply chain risk.

Lens 3: Dependency Context

Determines: how important is this dependency, is it directly used by the application, and what impact would upgrading introduce.

Stage 5: Fix Confidence Evaluation

After analyzing risk, Arguss evaluates whether remediation is safe. The Fix Confidence Engine considers:

- Version change size.
- Compatibility impact.
- Dependency conflicts.
- Historical stability.
- Security improvement.
- Testing results.

The engine produces a confidence classification:

Confidence Score	Recommendation
High Confidence	Automatically create remediation
Medium Confidence	Generate pull request for review
Low Confidence	Require manual investigation

Stage 6: Remediation and Reporting

The final stage provides actionable results. Depending on confidence, Arguss may:

- Automatically generate a remediation pull request.
- Provide upgrade recommendations.
- Generate an executive security summary.
- Explain remediation reasoning.

The final output is designed for both technical and non-technical audiences. Developers receive actionable fixes, while security leaders receive risk-focused summaries.

4.4 System Differentiation

Arguss differs from traditional vulnerability management tools by focusing on the transition between finding vulnerabilities and securely fixing vulnerabilities.

Traditional tools answer:

> *"What is vulnerable?"*

Arguss answers:

> *"What should we fix, why should we fix it, and how confident are we that the fix is safe?"*

This distinction allows Arguss to move vulnerability management from reactive detection toward proactive, intelligent remediation.

5. Multi-Lens Risk Assessment Framework

5.1 Overview

A central challenge in software supply chain security is that vulnerability severity alone does not provide enough information to determine appropriate remediation actions. A vulnerability may have a high severity score but represent limited practical risk depending on application context, dependency usage, or exploit availability.

Arguss addresses this challenge through a Three-Lens Risk Assessment Framework. Instead of evaluating dependencies using a single vulnerability score, Arguss analyzes each dependency through three complementary perspectives:

12. Vulnerability Analysis Lens — Evaluates technical security risk based on vulnerability intelligence.
13. Package Trust Lens — Evaluates the reliability and security posture of the dependency itself.
14. Dependency Context Lens — Evaluates how the dependency is used within the application environment.

Together, these three lenses provide a comprehensive assessment of both security risk and remediation confidence.

5.2 Lens 1: Vulnerability Analysis

5.2.1 Purpose

The Vulnerability Analysis Lens determines the immediate security impact of a vulnerable dependency.

Traditional vulnerability management systems often rely heavily on CVSS scores. While CVSS provides useful information about vulnerability characteristics, it does not fully represent real-world risk.

For example:

- A critical vulnerability may have no known exploitation activity.
- A moderate vulnerability may already be actively exploited.
- A vulnerable package may exist but never be executed by the application.

Therefore, Arguss combines multiple vulnerability intelligence sources to create a more accurate risk assessment.

5.2.2 CVE and CVSS Analysis

Arguss incorporates Common Vulnerabilities and Exposures (CVE) information to identify publicly disclosed vulnerabilities affecting dependencies.

Each vulnerability is evaluated using: vulnerability description, affected package versions, available remediation versions, severity rating, and attack characteristics.

CVSS scores provide insight into exploit complexity, required privileges, user interaction requirements, and potential impact. However, CVSS is treated as one input among many rather than the final decision factor.

5.2.3 Exploit Prediction Scoring System (EPSS)

Arguss incorporates EPSS to estimate the probability that a vulnerability will be exploited in the near future. This improves prioritization because vulnerabilities with active exploitation potential require faster response.

For example:

Vulnerability	CVSS	EPSS	Priority
Vulnerability A	Critical	Low exploitation probability	Medium
Vulnerability B	High	High exploitation probability	High

By combining severity and exploit likelihood, Arguss avoids prioritizing vulnerabilities based only on theoretical impact.

5.2.4 Known Exploited Vulnerabilities (KEV)

Arguss integrates the CISA Known Exploited Vulnerabilities catalog to identify vulnerabilities that have been observed in active attacks.

KEV status increases remediation priority because it provides evidence that attackers are already targeting the vulnerability. A vulnerability appearing in the KEV catalog receives additional risk weighting because:

- Exploitation has been confirmed.
- Attack techniques may already exist publicly.
- Delayed remediation increases exposure.

5.2.5 Vulnerability Risk Output

The Vulnerability Analysis Lens produces a security risk assessment containing: vulnerability severity, exploitation likelihood, active exploitation status, and recommended urgency.

Example output:

> *Package: example-library — Vulnerability: CVE-XXXX-XXXX — Severity: High — EPSS: Elevated exploitation probability — KEV Status: Active exploitation observed — Recommendation: Prioritize remediation*

5.3 Lens 2: Package Trust Analysis

5.3.1 Purpose

Security risk is not limited to known vulnerabilities. The security posture of the dependency itself is equally important.

A package may not currently contain a known vulnerability but still represent a future supply chain risk if it is abandoned, poorly maintained, controlled by unknown contributors, or missing security practices.

The Package Trust Lens evaluates the reliability and maturity of dependencies.

5.3.2 OpenSSF Scorecard Integration

Arguss incorporates OpenSSF Scorecard signals to evaluate open-source project health.

Relevant indicators include: repository activity, security practices, dependency management, code review processes, and maintainer activity.

These signals help distinguish between trusted and higher-risk dependencies:

Trusted dependency example	Higher-risk dependency example
Active maintainers.	No recent updates.
Regular releases.	Limited maintainers.
Security-focused practices.	Unknown ownership.
Transparent development.	Lack of security controls.

5.3.3 Repository Health Indicators

Arguss evaluates additional package trust factors:

Maintenance Activity

Questions evaluated: when was the last release, are security fixes regularly published, and is development ongoing.

Community Adoption

Indicators include: number of users, popularity within the ecosystem, and number of contributors. Widely adopted packages are not automatically secure, but adoption can provide additional evidence of community review.

Contributor and Ownership Signals

Arguss considers: number of maintainers, contributor diversity, and repository ownership history. These signals help identify potential risks associated with single-maintainer projects or abandoned packages.

5.3.4 Package Trust Output

The Package Trust Lens produces a trust assessment. Example:

> *Package: example-library — Repository Health: Strong — Maintenance: Active — Security Practices: Moderate — Trust Assessment: Low supply chain concern*

This information helps determine whether remediation should include additional investigation.

5.4 Lens 3: Dependency Context Analysis

5.4.1 Purpose

The same vulnerability can represent different levels of risk depending on how a dependency is used.

The Dependency Context Lens provides application-specific understanding by analyzing dependency relationships, direct versus transitive usage, application importance, and upgrade impact.

5.4.2 Software Bill of Materials (SBOM)

A key component of Arguss is SBOM generation. An SBOM creates an inventory of all software components within an application.

It provides visibility into package names, versions, dependency relationships, and component origins. SBOMs are essential because organizations cannot secure software they cannot fully identify.

5.4.3 Direct vs. Transitive Dependencies

Arguss distinguishes between direct dependencies (packages explicitly installed and managed by developers) and transitive dependencies (packages introduced indirectly through another dependency).

Direct dependency example

```
Application
    ↓
requests-library
```

Transitive dependency example

```
Application
    ↓
 framework
    ↓
security-library
```

Understanding this relationship helps determine remediation ownership. A direct dependency may require a developer update; a transitive dependency may require upgrading the parent package.

