SANS | Research Program

PRODUCT REVIEW

A New Era in Vulnerability Management: A SANS Review of the Seemplicity Platform

Written by **Dave Shackleford**

July 2025 – Updated April 2026

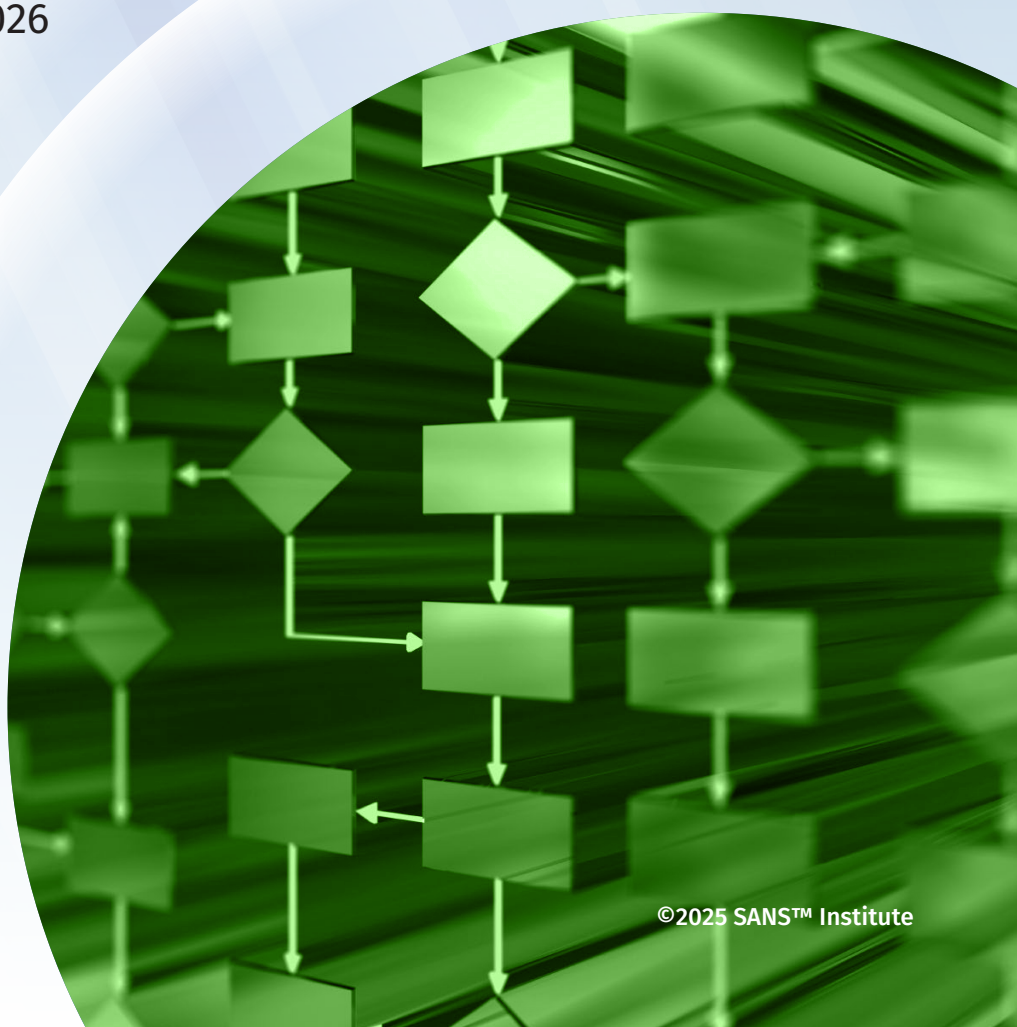

seemplicity

Introduction

Vulnerability management remains one of the most critical (and yet chronically underperforming) areas of cybersecurity in modern enterprises. According to the 2025 Verizon Data Breach Investigations Report, the exploitation of vulnerabilities has continued to grow as an initial access vector for breaches, reaching 20%.[1] Yet, despite advances in detection technologies, many organizations still take weeks or even months to remediate known, exploitable vulnerabilities. This delay is often due to fragmented tooling, poor asset inventory, limited prioritization strategies, and lack of alignment between security and IT operations. As attack surfaces expand with cloud, SaaS, and remote work, the inability to close gaps quickly is creating a persistent and compounding risk problem.

In many organizations, vulnerability management programs are broken not because of lack of visibility, but because of ineffective processes and bottlenecks in remediation. Common issues include an overreliance on outdated Common Vulnerability Scoring System (CVSS) scores, lack of business context, under-resourced patching teams, and "alert fatigue" caused by noisy or duplicative findings across scanning tools. Furthermore, vulnerabilities often require coordination across infrastructure, DevOps, and business teams—slowing down what should be a fast-moving response cycle. Given that threat actors now routinely exploit publicly disclosed vulnerabilities within days (or even hours), the current pace of remediation is simply not sufficient.

The need for faster and more intelligent vulnerability remediation is urgent. Threat actors are increasingly automating exploitation, leveraging AI to scan for weak points at scale. Without a shift toward continuous, risk-based prioritization and operational alignment, organizations will continue to fall behind. Improving this process means investing not only in better tools but also in better collaboration, contextual risk analysis, and real-time remediation workflows that reduce dwell time and limit exposure. The sooner organizations make this pivot, the better positioned they will be to defend against a rapidly evolving threat landscape.

[1] "2025 Data Breach Investigations Report," www.verizon.com/business/resources/reports/dbir

Unifying Assessment and Remediation Across Teams

Cyber risk remediation today is increasingly fragmented and slow, largely due to the overwhelming sprawl of risk domains organizations must manage. Modern enterprises must contend with vulnerabilities across cloud environments, traditional infrastructure, third-party services, source code repositories, SaaS platforms, penetration test reports, configuration baselines, and more. Each of these domains may report risk through different tools, formats, and severity scoring systems, making it difficult to correlate and prioritize findings. This siloed risk visibility leads to duplicate or conflicting remediation paths, inconsistent ownership, and, ultimately, a backlog of unresolved security issues—often measured in the tens or hundreds of thousands. The complexity of prioritizing what matters most is made worse by poor context about asset criticality or exploitability, leaving security teams unable to efficiently reduce the most impactful risks.

Equally challenging is the coordination required between multiple internal teams who must act on remediation tasks. Infrastructure, cloud, application development, DevOps, database, and business operations teams all play a role in resolving different categories of security issues. However, remediation efforts often break down because teams work on different schedules, follow different priorities, and may lack a shared understanding of what the risk means. Security teams are frequently left chasing owners, translating findings into technical tasks, and aligning on timelines—often without integrated tooling to manage tasks, statuses, and timelines across teams. Without a central remediation framework or governance model, organizations struggle to ensure accountability, measure remediation velocity, or track progress toward risk reduction goals. Coordination isn't just about who fixes what—it's about embedding security risk decisions into the planning and operational cycles of every team that can influence the attack surface.

The Seemplicity platform is focused on bridging the gap between vulnerability discovery domains (in code, scanners, cloud service, and any other source of vulnerability reporting information) and the actual "fixers" that are responsible for building and maintaining the various environments. Seemplicity automates the manual work across vulnerability management; application security; cloud and asset discovery; attack surface management; and governance, risk, and compliance (GRC). This allows teams to focus on reducing risk rather than spending time tracking and managing tickets. For this review, we were provisioned access into the platform in a demo account with a variety of real-world data sources and applications running across a wide range of different types of environments.

We were initially presented with the primary *Organizational* dashboard, which shows an overview of security findings detected in the environment; breakdowns of findings by data source, priority, and ticket status; and some overall categorization and classification information that includes the "top 10" most common findings noted in the environment (see Figure 1).

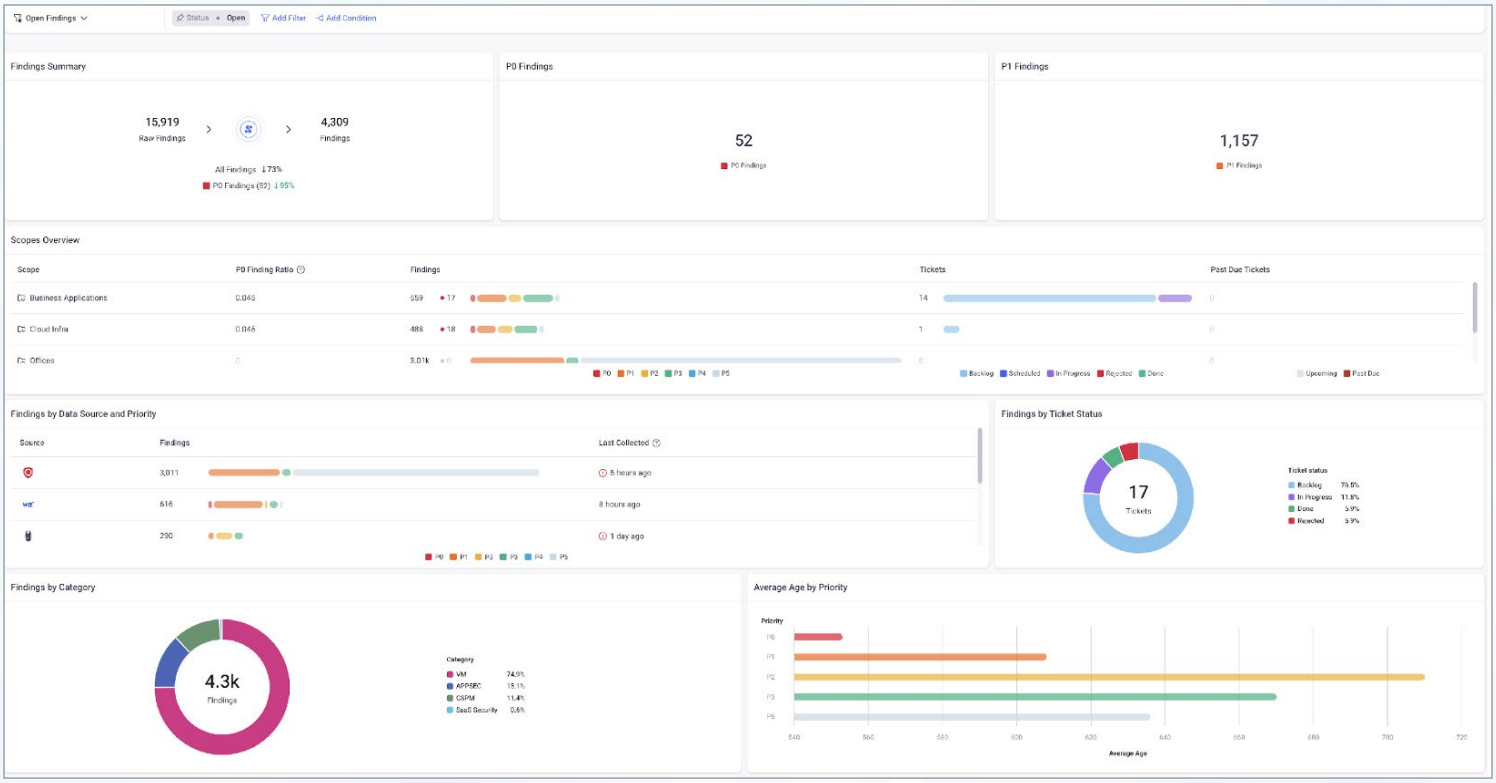

Figure 1. Our Initial Seemplicity Organizational Dashboard

Looking at this dashboard provided an enormous range of insights. First, we noted the number of "raw" findings that came in from our security event sources and how the platform distilled these into "actual" findings to reduce false positives and noise (more on this later). We also can see specific priority levels for findings, which helps save time by directing SOC analysts and other stakeholders to the most important issues noted. Findings broken down by data source, category (where they apply and the asset type), and age (how long they've been in the environment in this state) are also highlighted, helping analysts of all types to quickly see where the primary and most important issues may lie. Last but not least, an aggregate number of the most common findings can help to pinpoint the scope of vulnerabilities: How many do we have, and where are they found in our ecosystem?

Given the variety of different stakeholders who would likely want or need access to this platform, we reviewed the access controls available to us. First, we can define different scopes of access within the overall *Findings* area. Organizations can use the Seemplicity filtering model to create custom scopes of assets and findings across assets that users and groups can be granted access to. We created a simple scope with some different operating systems and project types, as shown in Figure 2.

