

ALPHA DEFENSE · WHITE PAPER

What *Others* Miss

Why deep security testing finds the risk that changes decisions.

CORE THESIS

The most valuable findings are not always the loudest. They are the findings that change what a buyer, security leader, product team, or engineering organization does next.

Prepared for CISOs, product security leaders, M&A diligence teams, and technical executives evaluating penetration testing depth.

WE FIND WHAT OTHERS MISS.



Executive Summary

Security leaders do not need another long report full of undifferentiated findings. They need confidence about what is exploitable, what is material, and what should change now.

That is the gap Alpha Defense is built to close. The strongest security work does not stop at a status code, a tool finding, an architecture claim, or a generic severity score. It tests the assumption behind the control and follows the evidence until the business decision becomes clearer.

This paper uses three anonymized representative engagements to show what “we find what others miss” looks like in practice. In one case, the missed issue was a hidden tenant-boundary failure. In another, it was critical product risk inside firmware and a custom protocol. In the third, it was the opposite of a new vulnerability: proof that only a small subset of flagged issues created immediate exploitable risk.



READER TAKEAWAY

*The point is not that deeper testing always produces more findings. The point is that deeper testing produces **better decisions**.*

The Assurance Gap

Most organizations already have some form of security assurance: annual penetration tests, automated scans, compliance assessments, third-party questionnaires, SBOM reviews, vulnerability reports, or product security checklists. Those activities can be useful. They can also create false confidence when the test is not shaped around how the system actually works.

The gap usually appears in places that are hard to standardize: authorization logic, tenant boundaries, custom protocols, embedded devices, business workflows, compensating controls, and vulnerability reachability. These areas require curiosity, judgment, and persistence from the person doing the work.

For a CISO or product leader, the practical question is not simply “Was this tested?” It is *“Was it tested deeply enough to answer the decision we actually need to make?”*

The Tenant Boundary That Wasn't

| ENGAGEMENT TYPE | DECISION AT STAKE | WHAT WAS MISSED |
|--|--|---|
| Application security and M&A due diligence | Whether the buyer could trust the target's data-isolation claims | A session-scoped value that looked like an ordinary token but also acted as an obfuscated data filter |

A client asked Alpha Defense to assess a web application they were considering acquiring. The target provided several years of prior penetration testing reports from a reputable security firm. Source code was out of scope, but Alpha Defense had application access and access to relevant product and engineering stakeholders.

At kickoff, the target explained that cross-customer access was technically impossible because customer records were separated at the infrastructure level. That assurance was important. For the buyer, tenant isolation was not an abstract control. It was part of the acquisition risk model.

○ WHERE SURFACE TESTING STOPS

A request returns 200, so the behavior looks normal

A tool can send a request, receive a successful response, and treat the behavior as normal. In this case, the application continued to work when a session-token value was removed. It still returned a normal-looking page of results, and pagination still showed the expected number of records per page. From a scanner's perspective, the request could look successful with or without the value present.

● WHAT ALPHA DEFENSE DID DIFFERENTLY

Alpha Defense manually compared the actual records returned across roles, sessions, and data contexts. The key difference was not that the page loaded. It was that the data changed. When the value was removed, the application returned more records than it should have returned, even though the page still appeared structurally normal.

That human review showed the value was not merely an authentication artifact. It was influencing how the application filtered data by domain, organization, and production instance. Once that was understood, Alpha Defense could demonstrate that the application's tenant boundary depended on inconsistent data-scoping behavior rather than the hard separation the buyer had been told to expect.

DECISION IMPACT

The finding changed the buyer's view of the target's risk. The value was not just a vulnerability report – it was better information before a business decision.

The Custom Protocol No Scanner Understood

| ENGAGEMENT TYPE | DECISION AT STAKE | WHAT WAS MISSED |
|---|--|---|
| Connected / IoT device, firmware, and custom protocol testing | Whether product risk was tested where standard tools could not reach | Critical vulnerabilities inside product-specific firmware and a custom communication protocol |

A client with a connected device product line brought Alpha Defense in to assess multiple products over a multi-year program. One product manager was initially reluctant to provide source code. The concern was understandable: previous testing had been conducted without it, and the argument was that a real attacker would not have source code either.

Alpha Defense started where the client was comfortable. The team assessed the device without source code, extracted and reverse-engineered firmware, analyzed the custom communication protocol, and built test tooling tailored to the product.

○ WHERE STANDARD METHODS FALL SHORT

A checklist cannot test what public tooling cannot see

Custom devices and embedded products often contain the most important risk in areas that generic scanners do not understand. If the protocol is proprietary, the behavior is device-specific, and the attack surface is not represented in public tooling, the test cannot be reduced to running a standard checklist.

● WHAT ALPHA DEFENSE DID DIFFERENTLY

Alpha Defense built the method the assessment required. The team reverse-engineered the relevant behavior, created custom tooling, and wrote a fuzzer for the protocol. That work produced multiple critical vulnerabilities with enough technical clarity for the client to act.

The engagement also changed the relationship. Once the product manager saw the quality and intent of the work, the client provided source code and connected Alpha Defense directly with senior engineers for the rest of the program — a trusted product security partner, not a testing vendor at the end of the release cycle.