5.4.4 Upgrade Impact Assessment

Before recommending automated remediation, Arguss evaluates potential upgrade impact. Factors include version difference, breaking changes, dependency conflicts, and compatibility concerns.

A patch-level update is generally lower risk than a major version upgrade:

Change	Risk
1.2.3 → 1.2.4	Low change risk
1.2.3 → 2.0.0	Higher change risk

This information directly contributes to Fix Confidence Engine decisions.

5.5 Combined Risk Assessment Model

The three lenses are combined to produce a complete dependency risk profile. Conceptually:

> *Overall Risk = Vulnerability Risk + Package Trust Risk + Dependency Context Risk*

However, Arguss does not simply calculate a numerical score. Instead, it uses these signals to determine security importance, remediation urgency, and automation confidence.

Figure 2. Three-Lens Risk Assessment Framework

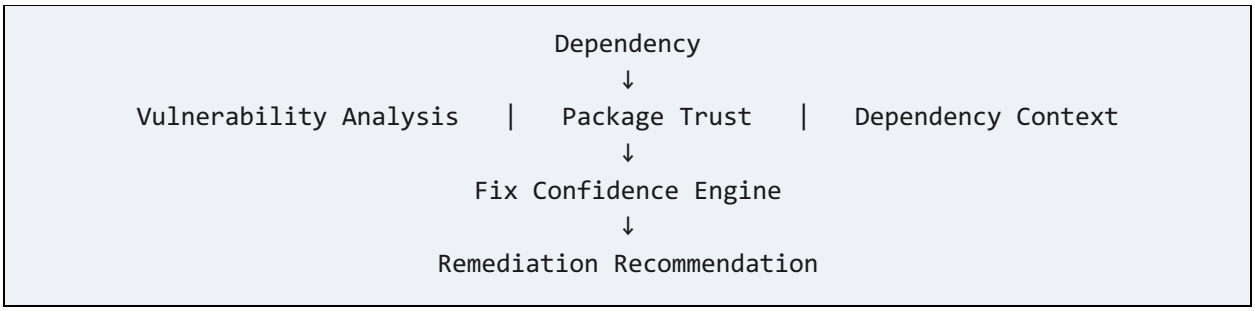

5.6 Advantages of the Three-Lens Approach

The Three-Lens Framework provides several advantages over traditional vulnerability scanning:

Improved Prioritization

Security teams focus on vulnerabilities that represent meaningful risk.

Reduced Alert Fatigue

Low-impact findings receive less attention.

Safer Automation

Remediation decisions consider operational impact.

Better Explainability

Users understand why a dependency was prioritized.

By combining vulnerability intelligence, package trust, and application context, Arguss creates a more complete understanding of software supply chain risk.

6. Fix Confidence Engine

6.1 Overview

The Fix Confidence Engine is the core innovation of Arguss. While traditional vulnerability management tools focus on identifying security issues, Arguss focuses on determining whether a remediation action can be performed safely and responsibly.

Automated remediation introduces a fundamental challenge: a security fix may successfully eliminate a vulnerability while unintentionally creating operational problems. For example, upgrading a dependency

may introduce breaking changes, remove required functionality, or create compatibility issues with other components.

Therefore, Arguss does not treat remediation as a binary decision of “fix” or “do not fix.” Instead, it evaluates the level of confidence that a proposed fix will improve security without negatively impacting application reliability.

The Fix Confidence Engine answers the following question:

> *“Based on available security and operational evidence, how confident are we that this remediation is safe to apply?”*

This enables Arguss to move beyond simple automation toward defensible autonomous remediation.

6.2 Motivation for Confidence-Based Automation

6.2.1 Limitations of Fully Automated Remediation

Many existing dependency automation systems operate under the assumption that newer package versions are preferable. While updates frequently improve security, this assumption does not always hold true.

Potential risks include: breaking API changes, dependency conflicts, unexpected application behavior, failed deployments, and reduced system availability.

For example, a major package upgrade may resolve a vulnerability but require significant code changes from developers. Automatically applying such changes without evaluation could introduce additional operational risk.

6.2.2 Limitations of Fully Manual Remediation

At the opposite extreme, requiring human review for every vulnerability creates scalability challenges.

Security teams often face thousands of vulnerability findings, limited engineering resources, increasing remediation deadlines, and repetitive analysis tasks.

Manual-only workflows slow response times and make it difficult to address large-scale software supply chain risks.

6.2.3 Arguss Approach

Arguss balances these challenges through confidence-based automation. Instead of asking:

> *“Can this update be automated?”*

Arguss evaluates:

> *“Do we have enough evidence to trust this automation?”*

This approach allows organizations to safely automate low-risk decisions while preserving human judgment for uncertain cases.

6.3 Fix Confidence Inputs

The Fix Confidence Engine combines signals from multiple sources.

6.3.1 Vulnerability Risk Signals

The engine evaluates CVSS severity, EPSS exploitation probability, KEV inclusion, vulnerability age, and available security patches. These signals determine how important remediation is from a security perspective.

Example: a high-severity vulnerability with active exploitation receives a higher remediation priority than a vulnerability with no known exploitation.

6.3.2 Upgrade Complexity Signals

The engine evaluates the technical impact of applying a fix. Factors include version change magnitude, breaking change indicators, dependency conflicts, package compatibility, and release stability.

Version changes are categorized based on semantic versioning:

Change Type	Example	Expected Risk
Patch Update	1.4.2 → 1.4.3	Low
Minor Update	1.4.2 → 1.5.0	Medium
Major Update	1.4.2 → 2.0.0	High

A patch-level security update generally receives higher automation confidence than a major version migration.

6.3.3 Package Trust Signals

The engine incorporates package reliability information, including OpenSSF Scorecard results, repository maintenance activity, release frequency, contributor activity, and project maturity.

A well-maintained package with a predictable release history provides greater confidence than an abandoned dependency.

6.3.4 Validation Signals

When available, Arguss considers validation evidence from development workflows, such as existing CI/CD tests, build success, dependency installation success, and automated test results.

Successful validation increases confidence that a remediation will not negatively impact application functionality.

6.4 Confidence Classification Model

The Fix Confidence Engine categorizes remediation opportunities into three levels.

6.4.1 High Confidence Remediation

High-confidence fixes meet several criteria: clear security improvement, low upgrade complexity, strong package reliability, successful validation, and minimal expected application impact.

Recommended action: automatic remediation.

> *Example: a vulnerable package requires a patch-level upgrade from version 3.2.1 to 3.2.2. The update fixes an actively exploited vulnerability, introduces no breaking changes, and passes automated validation.*

Arguss may automatically generate a remediation pull request.

6.4.2 Medium Confidence Remediation

Medium-confidence fixes provide security benefits but require additional review. Common characteristics: minor version upgrade, limited compatibility information, moderate dependency impact, and partial validation.

Recommended action: create pull request with developer approval.

> *Example: a dependency upgrade resolves a high-risk vulnerability but introduces several package changes requiring application owner review.*

6.4.3 Low Confidence Remediation

Low-confidence fixes contain significant uncertainty. Characteristics include major version changes, poor package maintenance, missing validation, and potential breaking changes.

Recommended action: human investigation required.

> *Example: updating the dependency resolves a vulnerability but requires migration changes and may impact application functionality.*

Arguss avoids automatic action in these scenarios.