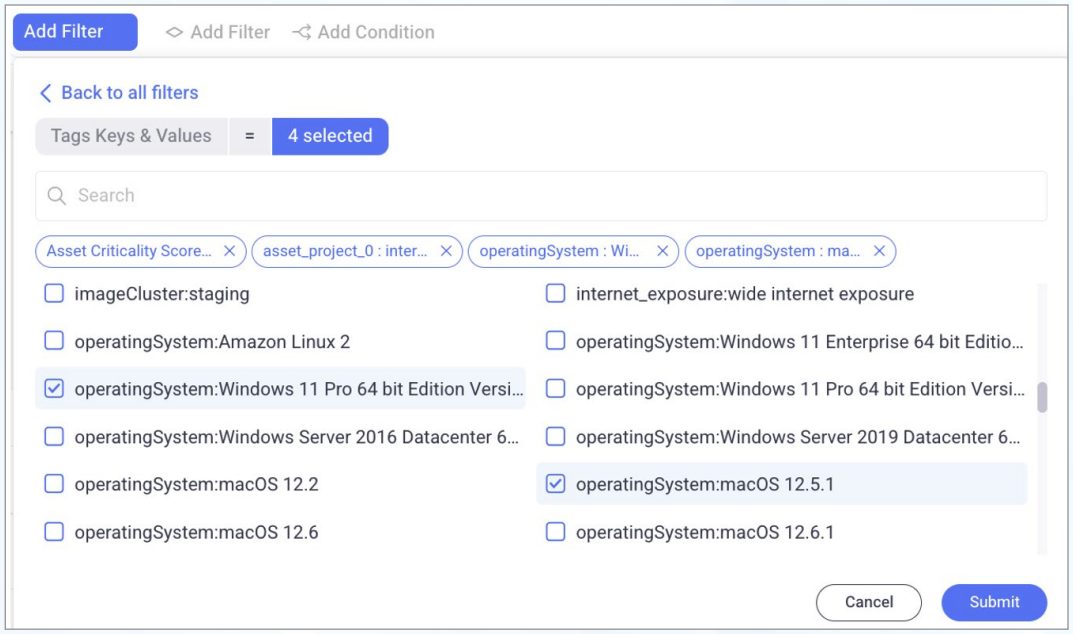

Figure 2. Creating a Custom Scope for Access

Scopes can have much more granular applicability to the variety of data sources and types ingested in the platform, too. Once a scope is created, you can assign a profile of users to it (or create a new one), like we did in Figure 3.

User and group integration through federated Security Assertion Markup Language (SAML) and API tokens is simple to configure. We simply reviewed these options, but they align with any modern identity and access management (IAM) service model today.

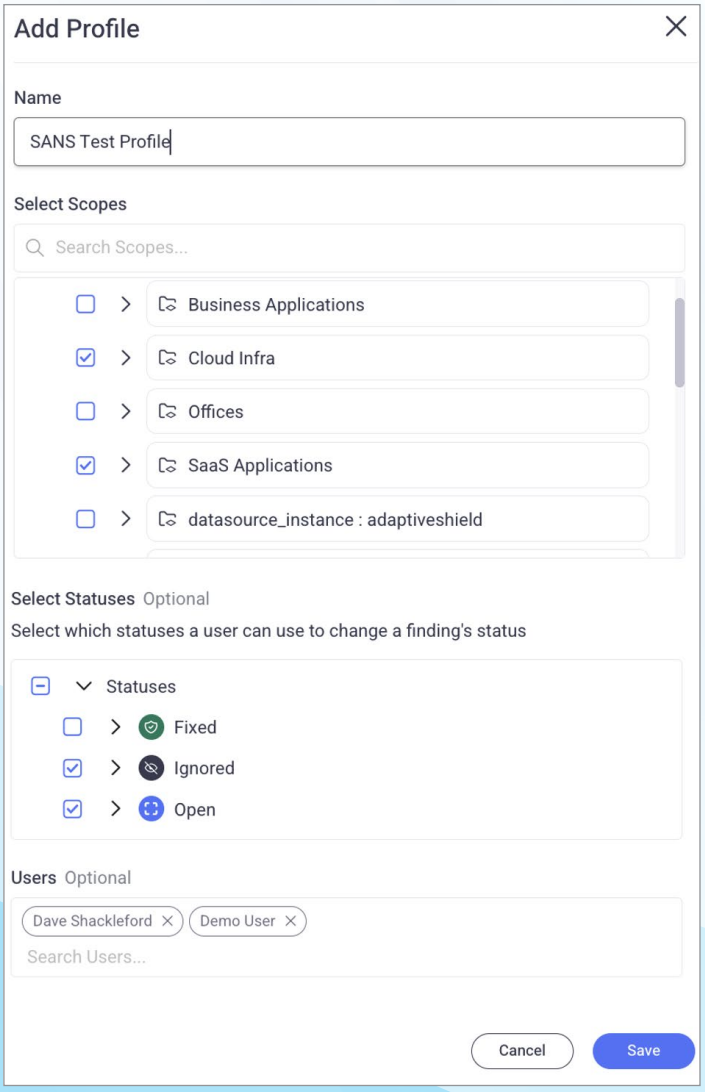

Figure 3. Assigning a Profile to a Scope for Access Management

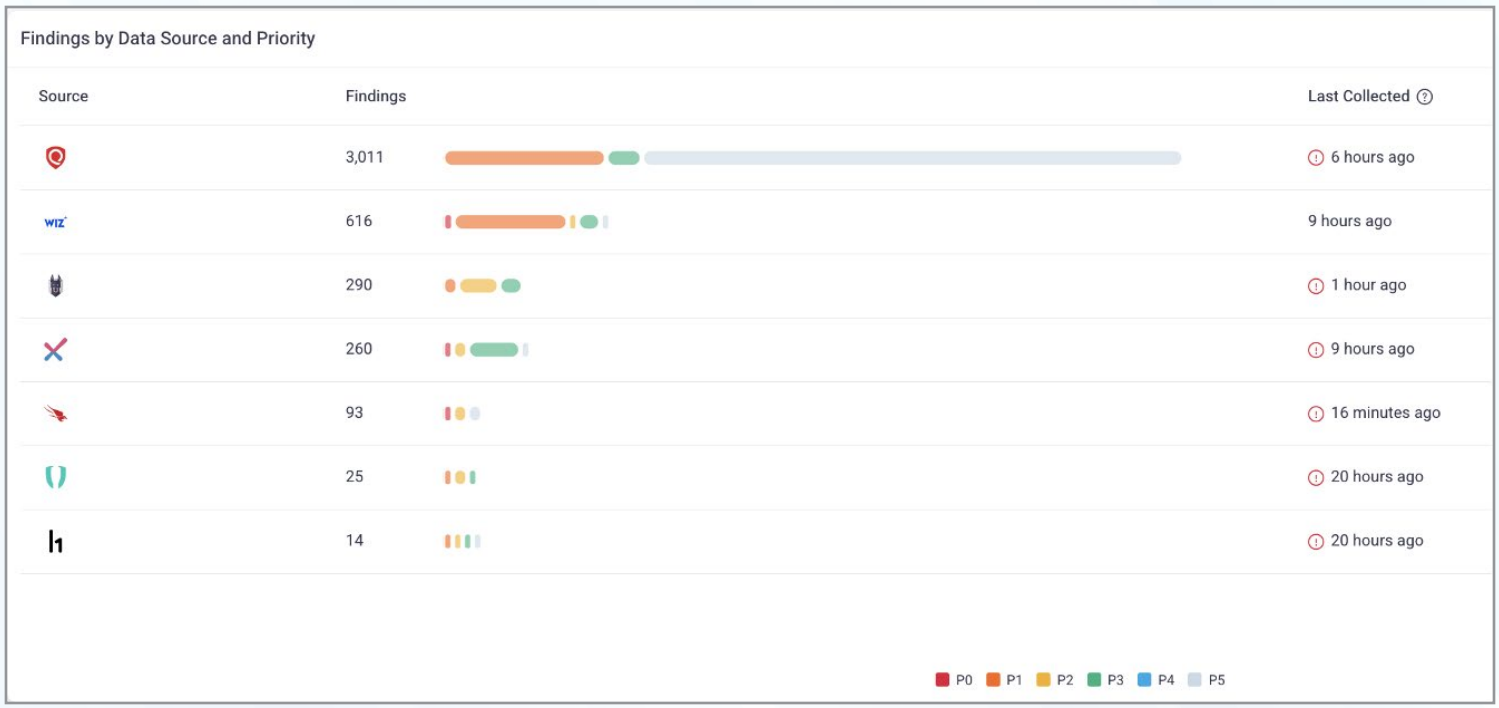

Figure 4. Seemplicity Data Sources

Once we got access to the platform and had the chance to explore a bit, the first feature we decided to look into was the ability to collect data from a wide variety of sources through the Seemplicity catalogue of connectors. In the *Dashboards* section, we saw a list of data sources that had been connected into the platform (including Qualys, Wiz, Snyk, CrowdStrike, and more), as shown in Figure 4.

When leveraging a platform that reconciles and streamlines vulnerability management, the number of integrated connectors is important. At the time of this review, Seemplicity had more than 150 available connectors, and the number is growing all the time. When the connectors begin feeding data into the platform, the events and findings are correlated and triaged to create a more meaningful and impactful range of findings to evaluate. In our test environment, we had roughly 16,000 findings that were analyzed and reduced to about 4,300 findings, thanks to deduplication or findings that shared a similar fix (see Figure 5). Only 47 of these are defined as P0s or deemed critical to the organization based on how they were defined. More on prioritization later in this paper.

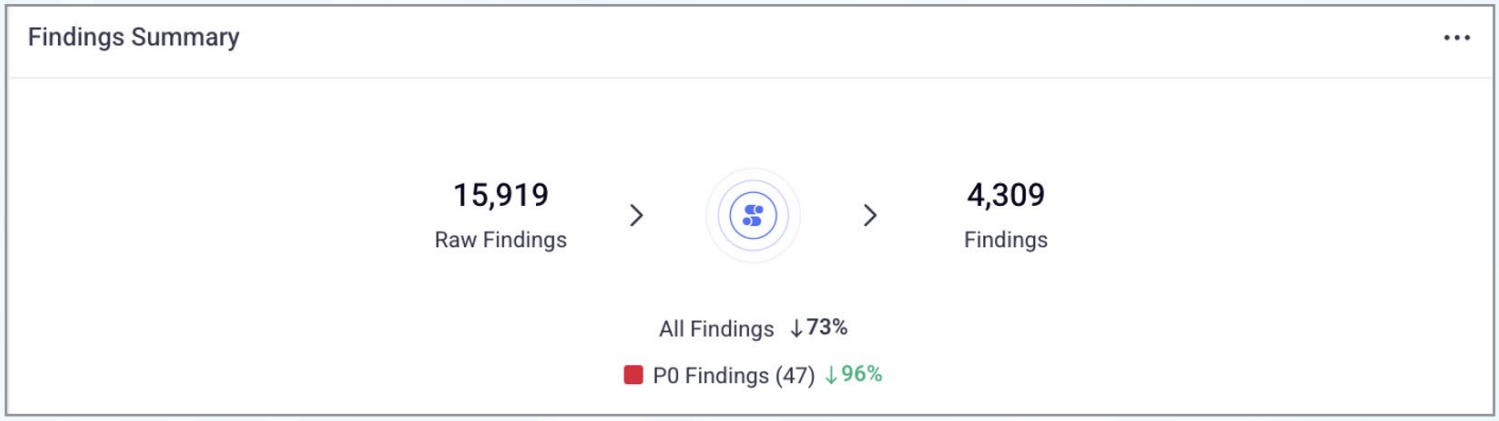

Figure 5. Seemplicity Data Triage

In some cases, these events are simply duplicates that need to be consolidated, whereas others have been aggregated into grouped findings because they share the same fix (and fixers) needed to resolve multiple vulnerabilities. Seemplicity can significantly help to normalize, deduplicate, and aggregate these events; this is critical to reduce alert fatigue and provide more intelligent datasets for analysis.

At the core, though, we still have to address the vulnerabilities noted in the platform. We went to dig into the findings by clicking into the *Findings* dashboard, as shown in Figure 6.