DECISION IMPACT

The client was able to remediate serious product risk before it became public, customer-facing, or product-threatening.

The 5% That Actually Mattered

| ENGAGEMENT TYPE | DECISION AT STAKE | WHAT WAS MISSED |
|--|---|--|
| Vulnerability prioritization and reachability analysis | Which findings needed immediate remediation, and which could wait | The difference between vulnerable code being present and being reachable and exploitable |

A client had invested in a commercial binary scanning platform and received a report identifying numerous vulnerable libraries and insecure components compiled into an application. The report provided useful visibility, but it did not answer the operational question the team most needed answered: which findings created immediate, exploitable risk in the actual product?

○ WHEN SEVERITY IS NOT ENOUGH

An accurate list can still leave the team triaging in the dark

A vulnerability list can be technically accurate and still be operationally unhelpful. If a report says a vulnerable component exists but does not analyze whether the affected code is actually called, reachable, or exploitable in the deployed application, the engineering team is left to triage in the dark.

● WHAT ALPHA DEFENSE DID DIFFERENTLY

Alpha Defense analyzed reachability and exploitability in the context of the actual product. The team separated theoretical exposure from practical risk and identified the subset of vulnerable code that could realistically be exercised.

The result was clear: only a small portion of the flagged vulnerable code created immediate exploitable risk. The client could patch that subset quickly while continuing to address the broader backlog in a more deliberate way.

~5%

of the flagged vulnerable code created immediate, exploitable risk — the subset the client patched quickly, instead of treating every finding as equally urgent.

DECISION IMPACT

The client issued a targeted patch within a day instead of treating the work as a months-long, undifferentiated remediation project.

The Pattern: What Others Miss Is Rarely Obvious

Across these engagements, the common thread is not a single vulnerability class. It is a way of working. Alpha Defense looks for the places where conventional assurance can appear complete while leaving the hardest question unanswered.

| WHERE TESTING OFTEN STALLS | WHY IT MATTERS | THE BETTER QUESTION |
|-------------------------------------|---|--|
| A request returns 200 | A successful response can still contain the wrong data or enforce the wrong boundary. | <i>Did the application return the right data for this role, tenant, and context?</i> |
| Architecture says separation exists | Design intent does not prove runtime behavior. | <i>Can the isolation claim be demonstrated under realistic use and abuse conditions?</i> |
| No scanner supports the protocol | Custom systems often hide risk outside generic tooling. | <i>What tooling must be built to test this product on its own terms?</i> |
| A library is flagged as vulnerable | Presence alone does not prove exploitability or priority. | <i>Is the vulnerable path reachable, callable, and exploitable in this application?</i> |
| The report lists many issues | A longer list can slow action if everything appears equally urgent. | <i>What should the team fix first to reduce real customer risk?</i> |

What Security Leaders Should Ask

When evaluating a penetration testing partner or reviewing an existing assessment, these questions reveal whether the work is likely to find what matters — or simply document what is easy to check.

-
- 01 Who actually runs the test.** Is the expertise demonstrated during scoping the same expertise applied during the technical work?

 - 02 How authorization is tested.** Does the team inspect the data actually returned across roles, tenants, workflows, and edge cases — treating architecture claims as hypotheses to prove rather than assurances to accept?

 - 03 What happens when tooling ends.** Can the team reverse-engineer, build harnesses, write fuzzers, or create custom methods when the target environment requires it?

 - 04 How exploitability is determined.** Does the report distinguish between vulnerable code that exists and vulnerable code that is reachable in the deployed product?

 - 05 How findings translate into action.** Does the output help the organization decide what to fix now, what to monitor, and what can wait?
-

Operating Principles

The stories in this paper are different, but the operating principles behind them are consistent.

Senior-led from scope through report

The expert who scopes the test is directly involved in running it, interpreting the evidence, and explaining the business impact.

Manual verification where it matters

Automated tools can help identify leads, but Alpha Defense validates behavior and follows the data when the result depends on context.

Custom tooling for custom systems

When a target uses firmware, embedded logic, proprietary protocols, or unusual workflows, the assessment method adapts to the product.

Exploitability before alarm

Findings are evaluated in terms of reachability, compensating controls, attack path, and practical impact.

Partnership over politics

The goal is to help product and engineering teams understand risk, fix what matters, and build trust around the work.

Conclusion

The phrase “we find what others miss” is not a promise to make every report longer or every finding more dramatic. It is a commitment to test the assumptions that matter, build the methods the system requires, and separate theoretical concern from exploitable risk.

For security leaders, that distinction is practical. It affects acquisition decisions, product release decisions, remediation priorities, customer risk, and the credibility of the security program itself.

NEXT STEP

Bring Alpha Defense the system, report, or concern that still feels unresolved. We will help determine what is exploitable, what matters most, and what your team should do next.

alphadefense.com info@alphadefense.com 888-275-1001



ABOUT ALPHA DEFENSE

Alpha Defense helps organizations understand and reduce real security risk through senior-led assessment work. The expert who scopes your test is the one who runs it. *We find what others miss.*

Note: The examples in this paper are anonymized representative engagements. Specific client names, product details, and exploit mechanics have been generalized to protect confidentiality and avoid publishing actionable attack instructions.