6.5 Fix Confidence Decision Workflow

Figure 3. Fix Confidence Engine Workflow

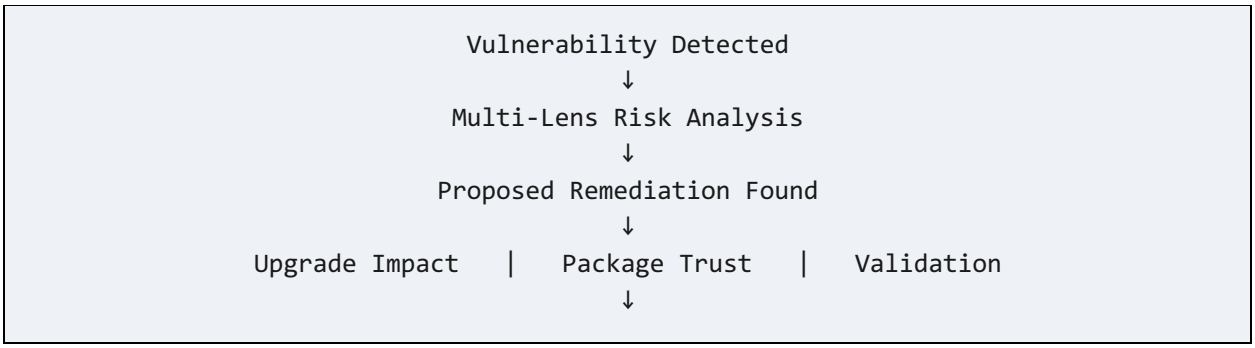

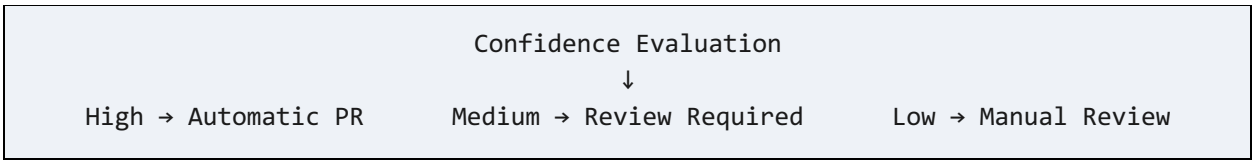

6.6 Explainability and Decision Transparency

A critical requirement for autonomous security systems is the ability to explain decisions. Arguss provides remediation reasoning by identifying the factors contributing to confidence levels.

Example — Recommendation: upgrade package lodash 4.17.20 → 4.17.21.

Security reasoning:

- Vulnerability severity: High
- Exploitation likelihood: Moderate
- KEV status: Not listed
- Package trust: Strong
- Upgrade type: Patch update
- Validation status: Passed

Decision: High Confidence — automatic remediation recommended.

This explanation allows developers and security teams to verify the decision before accepting changes.

6.7 Advantages of the Fix Confidence Engine

Reduced Developer Workload

Developers spend less time investigating low-value alerts.

Faster Response to Critical Threats

High-confidence vulnerabilities can be addressed quickly without waiting for manual review.

Safer Automation

The system avoids blindly applying updates that may introduce instability.

Increased Security Team Trust

Transparent reasoning allows security teams to understand and validate automated decisions.

6.8 Role of Human Oversight

Although Arguss enables autonomous remediation, human oversight remains an essential component.

Security decisions often involve business context that cannot be fully captured through automated analysis, such as regulatory requirements, application criticality, business deadlines, and customer impact.

Therefore, Arguss follows a principle of:

> *Automate confidently, escalate responsibly.*

This ensures that automation improves security operations without removing accountability.

6.9 Summary

The Fix Confidence Engine transforms Arguss from a vulnerability detection tool into a remediation decision platform. By combining vulnerability intelligence, package trust analysis, dependency context, upgrade impact analysis, and validation evidence, Arguss determines not only what should be fixed, but also how safely it can be fixed.

This confidence-driven approach represents the foundation of defensible autonomous remediation.

7. System Implementation

7.1 Overview

Arguss was implemented as an integrated cybersecurity platform that combines automated dependency analysis, security intelligence gathering, risk evaluation, and remediation recommendation. The implementation focuses on creating a practical workflow that fits into existing software development processes while providing advanced security analysis.

The system architecture follows a modular design where each component performs a specific security function:

15. Repository Analysis Layer — Identifies dependencies and application structure.
16. Security Intelligence Layer — Collects vulnerability, exploitability, and package trust information.
17. Risk Analysis Layer — Evaluates dependencies using the Three-Lens Risk Framework.
18. Decision Layer — Uses the Fix Confidence Engine to determine remediation actions.
19. Presentation Layer — Provides developers and security teams with actionable insights through dashboards and reports.

This modular architecture allows Arguss to scale as additional security data sources and analysis capabilities are integrated.

7.2 System Architecture Components

7.2.1 Repository Scanner

The Repository Scanner is the entry point of the Arguss workflow. Its primary responsibility is identifying software dependencies within a repository.

The scanner analyzes source code repositories, dependency configuration files, package manifests, version information, and dependency relationships.

Supported dependency ecosystems include common software development environments such as:

- JavaScript/Node.js (package.json)
- Python (requirements.txt, pyproject.toml)
- Java (pom.xml)
- Go (go.mod)

The scanner creates an initial dependency inventory that becomes the foundation for further security analysis.

7.2.2 Dependency Analysis Engine

After collecting dependency information, Arguss processes relationships between software components.

The Dependency Analysis Engine identifies direct dependencies, transitive dependencies, dependency versions, available upgrade paths, and potential dependency conflicts.

This analysis is important because vulnerabilities are often introduced indirectly. For example:

Indirect dependency inheritance

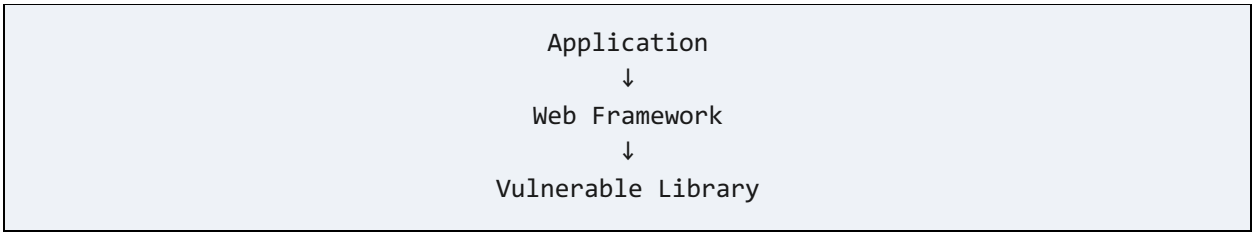

A developer may not directly install the vulnerable library but may still inherit its security risk through another package. Understanding these relationships allows Arguss to recommend the correct remediation strategy.

7.3 Software Bill of Materials Generation

7.3.1 Purpose of SBOM Integration

A Software Bill of Materials (SBOM) provides a complete inventory of software components contained within an application.

Arguss generates SBOM information to improve visibility into what dependencies exist, where they originate, which versions are deployed, and how components relate to one another.

SBOM generation supports both vulnerability management and software supply chain transparency.

7.3.2 SBOM Data Model

The Arguss SBOM includes information such as:

Field	Description
Package Name	Software component identifier
Version	Installed dependency version
Ecosystem	Package manager source
Dependency Type	Direct or transitive
Vulnerability Status	Known security issues
Upgrade Recommendation	Available remediation

This structured information allows Arguss to connect vulnerability findings directly to application dependencies.