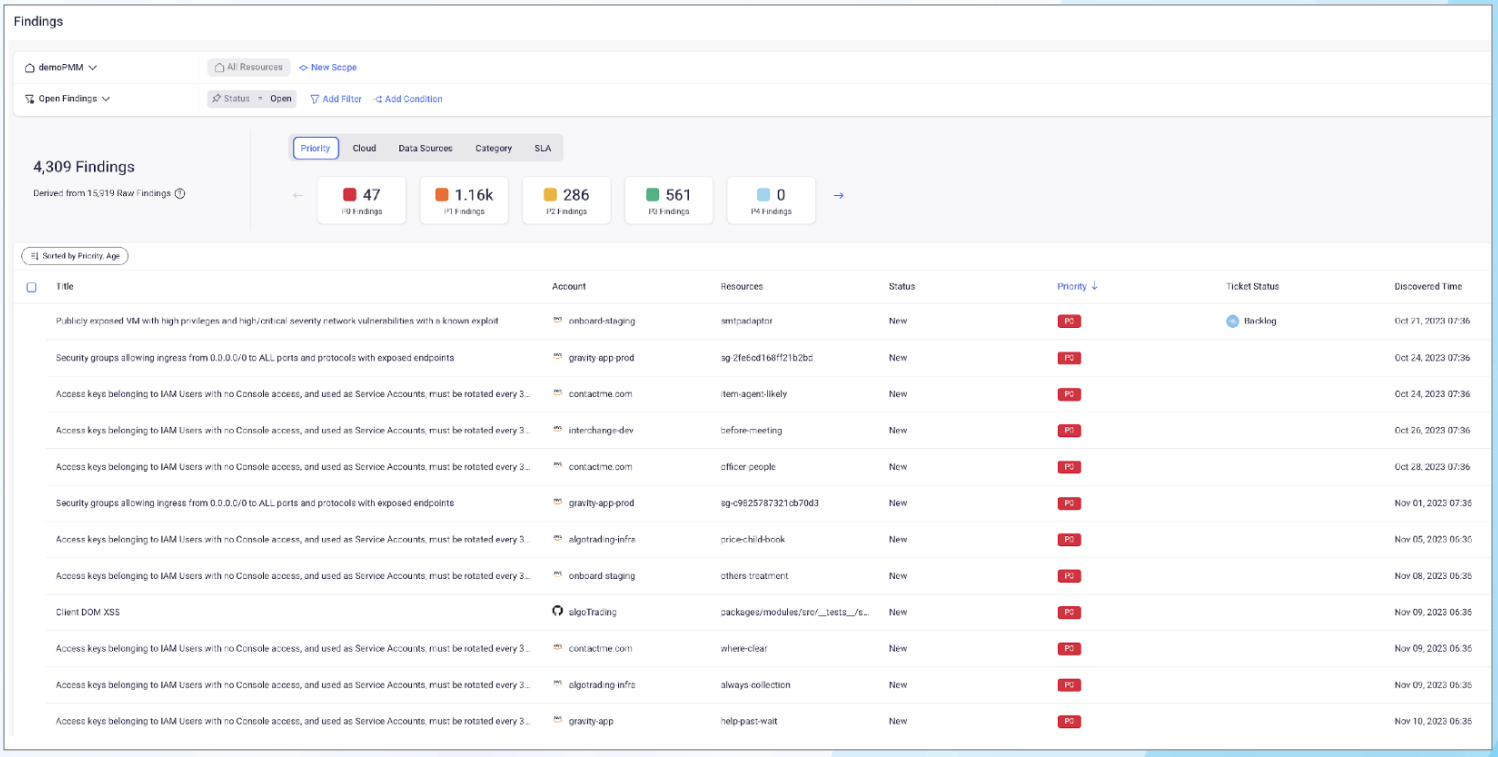

Figure 6. The Initial Findings Dashboard

We clicked into some of the findings for details about a vast array of data that included remediation suggestions and options to label the findings or generate tickets within Seemplicity or in integrated platforms such as ServiceNow. An excellent example of a finding that was aggregated into a single entry was a lack of MFA in Microsoft 365, shown in Figure 7. This view shows how Seemplicity aggregates data, with this particular finding linked to 30 underlying raw findings.

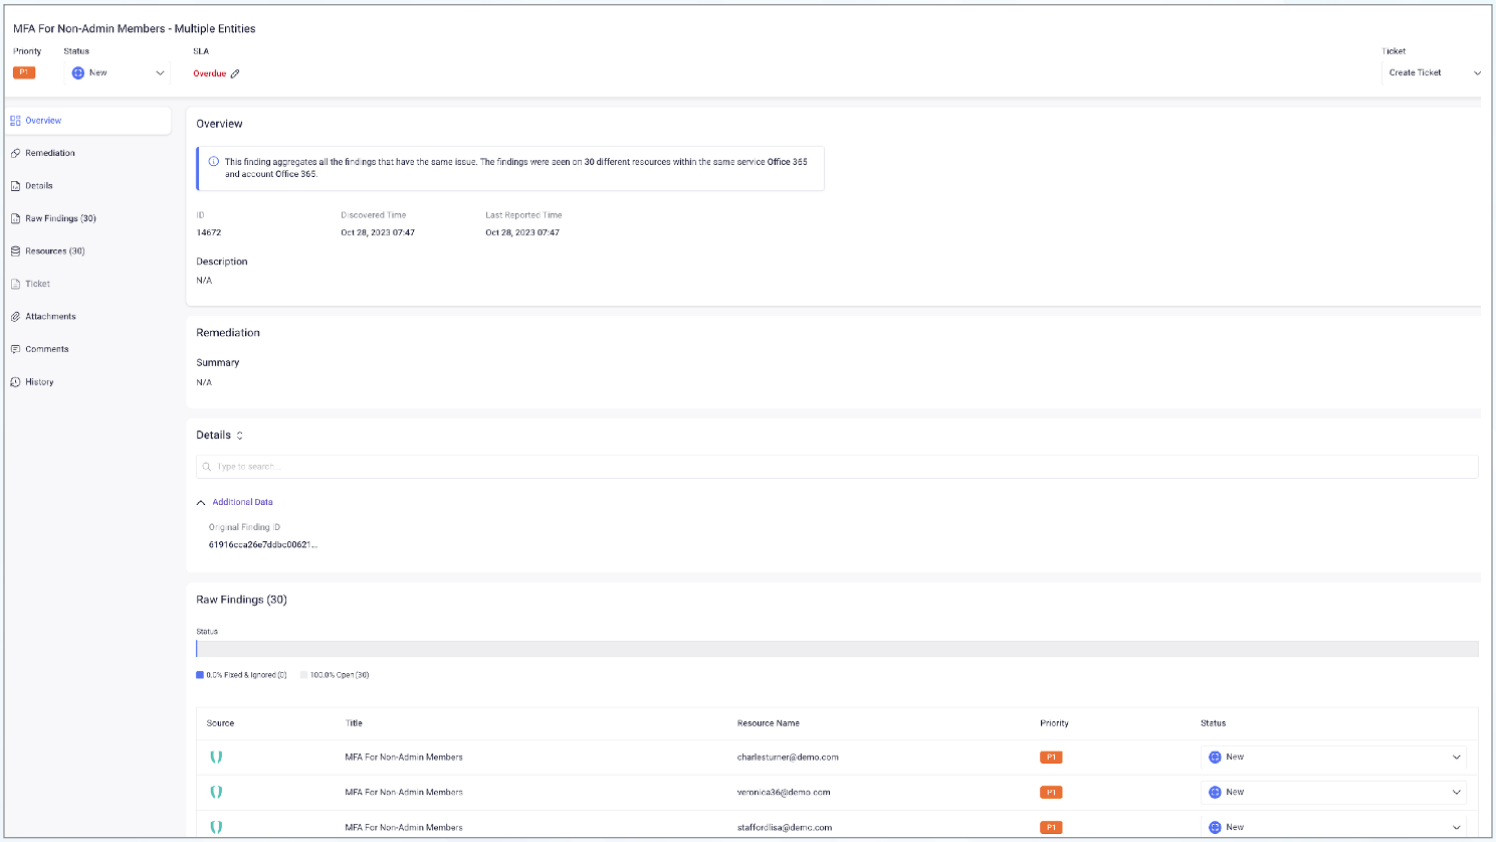

Figure 7. Finding Details in Seemplicity

For all findings, we can see the finding details' remediation options, including the capability to generate a new remediation plan with Seemplicity AI when and if needed; details on resources affected; and any case-relevant information, such as comments, case history, or even attachments related to the issue like network traffic captures, command line history, and more. Another example with detailed remediation guidance is shown in Figure 8.

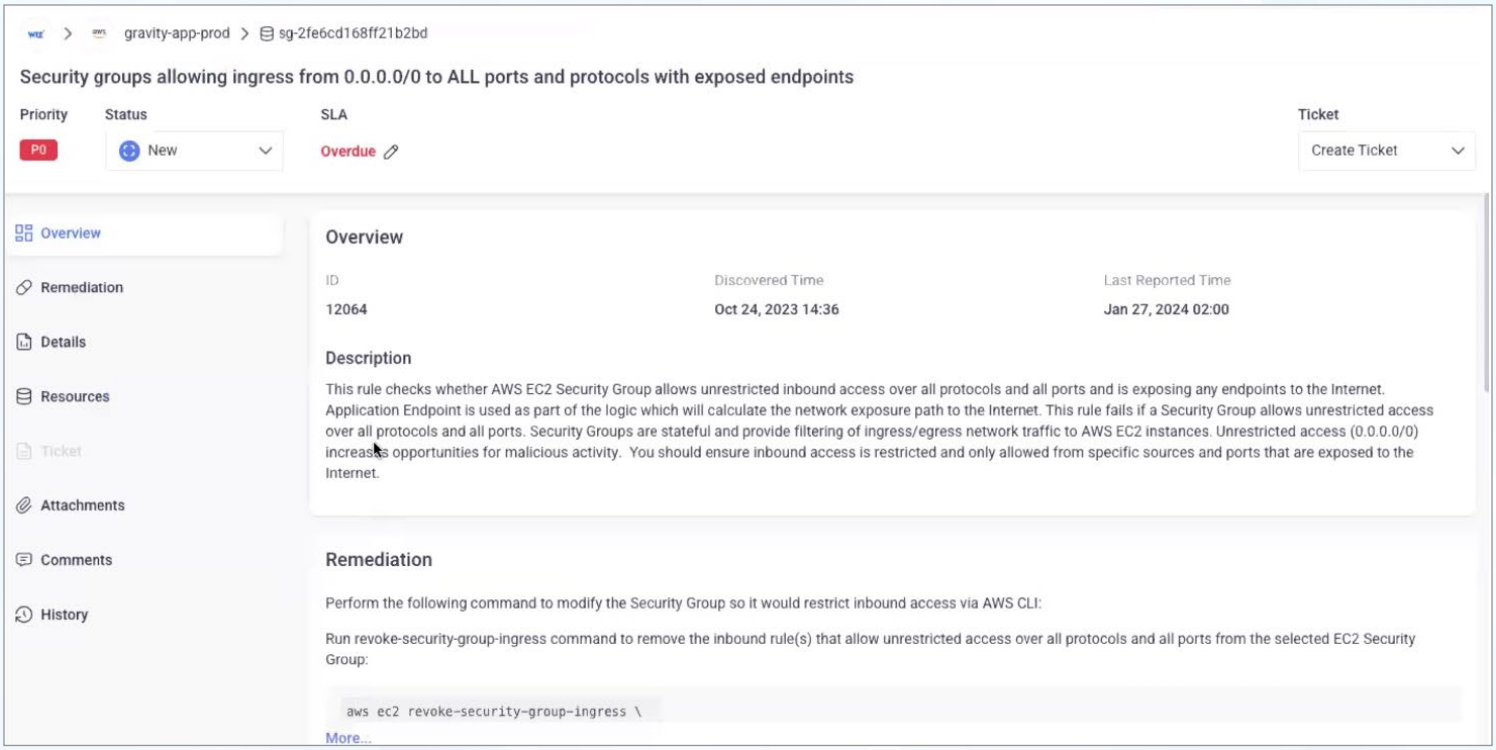

Figure 8. Remediation Details for a Finding

We also found it simple to modify the status of any given finding to New, Reviewed, Exception, False Positive, Inactive, Resolved, or more, as shown in Figure 9.

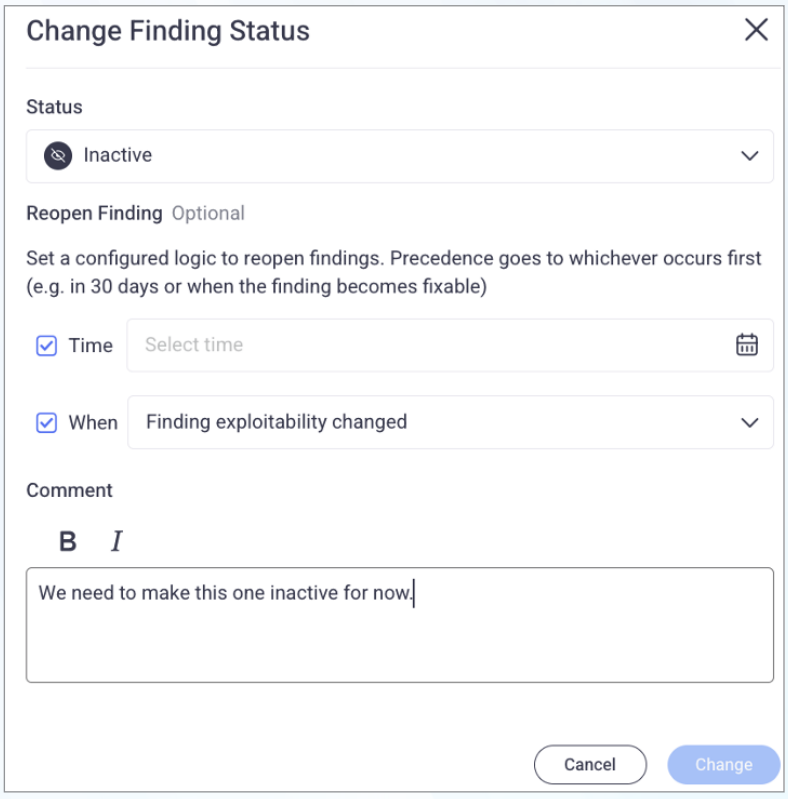

Figure 9. Updating Finding Status

Tickets can be sent to Jira and ServiceNow, as well, if organizations need to integrate Seemplicity into an existing IT service management (ITSM) platform and workflow, as seen in Figure 10. These bidirectional integrations enable teams to collaborate seamlessly while continuing to operate within their existing processes.

