



Whitepaper

Security and Application Teams Are Buried; It's Time to Dig Them Out

Organizations are suffering from too many noisy alerts and not enough essential context. A next-level application security posture management (ASPM) tool can help them break free.



Table of Contents

Application Risk on the Rise: Caution Without Context	3
Critical Questions, Unclear Answers	3
The Path to Proactive: Gaining a Full Picture of Risk	4
When Specialized Becomes Siloed	5
ASPM: Unify, Contextualize, Prioritize	6
Longbow: Unified Risk Remediation for Cloud-Native Applications	7
Next-Level Capability 1: See the Full Picture of Risk	9
Next-Level Capability 2: Identify Root Issues	12
Next-Level Capability 3: Remediate Risk with Targeted Solutions	13
Case Study: Financial Company Revolutionizes its Approach to Risk	15
Summary	16
Learn More and Get Started Today	16

Application Risk on the Rise: Caution Without Context

There may have been a time when security teams lacked strong detection capabilities, but that's no longer the case. Detection tools have proliferated, and these tools are effective—perhaps even *too* effective.

They produce too many alerts but not enough context. Rather than helping stretched security analysts remediate cloud-native application risks, these tools inundate them with more alerts than they can possibly handle, giving them no way to prioritize the most urgent ones.

Many large organizations now have backlogs with millions of vulnerabilities, and more added every day. Remediating them all is out of the question. The best they can do is tread water and try to keep the backlog from growing.

In the modern world of software development where innovation speed, quality and efficiency are all essential to stay competitive, this challenge is quickly becoming a crisis.

Critical Questions, Unclear Answers

When analysts receive an application security alert from a detection tool, they have many questions to answer, such as:

- Is the vulnerability legitimate?
- Where did the vulnerability come from?
- Which asset is at risk? Is it a high-value asset?
- Is the the asset in production or development?
- Who owns the asset within the organization?

Detection tools don't answer any of these questions, and that's by design. Their job is simply to detect application vulnerabilities; it's up to security analysts to decide how to respond. Analysts could answer these questions via manual research, but they don't have that kind of time—especially not for the volume of vulnerabilities that now demand their daily attention.

Instead, analysts are forced to do the best they can with the limited information they have. They'll make educated guesses about which vulnerabilities need to be addressed first and hand them off to developers to fix.

This leaves developers feeling frustrated by the steady flow of half-baked remediation tickets coming their way. Lacking the insights they need to execute fixes quickly, risk remediation becomes an unwanted chore that takes them away from their day jobs: building innovative new features, products and services.

Without a clear solution to this growing challenge, remediation rates will not keep up with detection, organizational security postures will weaken and business risk across industries will escalate.

The Path to Proactive: Gaining a Full Picture of Risk

To solve this challenge, security teams must rapidly gain clarity into cloud-native application risk.

Luckily, there's a very simple equation that defines application risk:



Likelihood x Impact = Risk

Unfortunately, many security teams struggle to use this equation to their advantage because they focus too much on the likelihood side of the equation. Their understanding of likelihood is based on models like the Common Vulnerability Scoring System (CVSS) and the Exploit Prediction Scoring System (EPSS), which rate the severity of vulnerabilities based on different factors, including whether there's exploit code available.

CVSS, EPSS and similar models can be helpful, but analysts shouldn't rely on them for everything. Just because an exploit is likely to occur doesn't mean it's urgent. To understand impact, analysts need context specific to their own assets and environment.

Consider a vulnerability that's rated as critical severity, but it's found in a low-value asset that's not in production. For instance, the asset could be a "Hello, World!" sandbox server that's used for testing and isn't accessible from the internet. Even if it were exploited, the impact would be minimal. Despite the severity, the vulnerability would be low-risk, and remediating it would not be urgent.

When Everything's an Emergency, Nothing Is

Severity analysis based on assumptions—how bad a vulnerability *could* be in the worst-case scenario—has been the basis of the entire cybersecurity industry for decades. However, with the massive volume of vulnerabilities that security teams have to manage these days, the limitations of severity analysis are clear.