7.4 Security Intelligence Integration

Arguss combines information from multiple external security sources to improve decision accuracy. Rather than relying on a single vulnerability database, the platform aggregates several security signals.

7.4.1 Vulnerability Intelligence Sources

Common Vulnerabilities and Exposures (CVE)

Provides vulnerability identifiers, technical descriptions, affected versions, and available fixes.

Common Vulnerability Scoring System (CVSS)

Provides severity ratings, attack complexity, required privileges, and potential impact.

Exploit Prediction Scoring System (EPSS)

Provides probability of exploitation and real-world exploit likelihood.

Known Exploited Vulnerabilities (KEV)

Provides evidence of active exploitation and prioritization signals for urgent remediation.

7.4.2 Package Trust Intelligence

Arguss evaluates package reliability using ecosystem health signals, including OpenSSF Scorecard metrics, repository activity, release frequency, maintenance history, and contributor activity.

These signals help determine whether a dependency represents potential supply chain risk beyond known vulnerabilities.

7.5 Backend Architecture

7.5.1 Backend Framework

The Arguss backend was developed using Python-based services designed to support security analysis workflows.

The backend manages repository scanning, dependency processing, security API integration, risk scoring, and remediation recommendations.

A modular backend design allows additional security analysis modules to be added without redesigning the entire system.

7.5.2 Data Processing Pipeline

The backend workflow follows this sequence:

Backend data processing pipeline

```
Repository Input → Dependency Extraction → Security Data Collection → Risk
Analysis → Confidence Evaluation → Recommendation Generation
```

Each stage produces structured output that is passed to the next component. This approach improves debugging, testing, transparency, and future expansion.

7.6 Frontend Dashboard

7.6.1 Purpose

The Arguss dashboard provides developers and security teams with a visual interface for understanding dependency risk.

The dashboard focuses on presenting security information in an actionable format rather than overwhelming users with raw vulnerability data.

7.6.2 Dashboard Capabilities

Vulnerability Overview

Displays the number of detected vulnerabilities, severity distribution, and priority findings.

Dependency Risk Visualization

Shows vulnerable packages, dependency relationships, and trust indicators.

Remediation Recommendations

Provides recommended upgrades, confidence level, and reasoning behind decisions.

Security Summary

AI-generated summaries translate technical findings into understandable security insights. Example:

> *"Three dependencies contain known vulnerabilities. Two have high-confidence fixes available and can be safely upgraded automatically. One requires manual review due to compatibility concerns."*

7.7 AI-Assisted Security Analysis

7.7.1 Role of AI Integration

Arguss incorporates AI capabilities to improve communication and analysis rather than replacing security decision-making.

The AI component assists with executive summaries, vulnerability explanations, remediation descriptions, and developer guidance.

7.7.2 Human-Centered AI Design

AI-generated recommendations are always grounded in collected security evidence. The system avoids making unsupported security claims by referencing vulnerability data, dependency information, confidence calculations, and analysis results.

This ensures AI enhances security workflows while maintaining transparency.

7.8 GitHub Workflow Integration

7.8.1 Development Workflow Integration

Arguss integrates with GitHub-based development workflows to reduce friction between security analysis and engineering processes.

The workflow enables repository scanning, automated analysis, pull request generation, and security review.

7.8.2 Automated Remediation Pull Requests

For high-confidence fixes, Arguss can generate remediation pull requests containing updated dependency versions, a security explanation, confidence rating, and validation information.

> *Example: Pull Request — Upgrade vulnerable-package 2.1.0 → 2.1.1. Reason: fixes actively exploited vulnerability CVE-XXXX. Confidence: High. Validation: dependency installation successful.*

This allows developers to review security changes using familiar workflows.

7.9 Technology Stack

Table 1. Arguss Technology Stack

Component	Technology
Frontend	React, TypeScript
Backend	Python, FastAPI
Database	PostgreSQL
AI Integration	Claude API
Repository Integration	GitHub API / Actions
Deployment	Fly.io
Containerization	Docker
Security Data	CVE, EPSS, KEV, OpenSSF

7.10 Deployment Architecture

Figure 4. Arguss Deployment Architecture

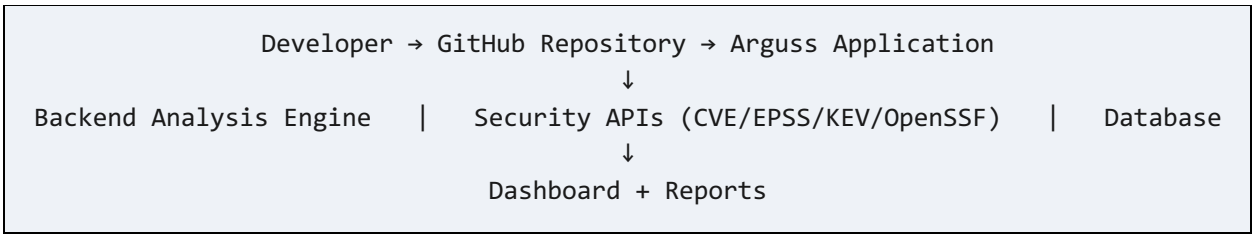

7.11 Implementation Challenges

Security Data Integration

Combining multiple security sources required normalization because different sources provide information using different formats and scoring systems.

Balancing Automation and Safety

A major design challenge was determining when automation should occur. Too little automation reduces efficiency; too much automation creates security and reliability risks. The Fix Confidence Engine was developed to address this balance.

Reliability of Automated Decisions

Security automation requires validation to avoid incorrect recommendations. Arguss addresses this through multiple independent signals, explainable decisions, and human approval for uncertain cases.

7.12 Summary

The Arguss implementation combines software security analysis, AI assistance, and automation into a unified platform.

Unlike traditional vulnerability scanners that stop at detection, Arguss provides an end-to-end workflow:

> *Identify → Analyze → Prioritize → Evaluate → Remediate*

Through modular architecture, multi-source intelligence, and confidence-based automation, Arguss demonstrates how autonomous security systems can improve software supply chain defense while maintaining developer control.

8. Evaluation Framework

8.1 Evaluation Overview

Evaluating an autonomous cybersecurity system requires more than measuring vulnerability detection accuracy. While identifying vulnerabilities is an important capability, the primary goal of Arguss is to improve the remediation decision-making process.

The evaluation of Arguss focuses on four key questions:

20. Can Arguss accurately identify software supply chain risks?
21. Can multi-source security intelligence improve vulnerability prioritization?
22. Can the Fix Confidence Engine safely determine when remediation should be automated?
23. Can Arguss reduce the operational burden placed on developers and security teams?

The evaluation framework measures both technical performance and practical usability.

Unlike traditional vulnerability scanners that are evaluated primarily on detection rates, Arguss is evaluated on its ability to produce accurate, explainable, and safe remediation recommendations.

8.2 Research Questions

RQ1: Vulnerability Detection Capability

Can Arguss accurately identify vulnerable dependencies within software repositories?

The first research question evaluates whether Arguss can successfully discover vulnerable components and associate them with relevant security information.

Evaluation focuses on dependency discovery accuracy, vulnerability identification, version matching accuracy, and SBOM completeness. A successful system should identify vulnerable dependencies while minimizing missed findings.

RQ2: Risk Prioritization Improvement

Does combining multiple security signals improve vulnerability prioritization compared with severity-only approaches?