In scenarios where you have a penetration test finding, or some other type of review or assessment that doesn't inherently offer automated event ingestion, Seemplicity makes it easy to add a manual finding or import a list of findings. An example of adding a *Manual Finding* is shown in Figure 11.

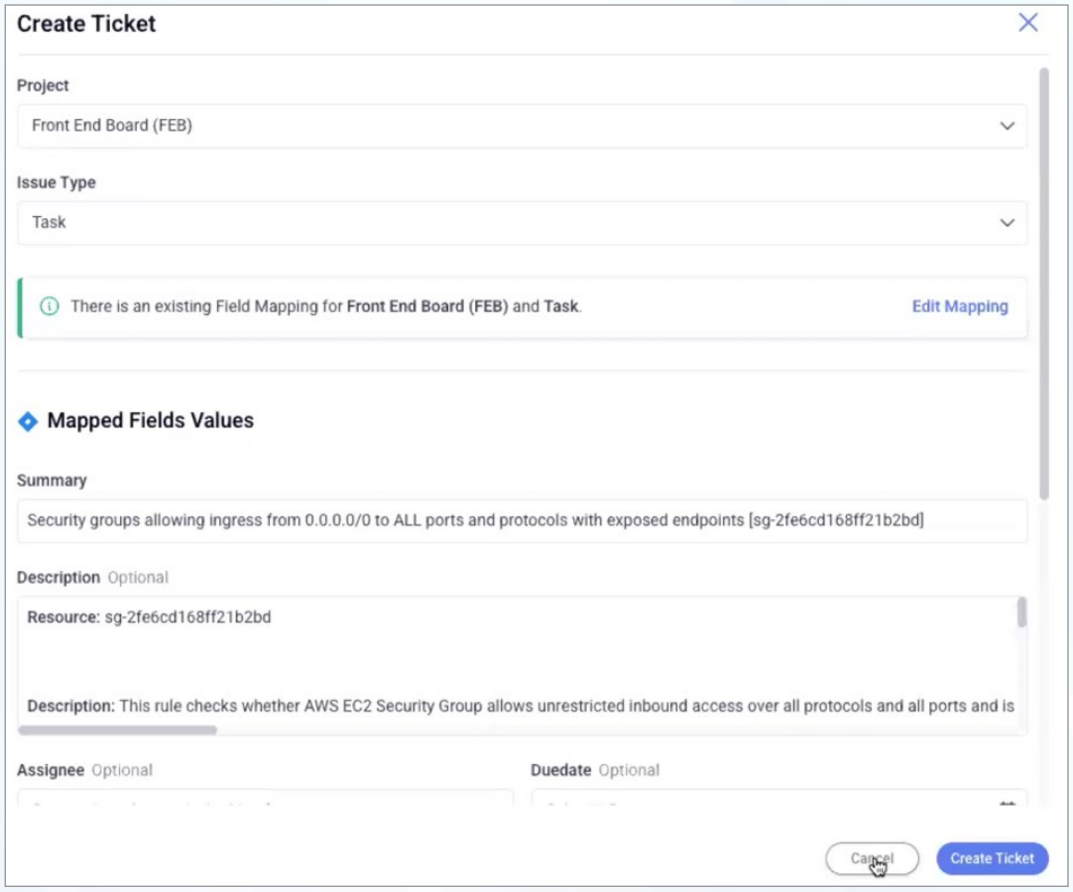

Figure 10. Generating a Finding Ticket from Seemplicity

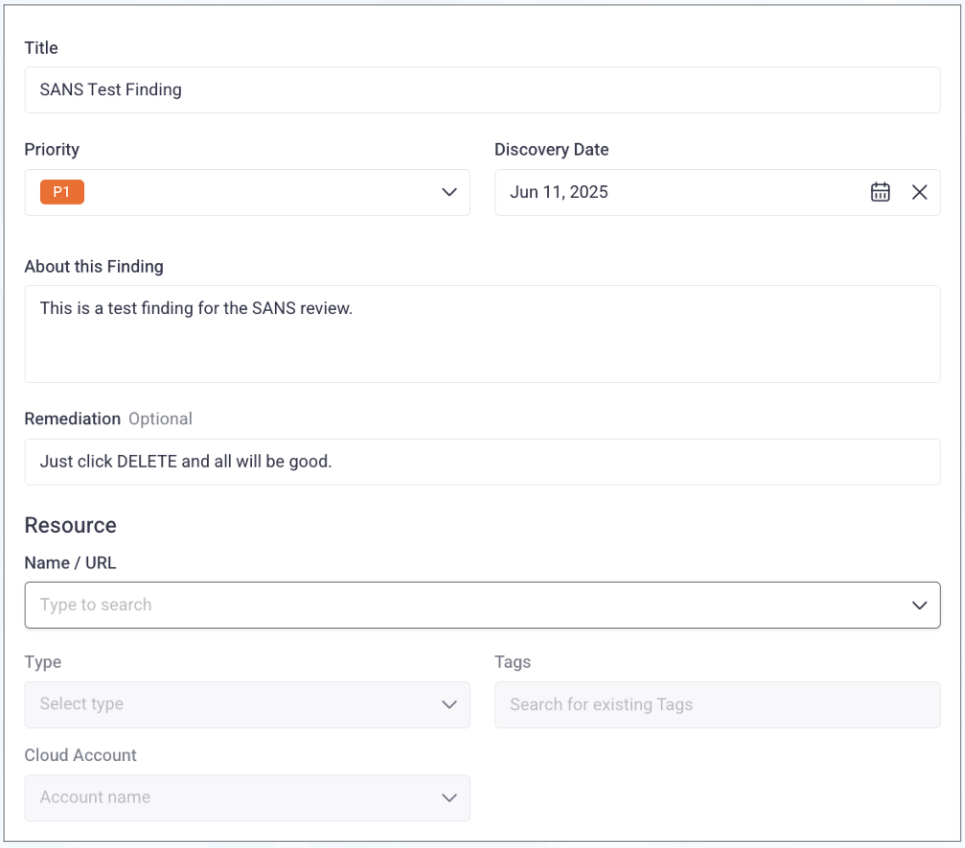

Figure 11. Adding a Manual Finding

We dug into some of the findings to identify where Seemplicity was correlating data between sources and adding insight and intelligence. One example, some vulnerabilities in components within an AWS EC2 instance, showed vulnerability data from both CrowdStrike and Wiz, but Wiz highlighted the level of internet exposure the asset had, where as CrowdStrike did not (see Figure 12).

The Seemplicity platform also aggregated 363 total vulnerabilities associated with this issue into a single finding, with an aggregated view of what to do in terms of remediation (see Figure 13).

This demonstrates how the platform takes common attributes of assets and events and helps to consolidate and correlate findings into more unified remediation action items, which should help to reduce the "noise" in both IT and security operations.

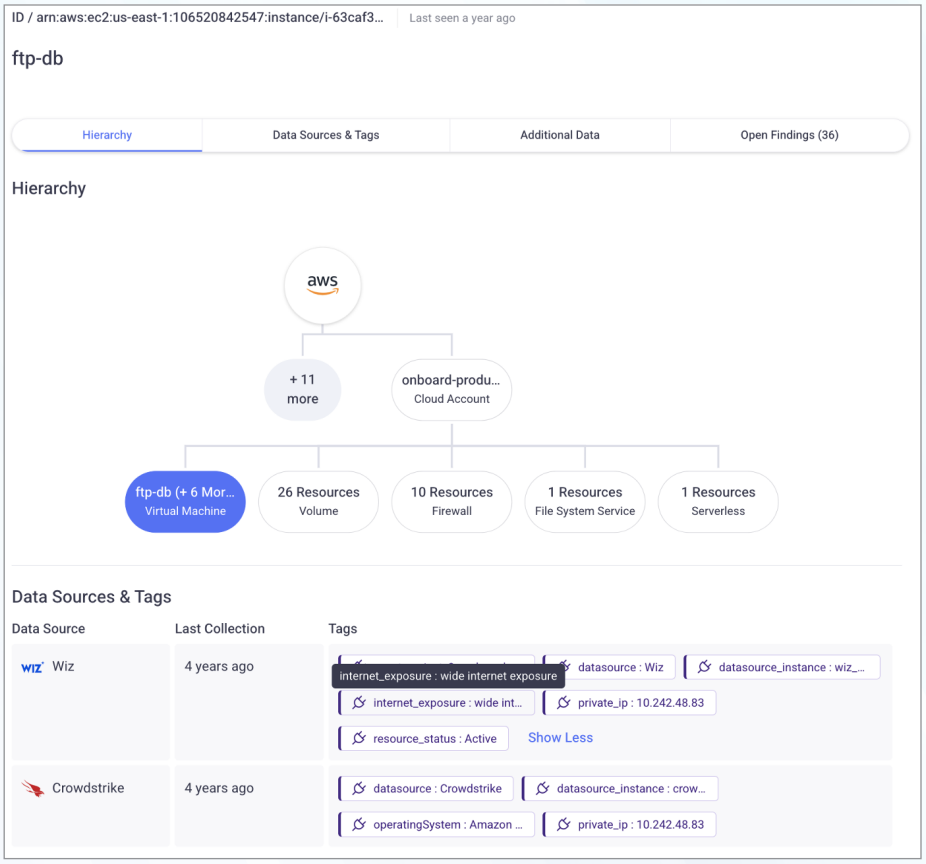

Figure 12. Event Source Correlation

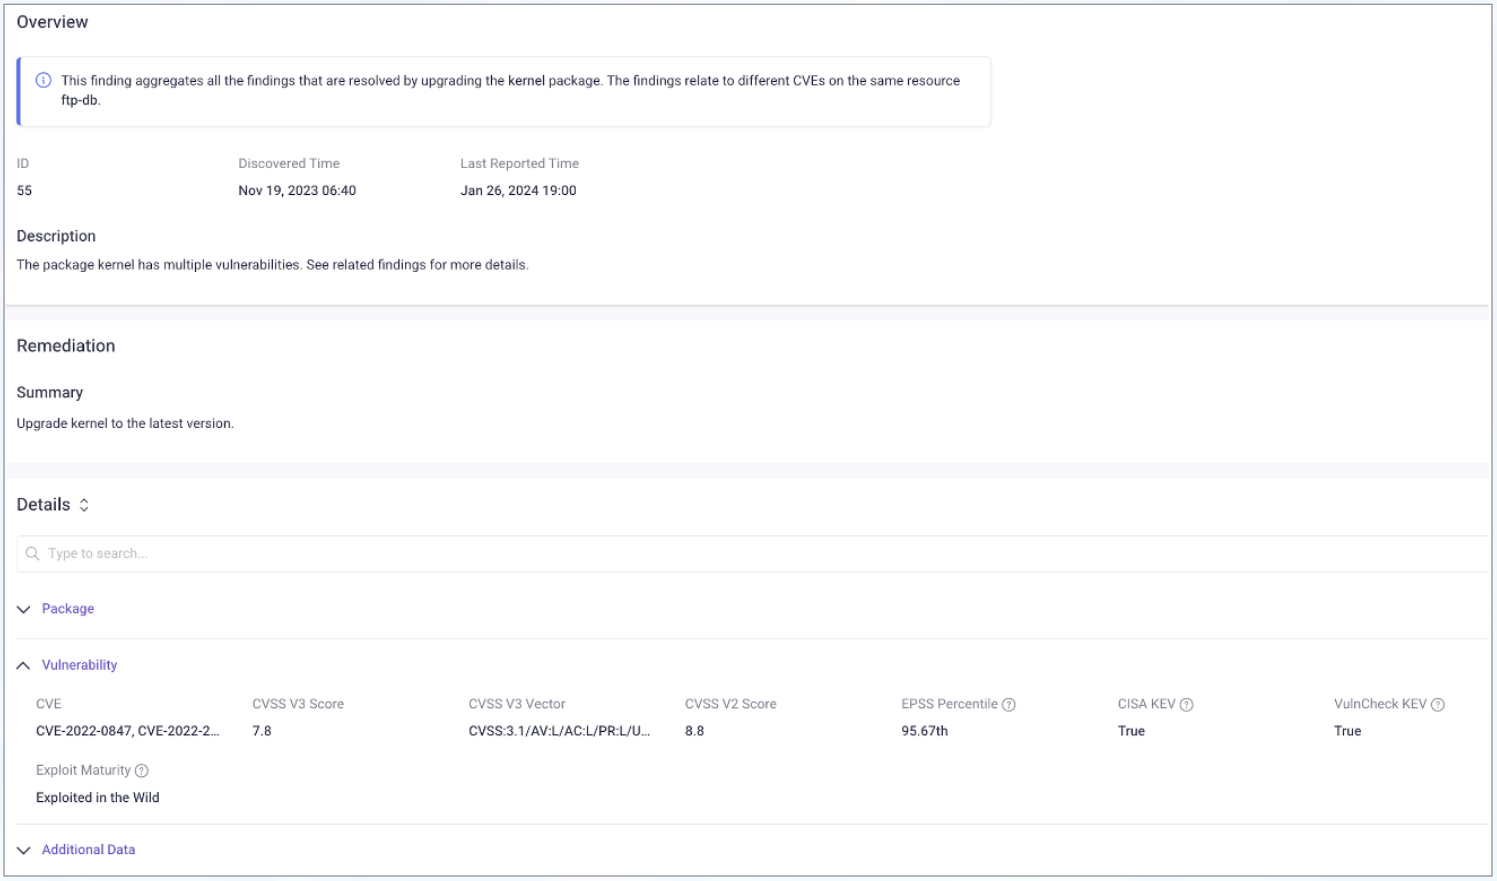

Figure 13. Vulnerability Aggregation in Seemplicity

Policies and Priorities in Seemplicity

A consistent vulnerability management challenge for security and operations teams is prioritizing vulnerabilities in their environments, particularly given the conflicting "scores" that are generated by vendors. Seemplicity allows you to create your own prioritization schemes that can be applied to specific asset scopes and use unique filters. For example, as shown in Figure 14, we have a priority rule called "External CISA KEV" that applies to all externally visible resources with a filter of "known exploitable vulnerability." Because these are likely the most exposed and vulnerable assets in the environment, we have a priority of P0, the highest available.