When so many vulnerabilities are potential emergencies, analysts can't prioritize their work, and they end up overwhelmed. Actual security emergencies—exploits that are both likely and impactful—may fall through the cracks.

When Specialized Becomes Siloed

Another challenge to gaining holistic risk clarity is fragmented application visibility. Security vendors have traditionally specialized in different stages of the application pipeline. For example:

Application security testing (AST) tools focus on the development stage.

Cloud-native application protection platform (CNAPP) tools focus on the runtime stage.

Neither of these can look outside their own environments:

- AST tools detect vulnerabilities in code but can't tell users which ones made it into production.
- CNAPP tools detect runtime issues but can't tell users where they originated.

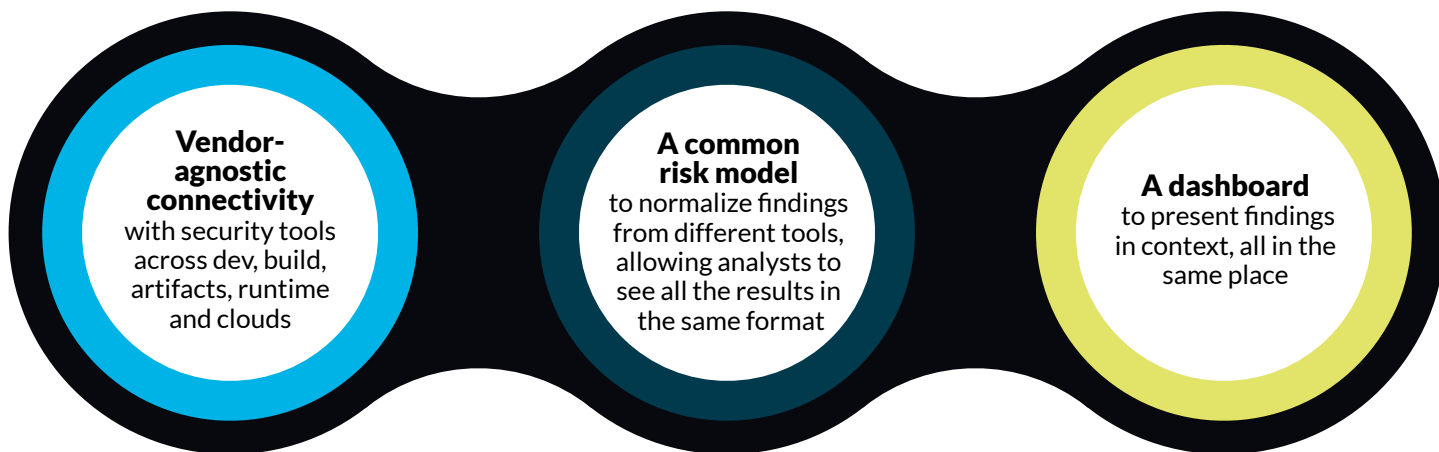
Container images, infrastructure as code (IaC) tools and cloud service provider security toolsets all add further complexity. It's easy to see why tracking vulnerabilities across the cloud-native application pipeline is so difficult, and why analysts often end up overwhelmed.

Security teams need to bridge the gaps between different detection tools. A new class of security tools, known as application security posture management (ASPM), has emerged to help them do that. A next-level ASPM tool can help security teams prioritize the right remediation solutions at the right time, while also empowering developers to work more efficiently.

ASPM: Unify, Contextualize, Prioritize

ASPM tools are designed to ingest data from different detection tools, unify that data, and add the context that the detection tools can't provide.

The ASPM market is still new, so the exact definition is in flux. However, there's basic consensus that all ASPM tools should include the following components:



It may also be helpful to consider what an ASPM tool is not. It's not intended to react to ongoing security events; instead, it helps security teams proactively manage risk and address vulnerabilities before they're exploited.

ASPM tools also don't replace detection tools. They pick up where those other tools leave off, using their findings as the basis for further analysis. By cutting through the clutter and identifying the most urgent vulnerabilities, ASPM makes detection tools useful again, giving companies a return on the security investments they've already made.