Traditional vulnerability management often relies heavily on CVSS scores. However, severity alone does not represent practical risk.

Arguss improves prioritization by combining CVSS severity, EPSS exploit probability, KEV exploitation status, package trust signals, and dependency context.

The evaluation compares the traditional approach (Vulnerability → CVSS Score → Priority) with the Arguss approach (Vulnerability → Multiple Security Signals → Context-Aware Risk Assessment → Priority Recommendation). The goal is to determine whether Arguss better identifies vulnerabilities requiring immediate attention.

RQ3: Remediation Confidence Accuracy

Can the Fix Confidence Engine correctly determine when remediation is safe to automate? This is the primary evaluation objective of Arguss.

The system must balance two competing goals: avoiding unsafe automation (incorrectly automating a risky update may introduce application failures, dependency conflicts, or security regressions) and avoiding excessive manual review (requiring approval for every change reduces the efficiency benefits of automation).

The evaluation examines whether confidence classifications align with actual remediation outcomes.

8.3 Evaluation Metrics

8.3.1 Vulnerability Detection Metrics

Precision

Measures how many identified findings represent valid security issues. High precision reduces unnecessary investigation.

> *Precision = True Positive Findings / All Identified Findings*

Recall

Measures how many existing vulnerabilities Arguss successfully identifies. High recall reduces the chance of missed vulnerabilities.

> *Recall = True Positive Findings / All Actual Vulnerabilities*

False Positive Rate

Measures the number of findings incorrectly identified as security risks. Reducing false positives is important because excessive alerts contribute to security fatigue.

8.3.2 Remediation Recommendation Metrics

Confidence Accuracy

Measures whether Arguss correctly predicts remediation safety. Example:

Arguss Prediction	Actual Result	Outcome
High Confidence	Successful update	Correct
High Confidence	Failed update	Incorrect
Low Confidence	Required manual changes	Correct

Automation Success Rate

Measures how often automated remediation successfully completes. Factors include successful dependency upgrade, successful build, passing tests, and no introduced conflicts.

Human Escalation Rate

Measures how often Arguss correctly identifies situations requiring human review. A strong system should avoid both excessive escalation and unsafe automation.

8.4 Evaluation Dataset and Test Scenarios

8.4.1 Vulnerable Repository Testing

Arguss can be evaluated using repositories containing known vulnerable dependencies. Test cases include applications with outdated packages, multiple vulnerable dependencies, transitive vulnerabilities, and conflicting upgrade paths.

Each repository is analyzed to determine whether Arguss correctly identifies and prioritizes risks.

8.4.2 Remediation Scenario Testing

Evaluation scenarios examine different remediation conditions.

Scenario	Package	Current → Fixed Version	Characteristics	Expected Result
1: Simple Patch Update	library-x	1.4.2 → 1.4.3	Small version change; no breaking changes; strong package maintenance	High Confidence
2: Minor Version Upgrade	library-y	2.1.0 → 2.2.0	Security improvement; possible compatibility concerns	Medium Confidence

Scenario	Package	Current → Fixed Version	Characteristics	Expected Result
3: Major Version Migration	library-z	3.x → 4.x	Significant API changes; migration required	Low Confidence

8.5 Comparative Evaluation

Arguss can be compared against traditional dependency security tools.

Table 2. Capability Comparison

Capability	Traditional SCA	Automated Updates	Arguss
Vulnerability Detection	✓	✓	✓
SBOM Generation	Some	Limited	✓
Package Trust Analysis	Limited	No	✓
Exploit Intelligence	Some	Limited	✓
Remediation Recommendation	Limited	✓	✓
Fix Safety Evaluation	No	No	✓
Explainable Automation	Limited	Limited	✓
Human-in-the-Loop Control	Limited	Limited	✓

8.6 Developer Efficiency Evaluation

A key goal of Arguss is reducing the amount of manual security analysis required from developers. The evaluation measures:

Time to Understand Vulnerability

Traditional workflow: alert received → research vulnerability → check exploitability → research package → determine fix → apply update.

Arguss workflow: alert received → Arguss analysis → review recommendation → apply or approve fix.

The expected improvement is a reduction in investigation time.

Alert Fatigue Reduction

Arguss reduces alert fatigue by prioritizing findings based on real-world exploitability, application context, package reliability, and remediation safety. Instead of presenting hundreds of equally weighted vulnerabilities, Arguss highlights the issues requiring action.

8.7 Security Validation Approach

Because Arguss performs autonomous security recommendations, validation is critical. The evaluation process includes:

Recommendation Review

Security experts review whether recommendations are appropriate, explainable, and technically justified.

Remediation Testing

Generated fixes are evaluated through dependency installation, application builds, automated tests, and deployment validation.

Failure Analysis

Incorrect recommendations are analyzed to identify missing signals, incorrect assumptions, and data quality issues. This feedback improves future versions of the system.

8.8 Evaluation Limitations

Limited Repository Diversity

Testing may not represent every programming ecosystem or application architecture.

Dependency Metadata Availability

Some packages may lack sufficient information for accurate trust analysis.

CI/CD Dependency

Automated validation depends on the quality and availability of existing testing pipelines.

Evolving Threat Landscape

Software supply chain threats continuously change. Evaluation results may require periodic updates as new attack methods emerge.

8.9 Summary

The evaluation framework demonstrates that Arguss should be measured not only by its ability to discover vulnerabilities, but by its ability to make reliable remediation decisions.

The success of Arguss depends on three outcomes: accurate identification of software supply chain risks, improved prioritization through contextual analysis, and safe automation through confidence-based remediation.

By evaluating detection accuracy, remediation confidence, and developer efficiency, Arguss provides a measurable approach toward safer autonomous cybersecurity operations.

9. Security and Ethical Considerations

9.1 Overview

As cybersecurity systems become increasingly automated, organizations must carefully consider the risks and responsibilities associated with autonomous decision-making. While automation can significantly improve security operations, poorly designed automation can introduce new risks, including incorrect remediation, lack of accountability, and reduced human oversight.

Arguss was designed around the principle that cybersecurity automation should be defensible, transparent, and human-centered. The platform does not attempt to replace security professionals or developers. Instead, it augments human decision-making by reducing repetitive analysis and providing evidence-based recommendations.

This section discusses the ethical and security considerations involved in deploying Arguss, including explainability, human oversight, responsible AI usage, privacy, and equitable access to security automation.

9.2 Human-in-the-Loop Security

9.2.1 Importance of Human Oversight

Autonomous security systems operate in complex environments where technical data alone may not capture every operational consideration.

For example, an automated system may determine that upgrading a dependency is technically safe, but a development team may know that the application is preparing for a major release, the dependency supports a business-critical feature, the organization has compliance requirements, or additional testing is required before deployment.

Human expertise remains essential for these contextual decisions. Arguss incorporates human oversight by ensuring that automation levels are determined by confidence.

9.2.2 Confidence-Based Human Involvement

Arguss divides remediation decisions into three categories:

Confidence Level	System Action	Human Role
High Confidence	Generate automatic remediation	Review results
Medium Confidence	Create recommended fix	Approve changes
Low Confidence	Require investigation	Make final decision

This approach avoids two extremes. Over-automation — automatically applying every fix may create operational instability. Under-automation — requiring manual review for every vulnerability prevents organizations from responding efficiently.

Arguss aims to establish a balanced model where automation is applied only when sufficient evidence exists.