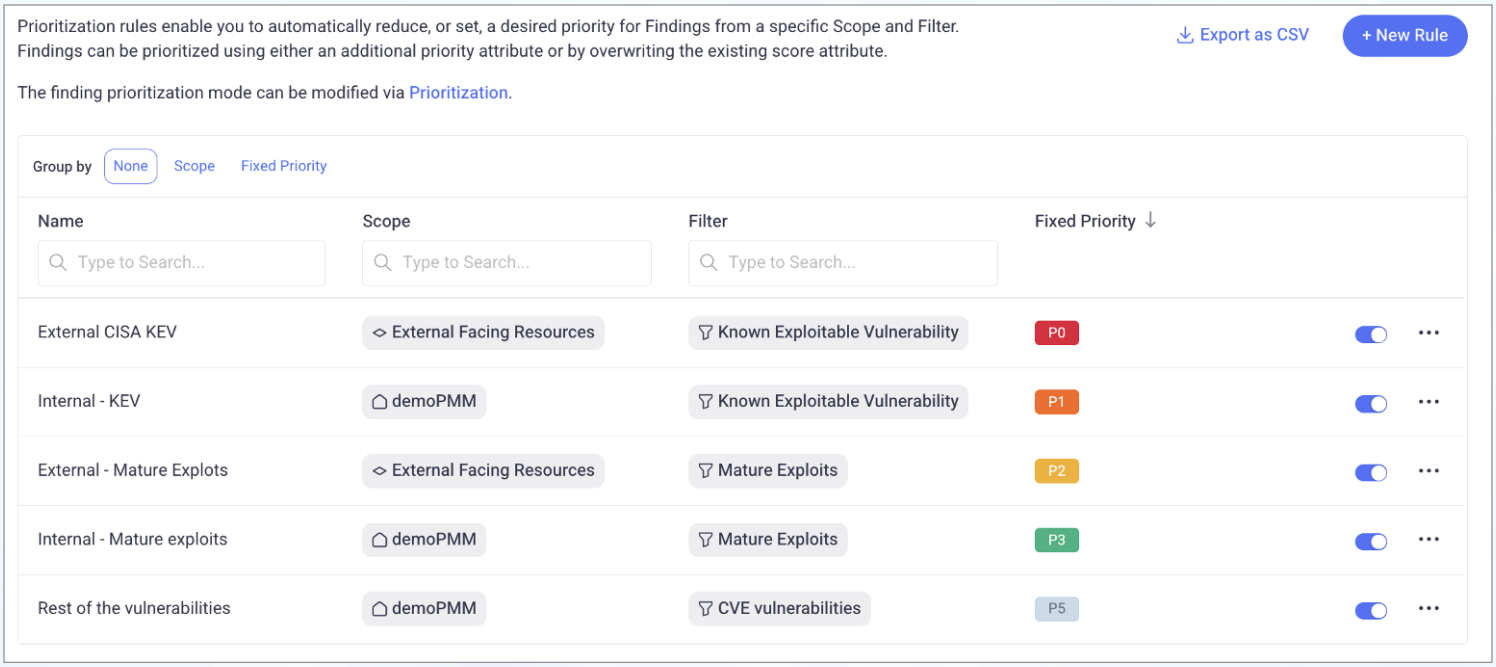

Figure 14. Priority Rules in Seemplicity

It was also easy to assign SLAs for remediation based on the scope and filters, allowing a highly flexible way to rank tickets in a remediation queue. Figure 15 shows an example rule we created.

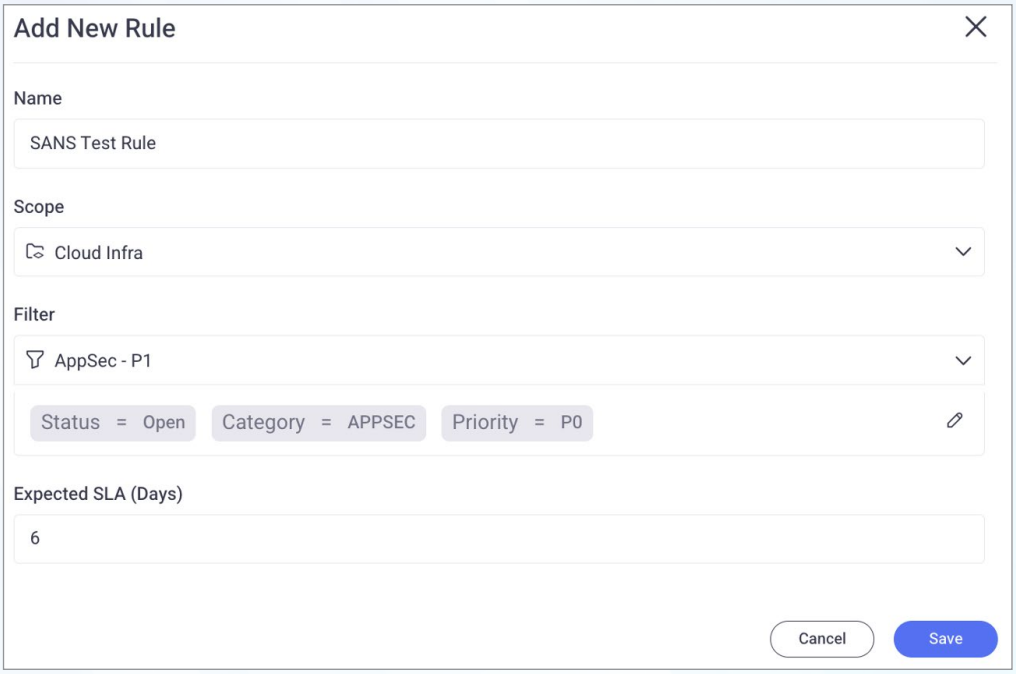

Figure 15. Creating a Remediation SLA Rule

The Seemplicity team had a number of preconfigured examples of remediation SLAs, as shown in Figure 16.

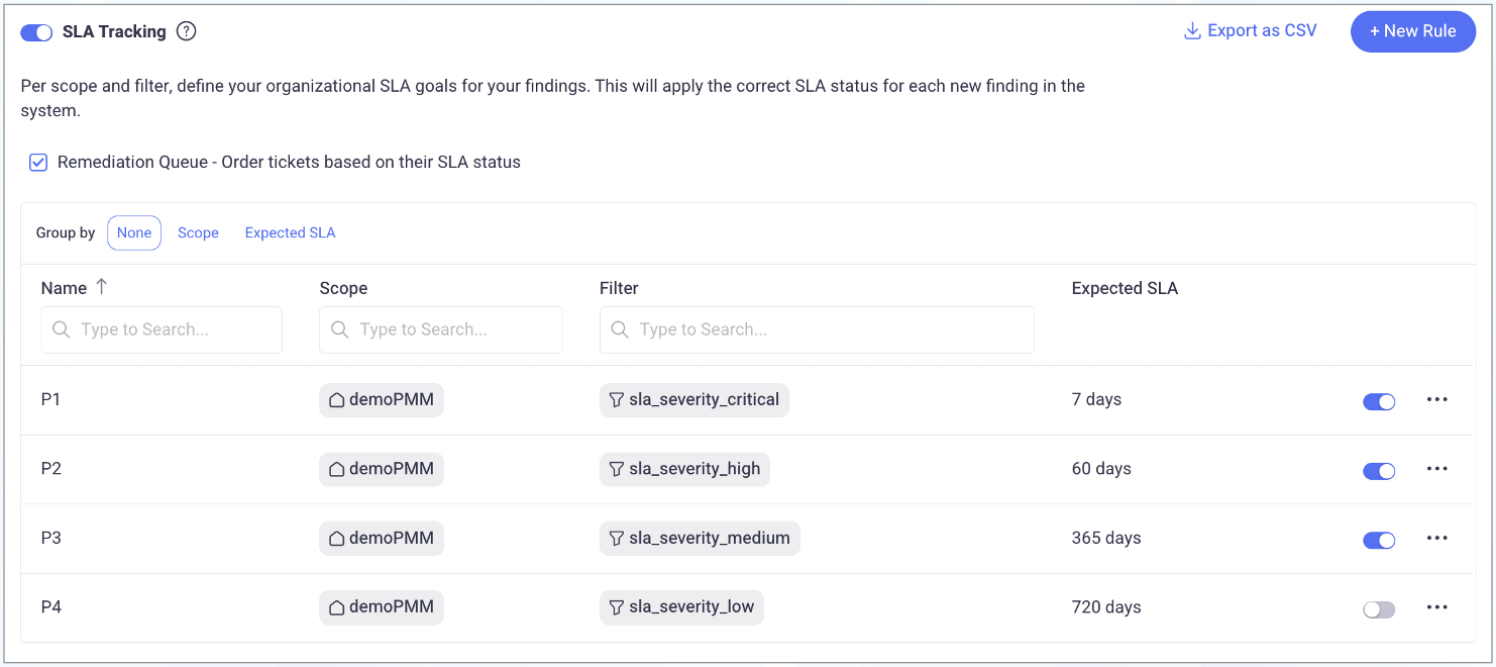

Figure 16. Severity-Based SLAs

In this set of examples, you can easily see that the SLAs for remediation are based on filters oriented toward vulnerability/posture severity, with a critical issue or risk mandating a seven-day remediation cycle, and an asset with low severity issues leaving roughly two years for remediation. (These are just examples, but they illustrate the flexible nature of how ticket assignments can be created.) You also can track SLA adherence through the *SLA Compliance* dashboard, as shown in Figure 17.