Security teams can benefit from choosing an ASPM tool that includes advanced features that go beyond the minimum capabilities mentioned above. These progressive tools provide more than just visibility; they deliver targeted solutions that remediate risk in a clear, quantifiable way.

This is what Longbow, the next-level ASPM tool from Veracode, offers.

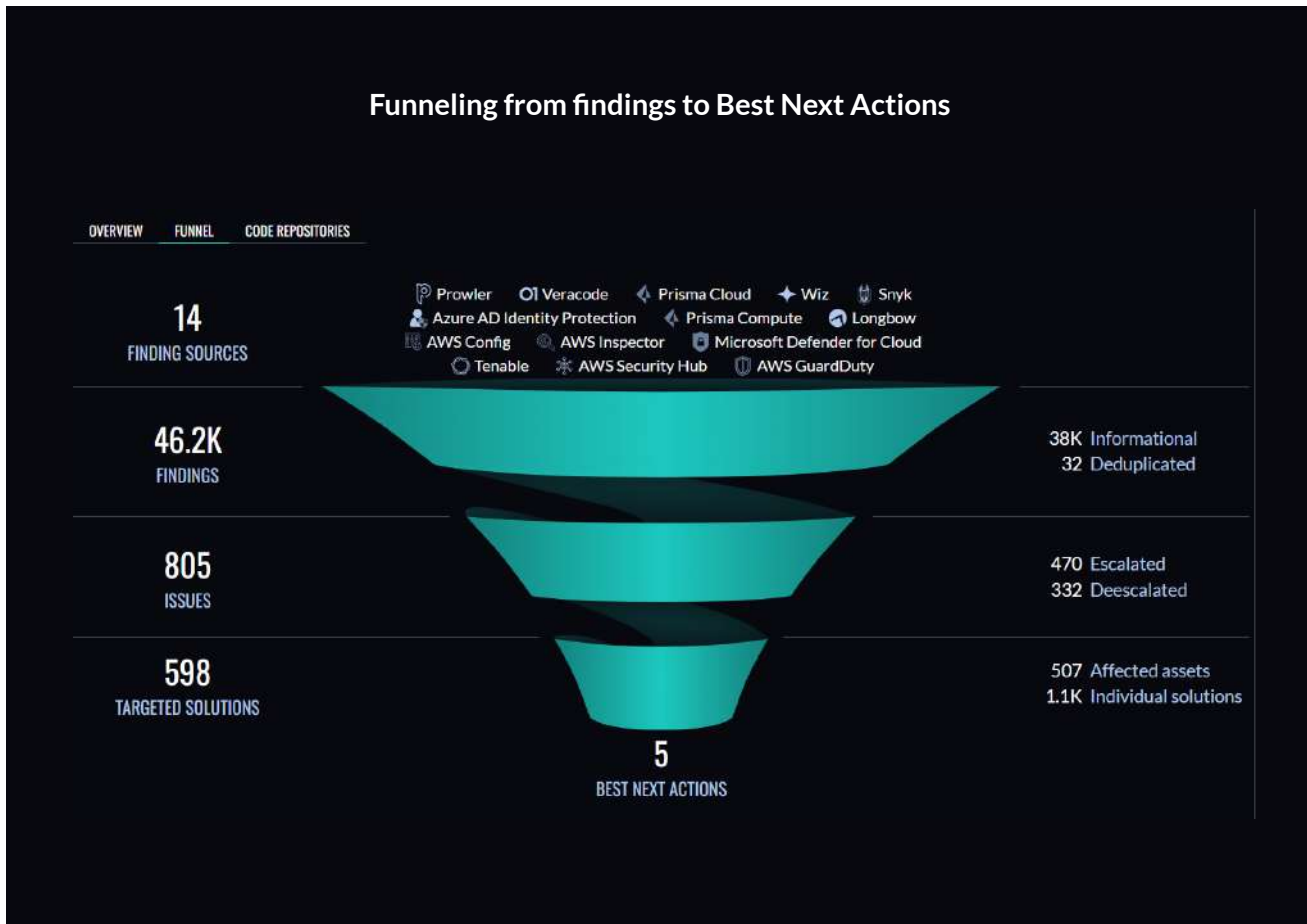


Longbow: Unified Risk Remediation for Cloud-Native Applications


Longbow was built for security people, by security people. It arms security teams to identify critical risks and prioritize them based on their application environment and business context.

