

How Happy Returns turned NodeZero® WebApp testing into verifiable proof developers could use

After about a dozen early NodeZero WebApp tests, the team expanded to nearly 30 autonomous WebApp pentesting campaigns across more than 20 applications, giving security and engineering teams a shared view of exploitable risk inside an AWS-heavy environment.

Customer *snapshot*

CUSTOMER

Happy Returns

INDUSTRY

E-commerce returns and
reverse logistics software

ENVIRONMENT

AWS-heavy infrastructure,
Kubernetes, customer-
facing web applications with
a Security Team spanning
cloud, infrastructure,
application security,
privacy, and compliance

CHAMPION

Phillip Walsh
Head of Information Security

SOLUTION

Horizon3 NodeZero for
AWS-focused validation
and NodeZero WebApp
for continuous application
testing

CHALLENGE

Operational context

Happy Returns helps merchants manage returns at scale, which means its customer-facing applications and supporting cloud workloads sit close to the revenue path. The business runs primarily in AWS, relies on modern application infrastructure, and asks the security team to cover cloud, infrastructure, AppSec, privacy, and compliance at the same time. That makes security risk less about one isolated bug and more about whether hidden weaknesses across applications, identities, and cloud services can combine into something with real business impact.

The security gap

That was the core tension. Many of the tools the team evaluated were built around Active Directory-focused attack paths that did not reflect how this environment was actually designed. The next bottleneck was not another on-prem assessment. It was web application coverage in an AWS-heavy business, plus a better way to prove what was truly exploitable. The team already had scanners, testing tools, and periodic pentests. What it still lacked was a continuous way to validate risk and bring developers evidence they could act on without another round of report arguments.



WHY NOW

Proving operational fit

The first proof point came from operational fit. Early NodeZero testing showed that the platform could operate credibly in the company's AWS-hosted environment, including Kubernetes and the connections between cloud assets. In the first few weeks alone, the team ran about a dozen tests and kept refining scope, runner placement, and segmentation to get cleaner answers. In one AWS pentest, NodeZero mapped about 2,000 credential access edges across roughly 10 credentials and 50 hosts, giving the team a concrete view of how access relationships could compound risk across the environment.

Driving continuous testing

A second moment came when that infrastructure validation pointed to a bigger opportunity. What the team wanted next was continuous, production-aware WebApp testing that fit a cloud-native operating model better than another static report. By May 2026, it had already run nearly 30 WebApp tests and identified more than 20 applications for early testing. That level of use mattered because it showed the team was not treating WebApp as a side experiment. It was using it to pressure-test how customer-facing risk should be validated going forward.

The power of evidence

The strongest moment, though, was how quickly the evidence model resonated with Phillip Walsh, their Head of Information Security. When Horizon3 showed an early visual approach for mapping findings to application routes and making proof easier to navigate, the response was immediate: it would be far more useful in conversations with developers than previous vendor findings. Phillip put it plainly: "Instead of just looking at a report, you can actually show what matters." That is the turning point in the story. The value was no longer just finding issues. It was giving security and engineering teams a shared view of what was actually happening so decisions could move faster.



WHY HORIZON3

Platform alignment

Horizon3 earned its place because NodeZero aligned with the environment the team actually had, not the one many security tools assume. It gave the organization a centralized way to validate cloud and application risk across an AWS-heavy estate, while reducing the manual follow-up that usually follows scanner output. Just as important, it made the findings easier to socialize. Instead of asking engineers to react to abstract severity ratings, the team could bring route-level or attack-path evidence into the conversation.

Evolving the operating model

That changed how the platform fit into the security program day to day. NodeZero started as an infrastructure fit check, then became part of a broader operating model for continuous validation. As WebApp matured, Happy Returns became an active design partner, shaping the product with feedback on authentication, route visibility, and proof. That kind of engagement is a sign that the platform was solving a real operating problem, not just adding another dashboard.

VERIFIABLE RESULTS

The results are strongest as program evidence rather than a single vanity number.

- The team moved from evaluating fit to sustained use, running about a dozen early NodeZero tests, nearly 30 WebApp tests by May 2026, and lining up more than 20 applications for early coverage.
- Security gained a clearer view of compounded cloud risk, including a mapped set of roughly 2,000 credential access relationships across 10 credentials and 50 hosts in one AWS pentest.
- Application security conversations became more evidence-led. The operating shift described in the source draft was from interpreting scanner noise to reviewing proof that developers and security teams could assess together.
- The relationship expanded beyond initial platform use into an active WebApp design partnership, a strong sign that Happy Returns saw long-term strategic value in the approach.



LOOKING AHEAD

Next steps

The next chapter is clear. Happy Returns is likely to deepen authenticated WebApp testing, extend production coverage across more of the application estate, and keep pushing for evidence that shortens the distance between finding, validation, and fix. What has changed already is more important: security is no longer relying on periodic reports to argue for action. It now has a stronger way to test how risk actually behaves in a cloud-native environment, and to show why it matters before attackers do.