9.3 Explainability and Transparency

9.3.1 Importance of Explainable Security Decisions

Security decisions often affect critical systems. Therefore, users must understand why an automated system recommends a particular action.

A remediation recommendation without explanation creates several problems: developers may not trust the system, security teams cannot verify decisions, errors become difficult to identify, and accountability becomes unclear.

Arguss addresses this issue by providing reasoning behind each recommendation.

9.3.2 Explainable Remediation Example

Instead of producing “Upgrade package X,” Arguss provides:

> *Recommendation: upgrade package X from version 2.4.1 to 2.4.2.*

Reasoning:

- The package contains a high-severity vulnerability.
- The vulnerability has active exploitation indicators.
- The update is a patch-level change.
- The package has strong maintenance signals.
- No dependency conflicts were detected.

Confidence Level: High Confidence. This explanation allows users to understand the evidence supporting the decision.

9.4 Responsible Use of Artificial Intelligence

9.4.1 Role of AI in Arguss

Arguss uses artificial intelligence to improve security analysis and communication. AI capabilities assist with generating security summaries, explaining vulnerability findings, translating technical information for different audiences, and providing remediation context.

However, AI is not responsible for independently making unrestricted security decisions. The final remediation decision remains controlled by the Fix Confidence Engine and human review process.

9.4.2 Preventing AI Over-Reliance

AI-generated security recommendations introduce potential risks: incorrect explanations, missing context, overconfident recommendations, and misinterpretation of security data.

To reduce these risks, Arguss follows several principles:

Evidence-Based Output

AI responses are generated from collected security data rather than unsupported assumptions.

Limited Authority

AI assists decision-making but does not override security policies.

Human Verification

Users can review and validate recommendations before applying significant changes.

9.5 Security of the Arguss Platform

9.5.1 Protecting Repository Information

Because Arguss analyzes software repositories, protecting source code and dependency information is a critical security requirement.

Potential risks include exposure of proprietary code, leakage of dependency information, and unauthorized repository access.

Security controls include limited repository permissions, secure authentication, least privilege access, and protected API communication.

9.5.2 Credential Management

Access credentials represent a significant security concern in developer tooling. Arguss follows secure credential management practices: avoiding unnecessary permissions, using scoped authentication, limiting access duration, and protecting sensitive tokens.

The principle of least privilege ensures that Arguss only receives the permissions required to perform security analysis.

9.6 Privacy Considerations

9.6.1 Minimizing Data Collection

Security tools must balance visibility with privacy. Arguss follows a data minimization approach by collecting only information necessary for analysis, such as dependency metadata, package versions, and security intelligence.

The platform does not require unnecessary access to unrelated application data.

9.6.2 Transparency with Users

Organizations using Arguss should understand what information is collected, how analysis is performed, what outputs are generated, and who can access security findings.

Transparency improves user trust and supports responsible deployment.

9.7 Equity and Accessibility Considerations

9.7.1 Making Security Automation Available

Advanced cybersecurity capabilities are often concentrated among organizations with large security teams and significant resources. Smaller organizations may struggle with limited security expertise, high vulnerability volumes, and a lack of dedicated security personnel.

Arguss aims to reduce this gap by providing automated security analysis and remediation assistance.

9.7.2 Avoiding Unequal Security Outcomes

Automation systems must avoid creating security advantages only for organizations with extensive resources.

Arguss supports broader security accessibility by reducing manual workload, providing explainable recommendations, and helping developers without deep security expertise understand risks.

9.8 Ethical Alignment with the E3 Framework

Arguss can be evaluated through the E3 framework: Ethics, Equity, and Equality.

Ethics

Arguss emphasizes responsible automation by maintaining human oversight, providing explainable recommendations, and avoiding unsafe automatic changes. The system prioritizes security improvement while maintaining accountability.

Equity

Arguss improves access to cybersecurity capabilities by reducing the expertise required to analyze dependency risks. Organizations with smaller security teams can benefit from automated vulnerability prioritization, security recommendations, and risk explanations.

Equality

Arguss provides consistent security analysis based on evidence rather than organizational size or resources. Every analyzed dependency is evaluated using the same security principles: vulnerability impact, exploitability, package trust, and remediation confidence.

9.9 Potential Risks and Mitigations

Risk 1: Incorrect Automated Remediation

Impact: a wrong update may introduce application failures.

Mitigation: confidence scoring, validation checks, and human approval for uncertain cases.

Risk 2: Incomplete Security Data

Impact: missing information may lead to inaccurate recommendations.

Mitigation: multiple intelligence sources, transparent uncertainty reporting, and conservative automation decisions.

Risk 3: Over-Trust in Automation

Impact: users may accept recommendations without understanding risks.

Mitigation: explainable outputs, human review requirements, and clear confidence indicators.

9.10 Summary

Autonomous cybersecurity systems require careful consideration of security, ethics, and accountability. Arguss addresses these challenges by implementing responsible automation principles: automate only when confidence is high, explain every recommendation, preserve human decision-making, protect sensitive information, and improve access to security capabilities.

By combining automation with transparency and oversight, Arguss demonstrates how AI-assisted cybersecurity systems can strengthen software supply chain security while maintaining trust and accountability.

10. Limitations

10.1 Overview

Although Arguss introduces a new approach to software supply chain security through confidence-based autonomous remediation, several limitations remain. These limitations reflect both technical constraints and broader challenges associated with building reliable cybersecurity automation systems.

The goal of Arguss is not to eliminate all human involvement in vulnerability management. Instead, the platform aims to reduce repetitive security analysis while improving the quality and speed of remediation decisions.

Understanding these limitations is essential for responsible deployment and provides direction for future improvements.

10.2 Dependency Ecosystem Coverage

10.2.1 Current Coverage Limitations

Modern software development relies on a wide variety of programming languages and package ecosystems. While Arguss supports common dependency environments, complete coverage across all ecosystems remains challenging.

Different ecosystems have unique package management systems, dependency structures, release practices, and security reporting mechanisms. For example, dependency resolution in Python differs significantly from dependency management in JavaScript, Java, or Go.

Expanding Arguss across additional ecosystems requires specialized analysis capabilities and deeper integration with ecosystem-specific security data sources.

10.2.2 Future Expansion Needs

Future versions of Arguss could expand support for:

- RubyGems.
- Rust Cargo packages.
- PHP Composer.
- .NET NuGet packages.
- Mobile application dependencies.

Broader ecosystem coverage would allow organizations with diverse technology stacks to adopt Arguss more effectively.

10.3 Dependency Metadata Limitations

10.3.1 Reliance on Available Security Information

Arguss relies on external security intelligence sources to evaluate dependency risk. These sources provide valuable information, but they may have limitations: vulnerabilities may be disclosed after a delay, package ownership information may be incomplete, repository health signals may not always reflect actual security practices, and exploit information may change over time.

As a result, Arguss recommendations are dependent on the quality and availability of external security data.

10.3.2 Incomplete Package Context

Some dependencies may not provide enough information to accurately determine trust. For example, a new package may appear risky due to limited history, a small but legitimate project may have limited contributors, and a popular package may still contain security weaknesses.

Arguss addresses uncertainty through confidence scoring, but complete certainty is not always possible.

10.4 Limited Runtime Exploitability Analysis

10.4.1 Static Analysis Focus

The current implementation of Arguss primarily focuses on dependency and metadata analysis, including package versions, vulnerability records, dependency relationships, and repository information.