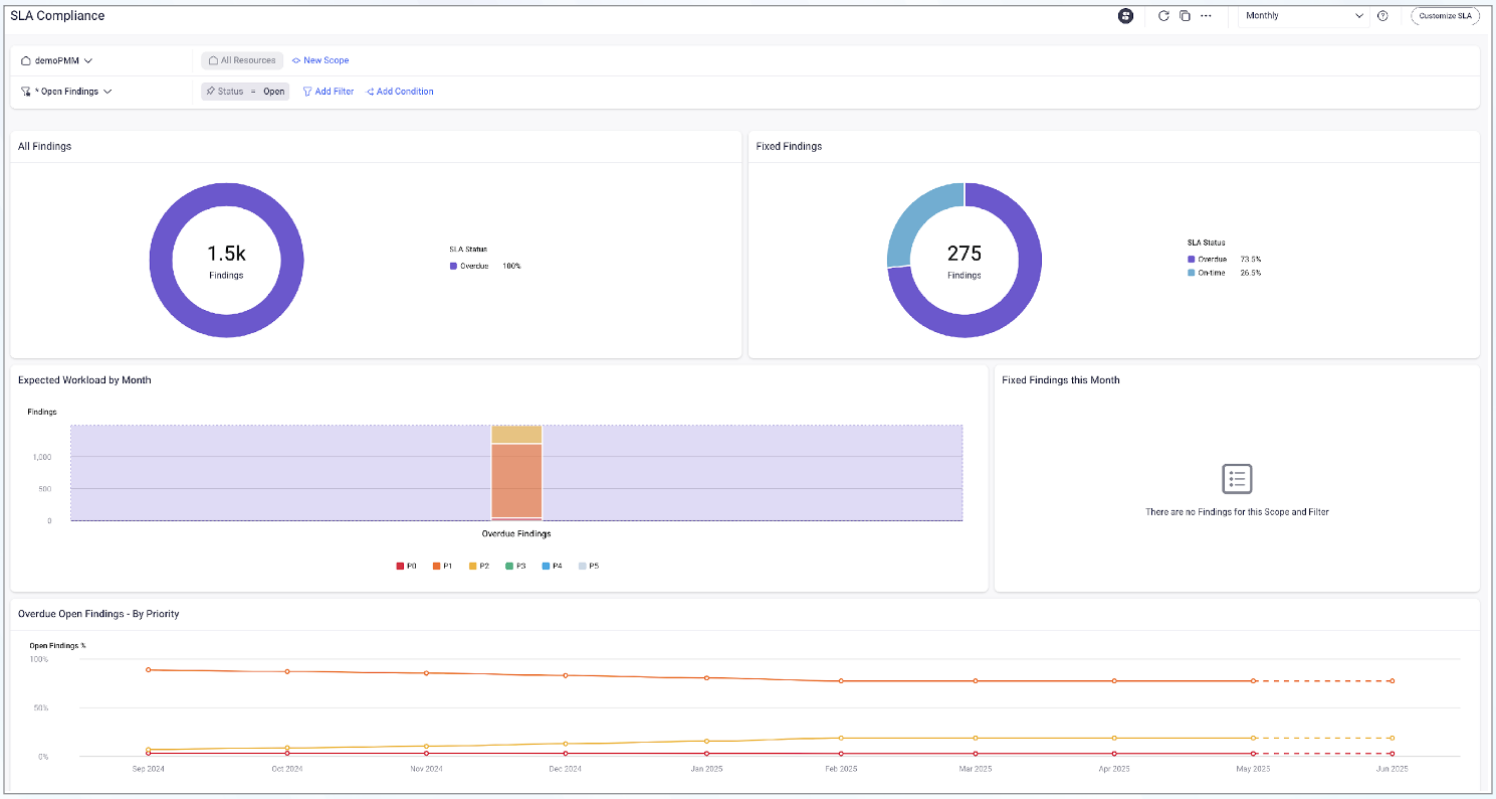

Figure 17. SLA Compliance Dashboard

This dashboard is great as a driver for both operational and security-oriented metrics, allowing you to track how well teams are performing in remediating issues labeled with specific priorities. You can easily see current issues that are being handled on time, those that are late, and longer-term metrics on overdue findings that have not yet been remediated.

When we have a specific application or scope that we need to dig into, we can look in the *Automations* category (for remediation information) to quickly see what open findings and open tickets are listed. In our test environment, we looked into an application called *Customer Portal* in our production environment. In the main dashboard, we found a number of open remediation items that led us to the Automations tab, where we selected the application environment associated with the *Front-End Team* and found there were 31 open findings (after consolidation), with 12 P0 and P1 tickets in the backlog and 2 in progress (see Figure 18).

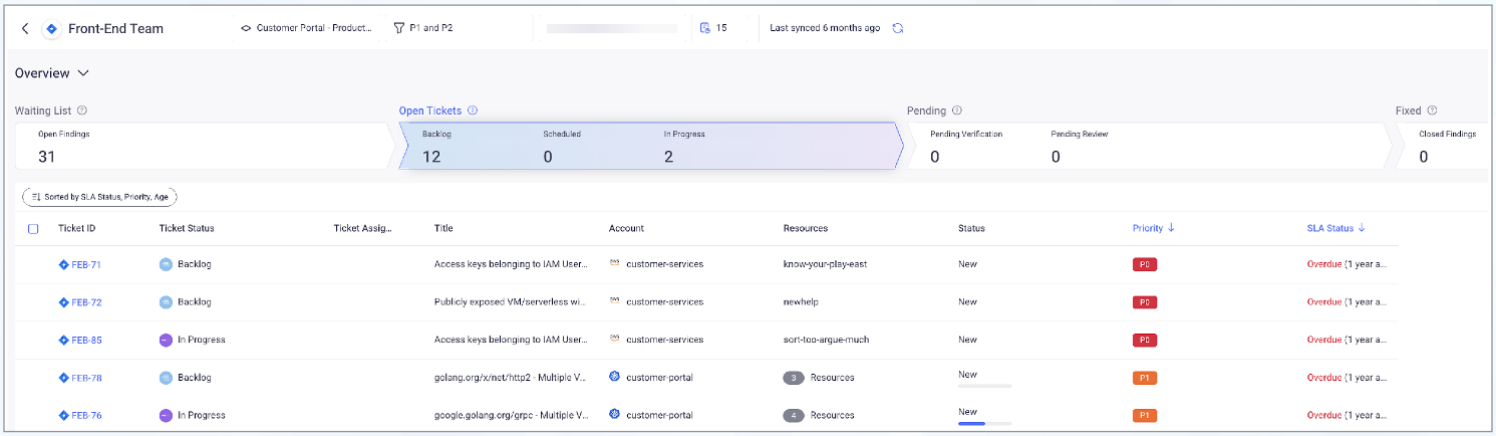

Figure 18. Customer Portal App Automation Reporting

There are a lot of important details in this view. First, we can see the total queue size of 15, which is the aggregated set of tickets and issues associated with this set of vulnerabilities. This is important, because it demonstrates a significant reduction in actual remediation items compared to the overall list of findings (31). Second, we see how Seemplicity has integrated into existing ticketing systems with unique ticket IDs, which helps track these issues across workflows. Finally, we can see at a glance what the current status of any remediation item is—many in backlog and some in progress, in this example. This is the kind of operational dashboard that can help numerous teams track and ultimately close vulnerabilities discovered in large, complex environments. In addition, it's possible to customize the queue size for individual teams so they're not overwhelmed with tickets when issues are discovered. This setting can be overridden for high-severity vulnerabilities, if desired. The ability to granularly, at a team-by-team level, customize the ticketing queues is important for improving collaboration among the IT and development teams and security operations.

Another key feature of the Seemplicity platform is the ability to customize ticket details. You can customize ticket templates with resource information, tagging labels, due dates, and many more fields, as shown in Figure 19.

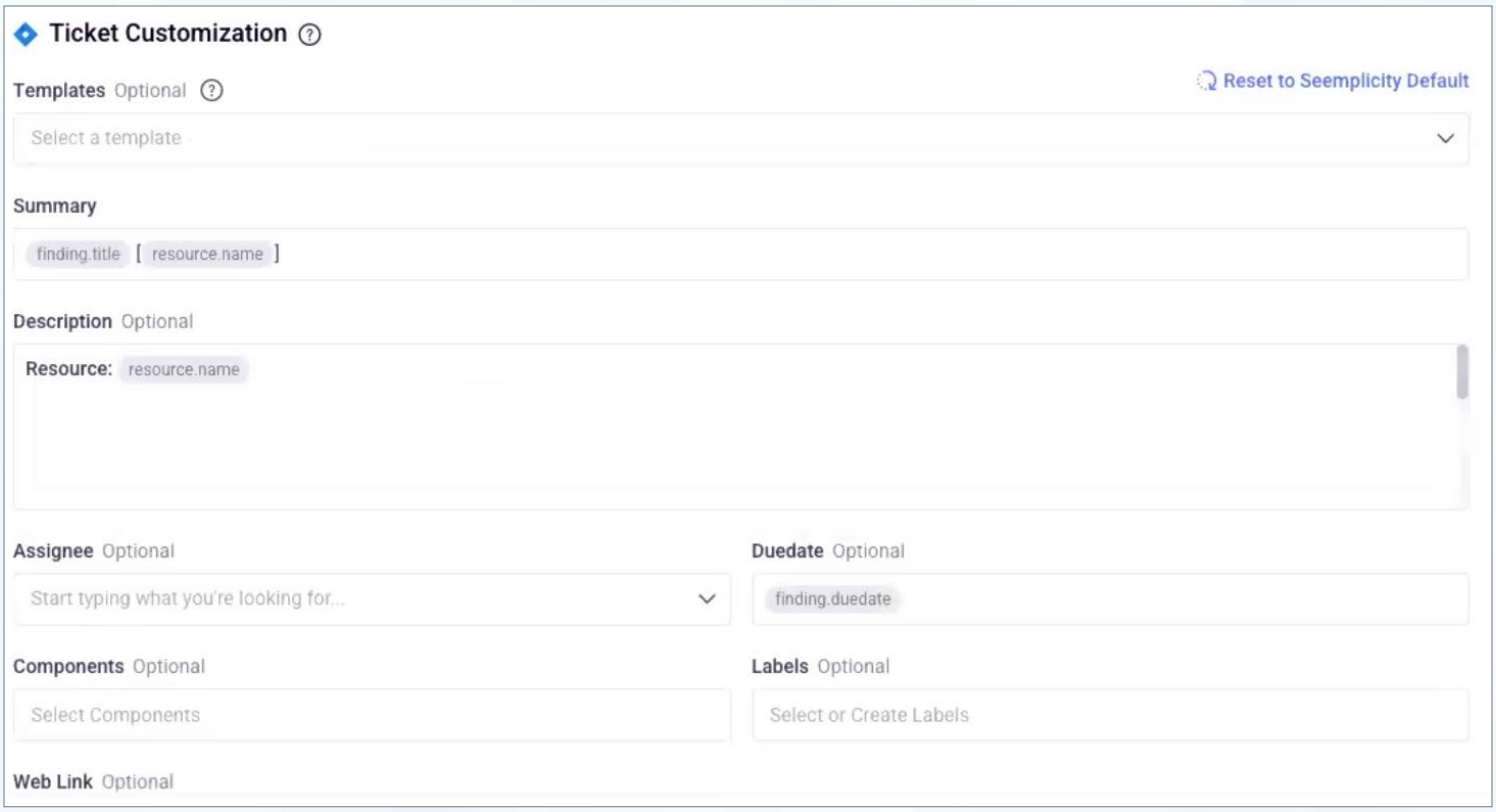

Figure 19. Customizing a Ticket for Remediation

It's important to note that you can create a variety of "remediation queues," which allow you to send tickets and manage findings and automation items for remediation to specific groups and asset owners.

Remediation Simplification and Prioritization

Within the Seemplicity platform, there is a significant focus on automation. Remediation opportunities can be presented visually to show what the scope of the remediation scenario is, what teams and roles are associated, and any filters that apply (such as P1 and P2 vulnerabilities), along with successful or failed criteria to remediate (as shown in Figure 20 for some application security findings).