However, static analysis alone cannot always determine whether vulnerable code paths are actively reachable during application execution.

10.4.2 Need for Runtime Context

Future improvements could incorporate runtime security signals such as application behavior monitoring, runtime vulnerability detection, code execution tracing, and production environment context.

For example, a vulnerable library may exist within an application but never execute the affected functionality. Runtime analysis could improve prioritization accuracy by determining actual exposure.

10.5 Dependence on CI/CD Validation

10.5.1 Importance of Automated Testing

The Fix Confidence Engine benefits from validation information such as build results, automated tests, and deployment checks. However, not all repositories maintain comprehensive testing pipelines.

Applications with limited testing coverage create uncertainty because Arguss has fewer signals to determine whether a remediation will be safe.

10.5.2 Impact on Automation Confidence

A repository with strong automated testing provides greater confidence. A repository with extensive unit tests, integration tests, and automated deployment validation can expect higher remediation confidence; a repository with limited or no automated tests and a manual deployment process can expect lower remediation confidence.

Future versions of Arguss could provide additional validation methods for repositories with limited CI/CD maturity.

10.6 AI-Generated Explanation Limitations

10.6.1 AI Reasoning Boundaries

Arguss uses AI assistance to generate summaries and explain security findings. While AI improves communication, it introduces potential limitations, including misinterpretation of technical information, overly generalized explanations, and missing application-specific context.

AI-generated outputs should therefore be considered assistance rather than the final security authority.

10.6.2 Mitigation Approach

Arguss reduces AI-related risks through grounding explanations in collected security data, limiting AI decision authority, maintaining deterministic confidence evaluation, and allowing human review.

The Fix Confidence Engine, rather than the AI component alone, determines automation decisions.

10.7 Limited Historical Remediation Data

10.7.1 Challenge

Determining whether a remediation is safe requires understanding historical outcomes. However, many organizations lack detailed records of previous dependency updates, failed upgrades, successful remediation patterns, and application-specific compatibility behavior.

Without sufficient historical data, confidence predictions may be conservative.

10.7.2 Future Improvement

Future versions could incorporate organization-specific remediation history, machine learning models trained on previous fixes, community remediation patterns, and dependency compatibility databases. This would allow Arguss to improve confidence estimates over time.

10.8 Human Decision Dependency

10.8.1 Remaining Need for Expertise

Although Arguss reduces manual security analysis, some decisions still require human expertise, such as business-critical applications, regulatory requirements, architectural changes, and complex dependency migrations.

Security professionals remain essential for evaluating broader organizational risk.

10.8.2 Balancing Automation and Control

A major challenge in autonomous security systems is determining the correct level of automation. Too little automation maintains excessive manual workload and slows vulnerability response. Too much automation creates operational risk and reduces human oversight.

Arguss addresses this challenge through confidence-based decision-making, but determining the ideal balance remains an ongoing research area.

10.9 Scalability Considerations

10.9.1 Large-Scale Repository Management

Enterprise organizations may manage thousands of repositories across multiple teams. Scaling Arguss across large environments introduces challenges such as processing large dependency graphs, managing frequent scans, prioritizing critical applications, and reducing unnecessary notifications.

10.9.2 Enterprise Integration

Future enterprise deployments may require integration with Security Information and Event Management (SIEM) systems, vulnerability management platforms, enterprise identity providers, and governance and compliance systems. These integrations would improve Arguss adoption within larger organizations.

10.10 Summary

Arguss represents a significant step toward safer software supply chain automation, but several limitations remain. Current challenges include limited ecosystem coverage, dependence on external security intelligence, lack of complete runtime context, CI/CD dependency, AI explanation limitations, and the need for human judgment in complex cases.

These limitations do not prevent Arguss from providing value. Instead, they highlight areas for continued research and development.

The platform is designed around the principle that effective cybersecurity automation does not require eliminating uncertainty — it requires recognizing uncertainty and responding appropriately.

By identifying when automation is safe and when human expertise is required, Arguss provides a foundation for more responsible autonomous security operations.

11. Future Work

11.1 Overview

Arguss establishes a foundation for defensible autonomous remediation by combining vulnerability intelligence, package trust analysis, dependency context, and confidence-based automation. However, software supply chain security continues to evolve rapidly, requiring continuous improvements in analysis capabilities, automation accuracy, and integration with modern development environments.

Future development of Arguss will focus on expanding technical capabilities while maintaining the platform's core principles:

- Explainable security decisions.
- Human-centered automation.
- Evidence-based remediation.
- Safe autonomous action.

The future roadmap focuses on improving ecosystem coverage, increasing contextual awareness, strengthening AI capabilities, and enabling enterprise-scale adoption.

11.2 Expanded Programming Language and Ecosystem Support

11.2.1 Broader Dependency Analysis

A major area of future development is expanding support for additional programming languages and package ecosystems.

Currently, software projects rely on a diverse range of dependency management systems. Supporting more ecosystems would allow Arguss to provide broader software supply chain visibility. Future support may include Rust Cargo, RubyGems, PHP Composer, .NET NuGet, Swift Package Manager, and mobile application dependencies.

11.2.2 Cross-Ecosystem Risk Analysis

Future versions of Arguss could analyze applications containing multiple technology stacks. Modern applications frequently combine backend services, frontend frameworks, mobile applications, cloud infrastructure, and third-party APIs.

A cross-ecosystem analysis capability would allow Arguss to evaluate the complete application supply chain rather than isolated dependency groups.

11.3 Runtime-Aware Security Analysis

11.3.1 Moving Beyond Static Dependency Analysis

The current Arguss implementation focuses primarily on dependency metadata and repository-level analysis. Future versions could incorporate runtime security signals to better determine whether vulnerabilities represent actual exposure, including runtime dependency monitoring, code execution tracing, application behavior analysis, and production environment awareness.

11.3.2 Reachability Analysis

A major challenge in vulnerability management is determining whether vulnerable code is actually reachable. For example, a vulnerable package may exist within an application but contain functionality that is never executed.

Future Arguss capabilities could determine whether vulnerable functions are called, which application components are affected, and whether exploitation paths exist — improving prioritization by focusing attention on vulnerabilities with real application impact.

11.4 Advanced Fix Confidence Modeling

11.4.1 Machine Learning-Based Confidence Prediction

The current Fix Confidence Engine uses security signals and remediation characteristics to determine automation confidence. Future versions could incorporate machine learning models trained on historical remediation outcomes, including previous successful upgrades, failed remediation attempts, dependency compatibility patterns, and organization-specific application behavior.

This could allow Arguss to make increasingly accurate predictions over time.

11.4.2 Organization-Specific Learning

Different organizations have different risk tolerances and development practices. For example, a financial organization may require stricter approval policies than a small development team.

Future versions could allow organizations to customize automation thresholds, approval requirements, risk tolerance, and compliance policies, making Arguss adaptable across different environments.

11.5 Enhanced AI Security Assistant

11.5.1 More Advanced Security Reasoning

Future AI improvements could expand Arguss beyond summary generation into deeper security assistance, including explaining complex dependency relationships, answering developer security questions, suggesting alternative remediation paths, and assisting with migration planning.

11.5.2 Security Policy-Aware Recommendations

Future AI capabilities could incorporate organizational security policies. For example, a company may define that production systems require manual approval, that critical vulnerabilities must be fixed within a specific timeframe, or that certain dependencies are prohibited.