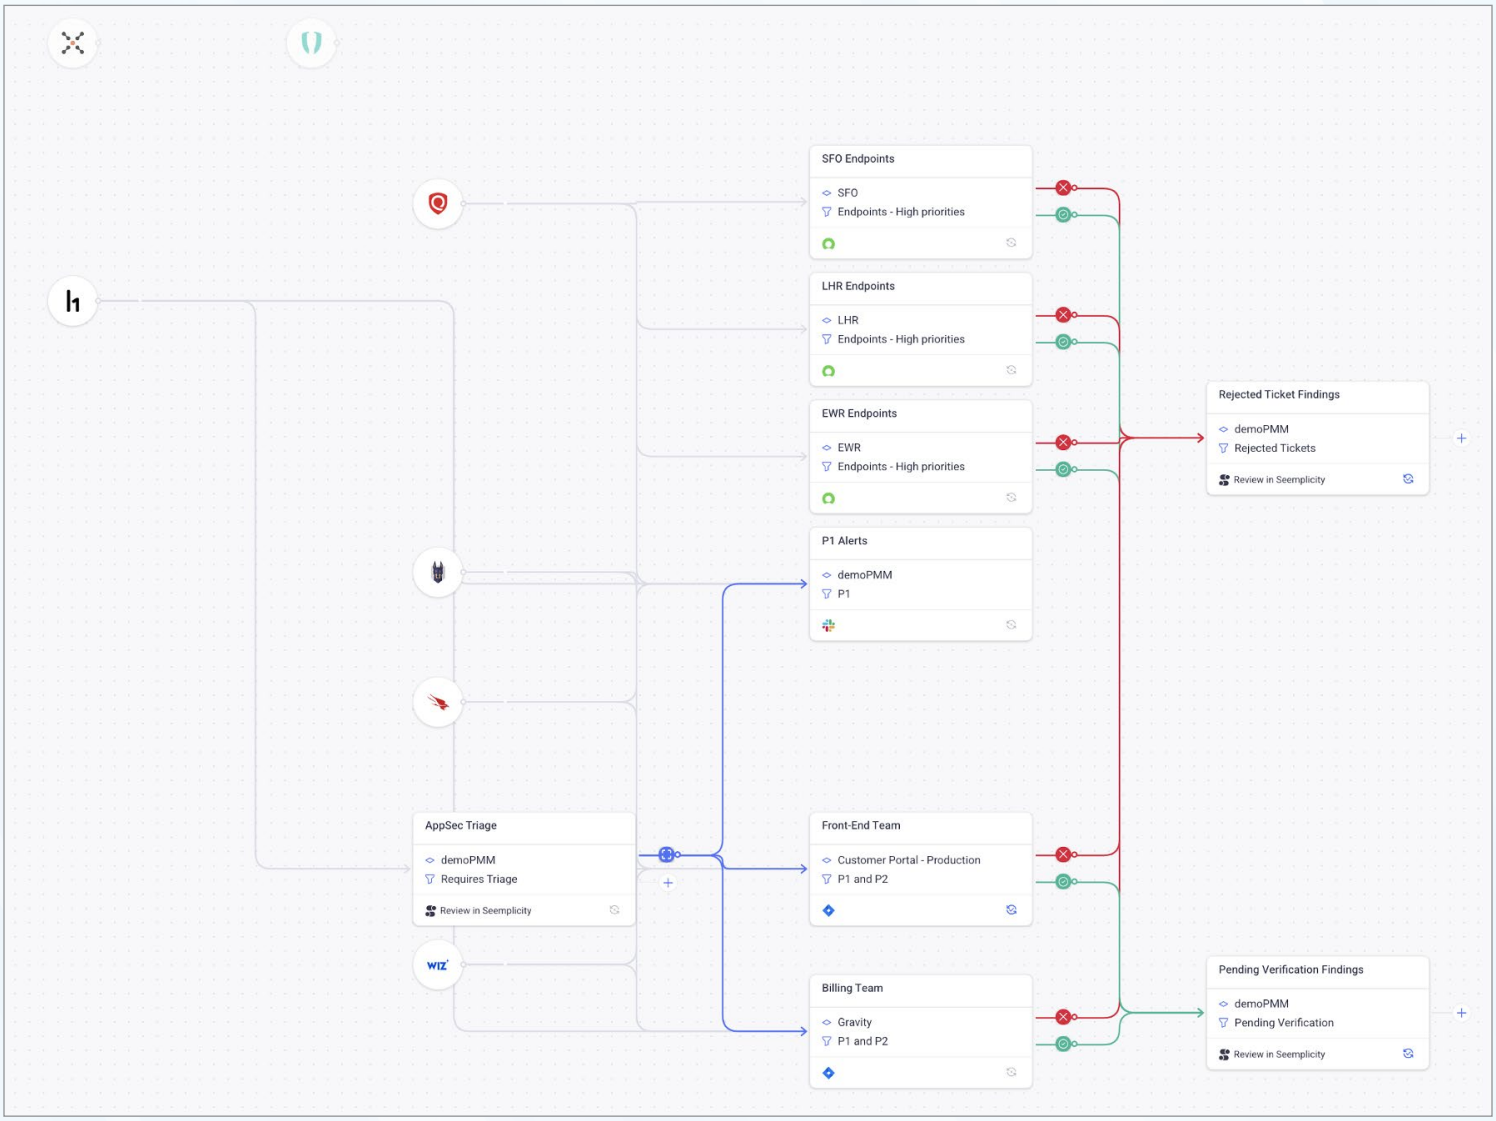

Figure 20. An Automation Workflow for Remediation

In this scenario, we can highlight a major input vector to find out where high-priority alerts are focused. Here, we highlight the vulnerability data from Qualys to show that several different locations (San Francisco, London, and Newark) need some immediate attention on high-priority alerts (see Figure 21).

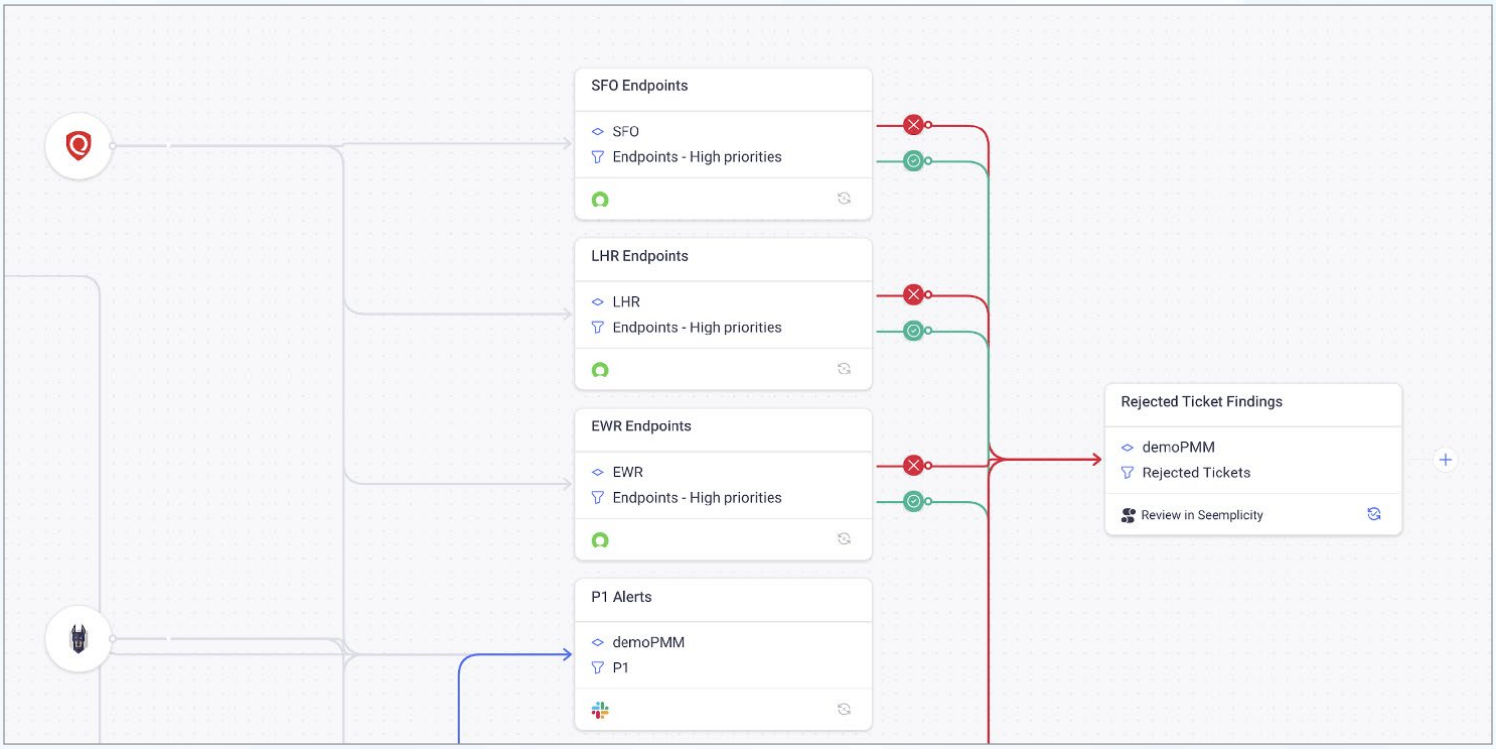

Figure 21. Escalation Paths for Qualys Findings

What really hits home in this scenario, however, is the additional range of different events and alerting that is occurring in a relatively small organization. There are a wide range of cloud events and alerts coming from tools like Snyk and Wiz, workload events reported by CrowdStrike, and even bug bounty issues being reported via services like Hacker1 (see Figure 22).

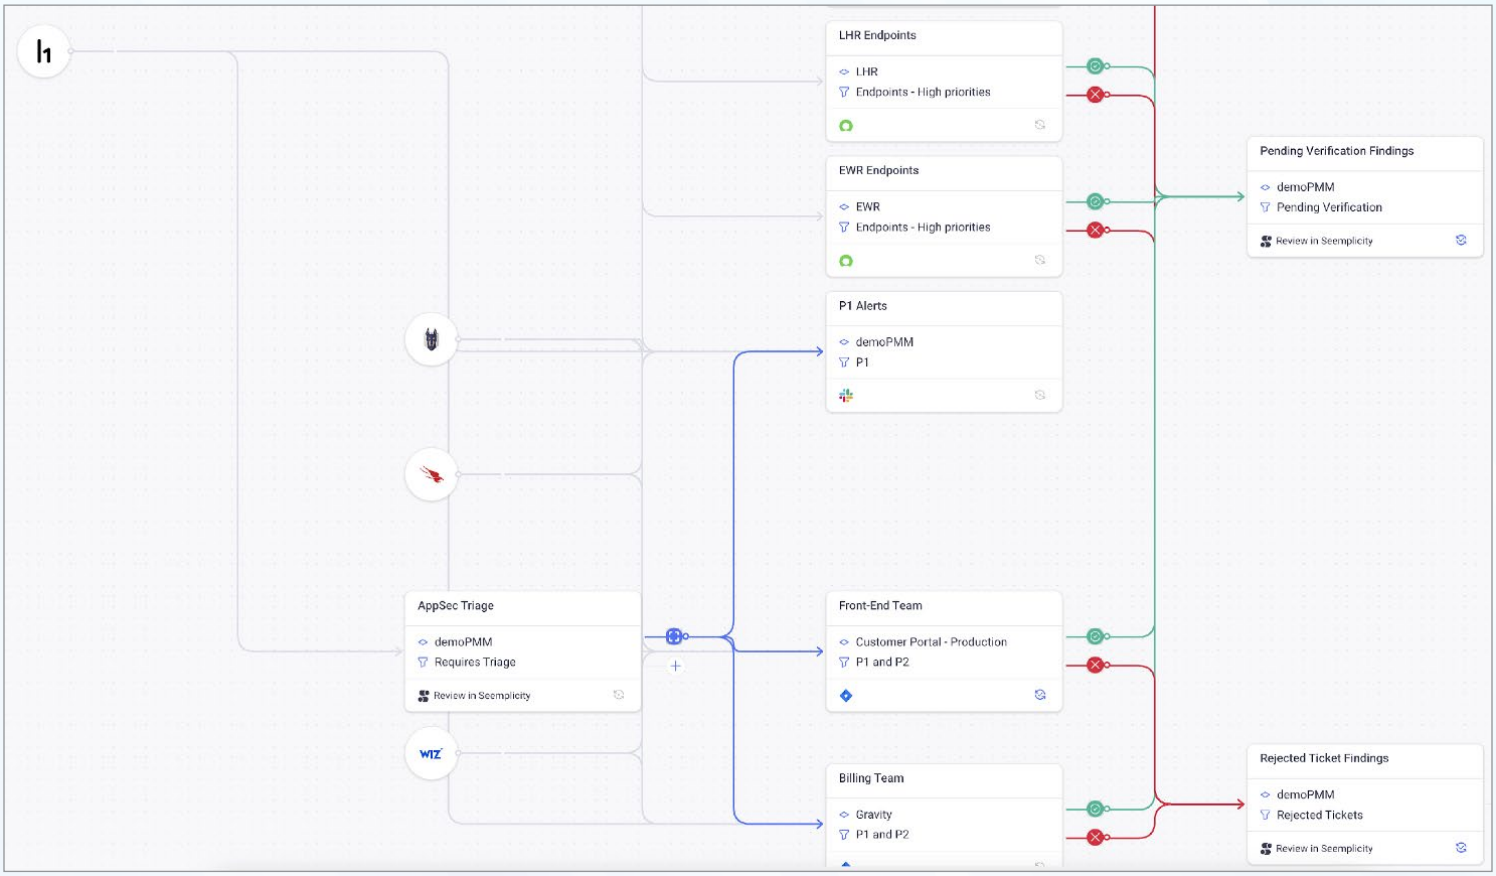

Figure 22. Additional Findings in the Seemplicity Workflow

In these workflows, teams assigned tickets can mark them as accepted or rejected, which can lead to a number of secondary workflows. Some tickets may need to have findings verified, such as when a vulnerability scanner has reported something and the security team needs to validate the alert. Others may be rejected outright, requiring an exception. In these cases, you can create new tickets, send dashboard reports to stakeholders, and/or send finding data to stakeholders (see Figure 23).