Arguss could use these policies to generate recommendations aligned with organizational requirements.

11.6 Continuous Software Supply Chain Monitoring

11.6.1 From Periodic Scanning to Continuous Protection

Many vulnerability management workflows rely on scheduled scans. However, software supply chain risks change continuously: new vulnerabilities are disclosed, packages are updated, exploits become available, and dependencies change.

Future Arguss versions could provide continuous monitoring by automatically detecting newly disclosed vulnerabilities, dependency changes, emerging exploit activity, and package trust degradation.

11.6.2 Real-Time Risk Updates

A future monitoring system could automatically update dependency risk assessments when new information becomes available. For example, an initial assessment stating that a package has no known exploitation activity could later be updated after the vulnerability is added to the KEV catalog — at which point Arguss could automatically increase priority and notify developers.

11.7 Enterprise Security Integration

11.7.1 Security Operations Integration

For enterprise adoption, Arguss could integrate with existing security platforms, including SIEM platforms, vulnerability management systems, governance platforms, and ticketing systems. This would allow Arguss findings to become part of broader security operations workflows.

11.7.2 Identity and Access Integration

Enterprise environments require strong access controls. Future improvements could include integration with Single Sign-On (SSO), identity providers, role-based access control, and audit logging — ensuring that remediation actions follow organizational security policies.

11.8 Improved Validation and Testing Capabilities

11.8.1 Automated Remediation Verification

Future versions could improve confidence evaluation through deeper validation, including automated testing, build verification, deployment simulation, and security regression testing.

11.8.2 Pre-Merge Security Validation

Arguss could integrate earlier into the software development lifecycle. For example, before merging a pull request, Arguss could analyze dependency changes, evaluate security impact, predict remediation confidence, and approve or request review — moving security analysis closer to the development process.

11.9 Community-Driven Security Intelligence

11.9.1 Shared Remediation Knowledge

Future Arguss development could explore community-driven intelligence sharing. Organizations could contribute anonymized information about successful dependency upgrades, failed remediation attempts, compatibility issues, and package reliability trends — improving remediation confidence models across the ecosystem.

11.9.2 Balancing Collaboration and Privacy

Any community intelligence system must protect sensitive information. Future implementations should ensure data anonymization, opt-in participation, clear privacy controls, and secure data handling.

11.10 Long-Term Vision

The long-term goal of Arguss is to create a more intelligent and responsible approach to software security automation.

The future of vulnerability management should not require organizations to choose between slow manual remediation and unsafe full automation. Instead, security systems should provide continuous awareness, context-driven decisions, explainable recommendations, and safe autonomous action.

Arguss aims to contribute to this future by creating security automation that understands uncertainty and responds appropriately.

11.11 Summary

Future development of Arguss will focus on expanding capabilities while maintaining responsible automation principles. Key future directions include expanded language ecosystem support, runtime vulnerability analysis, improved confidence prediction, advanced AI security assistance, continuous monitoring, enterprise integration, and enhanced validation workflows.

Through these improvements, Arguss can evolve from a remediation assistant into a comprehensive software supply chain security platform capable of supporting organizations throughout the entire software lifecycle.

12. Conclusion

12.1 Overview

The rapid growth of software supply chains has created new challenges for organizations attempting to secure modern applications. Applications increasingly depend on thousands of open-source components, creating complex dependency relationships and expanding the potential attack surface.

Traditional vulnerability management approaches have primarily focused on identifying security issues. While vulnerability detection remains essential, organizations face a more difficult challenge: determining which vulnerabilities require action, which fixes are safe to apply, and when automation can be trusted.

Arguss addresses this challenge by introducing a defensible autonomous remediation framework that combines security intelligence, dependency context, package trust analysis, and confidence-based automation.

12.2 Key Contributions

12.2.1 Moving Beyond Vulnerability Detection

The primary contribution of Arguss is shifting the focus of software security from detection to remediation decision-making.

Traditional security tools answer “What vulnerabilities exist?” Arguss expands this question to “Which vulnerabilities matter most, what remediation should be performed, and how confident are we that the fix is safe?”

This approach helps organizations move from reactive vulnerability management toward proactive risk reduction.

12.2.2 Multi-Lens Security Analysis

Arguss introduces a Three-Lens Risk Assessment Framework that evaluates dependencies through multiple perspectives: vulnerability analysis (severity, exploitability, active exploitation status, available security fixes), package trust analysis (open-source project health, maintenance activity, repository security signals, supply chain reliability), and dependency context analysis (dependency relationships, application importance, direct versus transitive usage, upgrade impact).

By combining these perspectives, Arguss provides a more complete understanding of software risk than vulnerability severity alone.

12.2.3 Confidence-Based Autonomous Remediation

A central innovation of Arguss is the Fix Confidence Engine. Instead of automatically applying every available update, Arguss determines whether sufficient evidence exists to safely automate remediation.

The system enables automatic remediation for high-confidence fixes, developer approval for medium-confidence fixes, and manual investigation for uncertain scenarios — balancing efficiency with security responsibility.

12.3 Impact on Cybersecurity Operations

12.3.1 Reducing Security Team Workload

Security teams often struggle with large volumes of vulnerability findings. Arguss reduces this burden by prioritizing meaningful risks, reducing unnecessary investigation, providing evidence-based recommendations, and automating low-risk remediation tasks — allowing security professionals to focus on complex threats requiring human expertise.

12.3.2 Improving Developer Security Experience

Developers frequently receive security alerts without sufficient context or guidance. Arguss improves developer workflows by providing clear explanations, recommended fixes, confidence ratings, and integration into existing development processes.

Instead of requiring developers to become vulnerability management experts, Arguss provides security guidance directly within their workflow.

12.3.3 Strengthening Software Supply Chain Security

Open-source dependencies are essential to modern software development but introduce significant supply chain risks. Arguss improves supply chain visibility through SBOM generation, dependency analysis, package trust evaluation, and continuous security intelligence — enabling organizations to better understand and manage the software components they rely on.

12.4 Responsible Automation as a Security Principle

A key lesson from Arguss is that effective cybersecurity automation is not defined by how many decisions a system can make independently.

Instead, successful automation depends on knowing when automation is appropriate, when additional evidence is required, and when human expertise is necessary. Arguss follows the principle:

> *Automate confidently, escalate responsibly.*

This ensures that automation strengthens security operations without creating new risks.

12.5 Broader Implications for AI in Cybersecurity

Arguss demonstrates a broader direction for artificial intelligence in cybersecurity: AI systems should function as decision-support partners rather than uncontrolled autonomous agents.

Future cybersecurity systems should prioritize explainability, transparency, human oversight, and evidence-based reasoning. By combining AI assistance with deterministic security analysis, Arguss provides a model for responsible AI adoption in security operations.

12.6 Final Perspective

Software security challenges will continue to grow as applications become more distributed, dependencies become more complex, and attackers increasingly target software supply chains.

Organizations need security solutions that can operate at modern software scale while maintaining trust and accountability.

Arguss provides a foundation for this future by combining vulnerability intelligence, package trust analysis, dependency awareness, AI-assisted explanations, and confidence-based remediation.

Rather than simply identifying security problems, Arguss helps organizations make better security decisions.

The future of cybersecurity automation is not replacing human judgment — it is enhancing human capability with systems that understand risk, explain decisions, and act safely.