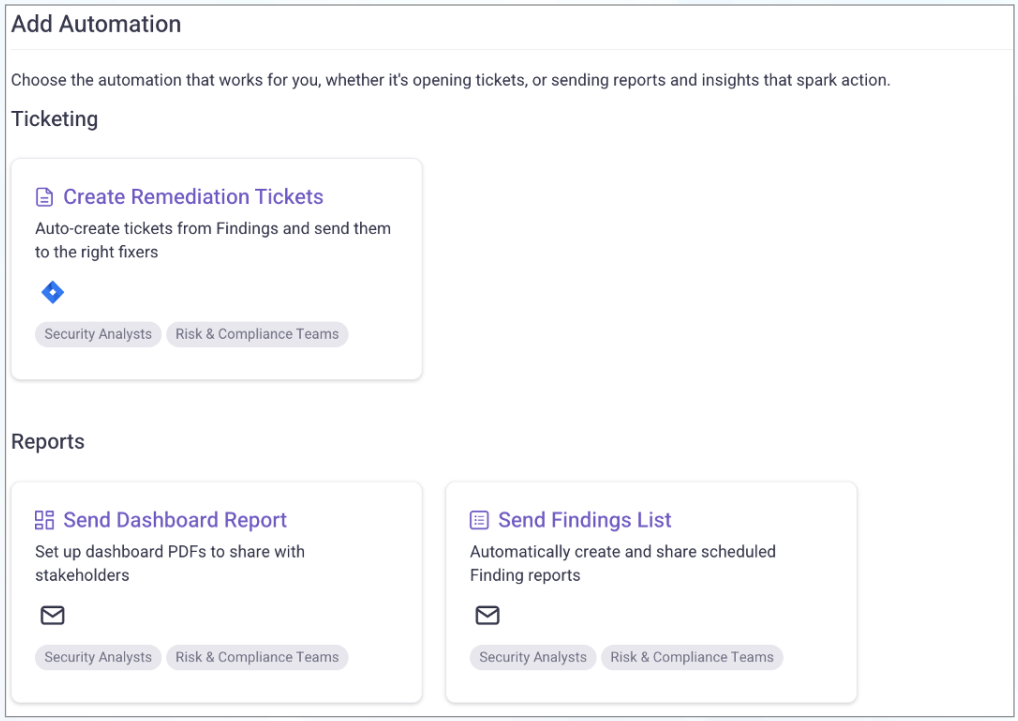

Figure 23. Exceptions Automations

As with any vulnerability management and remediation orchestration solution, Seemplicity offers a flexible reporting engine that is highly customizable. Teams can easily create new dashboards, and a variety of default reports are also available. Reports can be automatically distributed via email to stakeholders outside the Seemplicity platform, ensuring they still have access to key data and insights.

AI in the Seemplicity Platform

Seemplicity's recent platform enhancements introduce a meaningful expansion of its AI-driven capabilities, focused on improving how organizations operationalize exposure management data. As described during the product walkthrough by co-founder Ravid Circus, these updates are less about introducing standalone AI features and more about embedding AI across key workflow friction points—from data analysis and prioritization to remediation and executive reporting.

At the center of this evolution is Seema, an AI-powered exposure management assistant that goes well beyond a traditional chatbot interface. Seema is tightly integrated with Seemplicity's normalized data model and is capable of translating natural language queries into structured queries against the platform's underlying exposure data. In practical terms, this allows users to ask high-level questions—such as whether their environment is affected by a specific vulnerability like Log4j—and receive precise, contextualized answers. These responses include affected assets, associated vulnerabilities, and prioritized insights, eliminating the need for manual querying or dashboard navigation (see Figure 24 on the next page).

Importantly, Seema's capabilities extend beyond simple query-and-response interactions. The platform incorporates AI reasoning and analytical capabilities that enable users to derive actionable insights from their data. For example, Seema can analyze trends in vulnerability categories (such as common weakness enumerations [CWEs]) and map them to recommended developer training programs, effectively bridging the gap between security findings and remediation strategy. This represents a shift from passive reporting to active operational guidance, where AI helps translate raw security data into concrete next steps.

Another major enhancement is the use of AI to aggregate vulnerabilities and simplify remediation workflows. Rather than presenting large volumes of individual findings, Seemplicity now groups related vulnerabilities (such as multiple CVEs affecting a single system component) and identifies a common remediation path. For example, multiple kernel-related vulnerabilities may be consolidated into a single recommendation to upgrade to a specific version. This significantly reduces noise and enables security and IT teams to focus on remediation actions rather than triaging individual findings. Figure 25 shows this aggregation of 173 raw findings into a single set of unified recommendations.

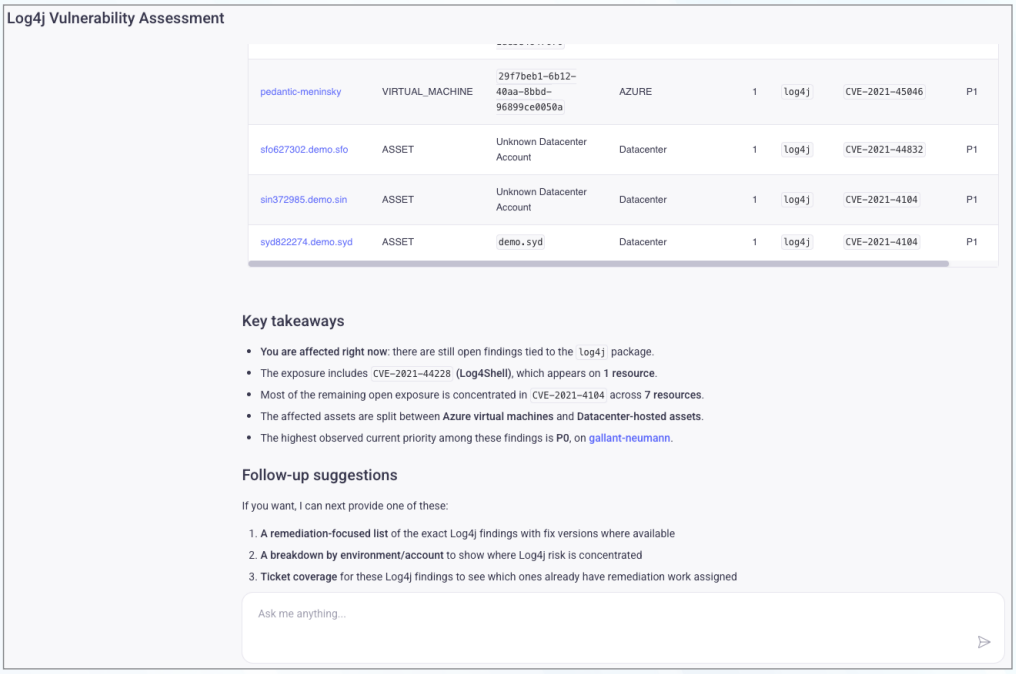

Figure 24. Seema Log4j Exposure Summary

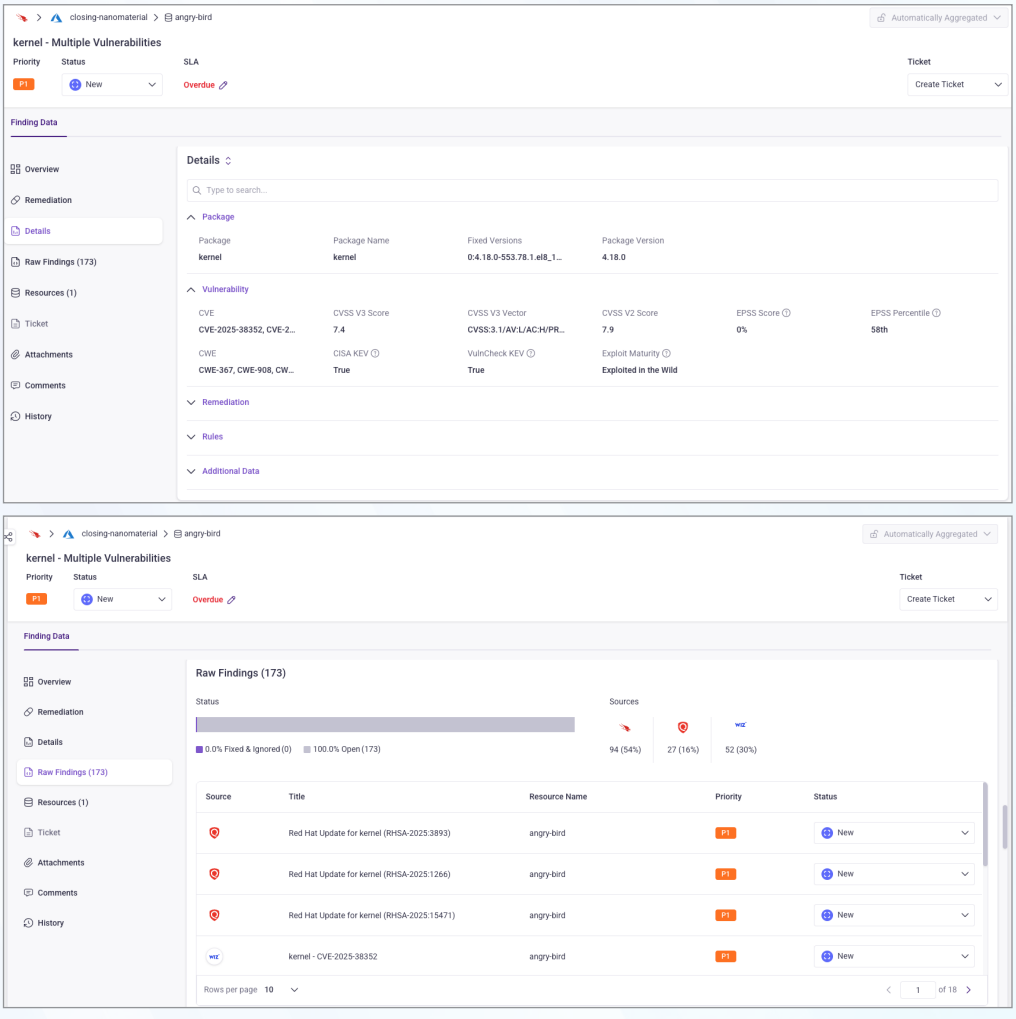

Figure 25. Vulnerability Aggregation and Unified Remediation

Figure 26. Detailed AI-Driven Remediation Guidance

Complementing this aggregation capability is the introduction of AI-generated remediation guidance. The platform can automatically produce detailed, step-by-step instructions for resolving vulnerabilities, including specific commands and operational steps. These instructions are embedded directly into tickets and reports, ensuring that remediation teams receive actionable guidance without needing to conduct additional research. This capability directly addresses one of the most common bottlenecks in vulnerability management: the gap between identifying issues and executing fixes. Figure 26 shows an example of more detailed remediation steps to address an issue, including command-line arguments.

Seemplicity has also expanded its use of AI in asset ownership and organizational context mapping. Through its “smart tags” functionality, the platform analyzes metadata from sources such as configuration management databases (CMDBs) and asset management systems to infer ownership, organizational hierarchy, and responsibility for remediation. This enables automated alignment of findings, dashboards, and workflows to real-world organizational structures, such as geographic regions or business units. As a result, organizations can more easily track performance and accountability across distributed environments without requiring extensive manual tagging or configuration.

A particularly notable addition is the AI Insights Feed, which addresses the growing challenge of "dashboard fatigue." Instead of requiring users to manually interpret dozens of KPIs, the platform continuously analyzes available metrics and surfaces key insights, typically highlighting both positive trends and areas of concern. For example, it may identify improvements in new findings while simultaneously flagging declining SLA performance or aging critical vulnerabilities (see Figure 27).

Figure 27. AI Insights

These insights can be scoped to specific environments, such as a GitHub environment or business unit, enabling more targeted operational awareness.

From an operational perspective, the overarching theme of these enhancements is reducing friction across the exposure management lifecycle. Seemplicity is using AI to address three primary challenges:

1. **Data access and interpretation—**Simplifying how users query and understand exposure data
2. **Remediation prioritization and execution—**Consolidating findings and providing actionable guidance
3. **Visibility and reporting—**Proactively surfacing insights rather than relying on static dashboards

This approach reflects a broader trend in security tooling, where AI is increasingly used not just for detection, but for workflow acceleration and decision support. In Seemplicity's case, the integration is tightly coupled to real operational use cases, particularly around vulnerability management, application security, and exposure prioritization. Overall, the platform's AI enhancements represent a maturation of exposure management from a data aggregation problem to an operational efficiency problem. By embedding AI across multiple stages of the workflow, Seemplicity is aiming to reduce manual effort, improve remediation speed, and provide more actionable intelligence to both technical teams and leadership.

Conclusion

As a security operations analyst, we could see Seemplicity is a practical solution that tackles one of the biggest pain points in modern security operations: the overwhelming backlog of vulnerabilities, misconfigurations, and security findings scattered across dozens of tools. Its core strength lies in how it automatically aggregates, normalizes, and deduplicates findings from various scanners, cloud security posture management (CSPM), vulnerability management, and pen test reports into a single, actionable remediation queue. This unified risk inbox reduces the noise and manual sorting that eats up analysts' time and ensures that only relevant, prioritized tasks are sent to the teams responsible for fixing them, through integrations with popular ticketing systems like Jira and ServiceNow.

Seemplicity's automation of workflow handoffs, real-time tracking, and built-in metrics around mean time to remediate (MTTR) help security teams bridge the gap between detection and resolution—one of the most common operational bottlenecks. By doing so, it aligns security, IT, DevOps, and cloud teams under one consistent process, eliminating the reliance on endless spreadsheets and fragmented communication. For organizations struggling with growing vulnerability debt, compliance audit gaps, or slow patch cycles, Seemplicity provides clear visibility into what's getting fixed, what's overdue, and where bottlenecks occur, empowering security teams to demonstrate measurable risk reduction and improve overall cyber hygiene without adding more manual workload.