While Longbow does include basic ASPM features, it goes beyond existing ideas of what ASPM tools are and what capabilities they should provide. The tool's sole focus is on giving security analysts the insights they need to prioritize the right work at the right time and partner with developers to remediate cloud-native application risk quickly.


Longbow works like a funnel, taking a massive number of findings and condensing them down to a few Best Next Actions™—the select solutions that will remediate the most risk with the least effort.




To help analysts do *more*, Longbow gives them *more* capabilities—ones that basic ASPM tools can't match. With Longbow, analysts are empowered to:

1 

See the full picture of risk across tools, grouping and filtering how they want

2 

Identify root issues across the application pipeline

3 

Remediate risk and decrease the volume of developer tickets with targeted solutions

SOLUTIONS 598 Targeted solutions

805
ISSUES

URGENCY

Urgent	81
High	115
Medium	401
Low	208

507
AFFECTED
ASSETS

CLOUD ACCOUNT

221433242586	283
demo	151
test-account	61
AltaTestEnvironment	9
Other	3

BEST NEXT ACTIONS View All

These solutions reduce the most cloud risk:

			Recommended Changes	Addressable Issues	Risk Reduction
1	> REMEDIATION	Update IaC Template: GITHUB: prod_mobile_app_module	9	40	1,986
2	> REMEDIATION	Update File in Source Repo: GitHub Repository: prod-website-container	4	17	1,251
3	> REMEDIATION	Update IaC Template: GITHUB: prod_website_module	6	19	818
4	> REMEDIATION	Update File in Source Repo: GitHub Repository: platform-pyngin	3	5	456
5	> REMEDIATION	Update File in Source Repo: GitHub Repository: platform-terraform	2	6	294

Let's examine each of these three advanced ASPM capabilities in more detail.



Next-Level Capability 1: See the Full Picture of Risk

The Longbow contextual engine links findings from detection tools with specific assets. This is the first step toward providing the context and specificity many security teams lack. It doesn't just tell analysts which vulnerabilities are worse; it tells them which ones are worse for their situation.

By understanding both the severity of a vulnerability and its potential impact, analysts get a much clearer picture of the risk it presents. In turn, this determines the urgency of remediating it.

The tool assigns each finding a baseline risk score using normalized severity values. This allows analysts to sort findings by severity, without having to manually track down this information across multiple tools. Hazardous findings are promoted to issues, which are scored using additional security factors pulled from EPSS and similar sources.

Longbow also detects when findings from different tools provide overlapping information about the same vulnerability, merging them into a single issue. The common risk model helps reconcile differences between the findings and puts them all in the same format.

Automated Factor Analysis

Starting with the baseline scores, each issue goes through an automated factor analysis to add additional context around urgency. The model adjusts risk ratings up or down based on numerous factors, such as where the asset is deployed (production or development), how valuable the asset is, how accessible it is, and whether its data is backed up and encrypted.

Issues associated with crown-jewel assets in production require urgent remediation, and the tool rates them accordingly. Issues linked to assets in non-production sandbox environments receive lower risk scores, even for critical vulnerabilities.

Key Terms

Findings

The results provided by detection tools.

Issues

Specific hazardous findings as they apply to specific assets. Longbow automatically promotes findings to issues after confirming their severity.

Asset

Anything in an application environment that faces risk and benefits from monitoring. This could be something as small as a single Docker file all the way up to an entire virtual private cloud.

Longbow’s analysis extends beyond just vulnerabilities and misconfigurations. It can also incorporate indicators of compromise, malware and data exposures. This supports the goal of helping security teams understand urgency better. For instance, the tool might identify a severe vulnerability impacting an asset that may have already been compromised. This would indicate much greater urgency than an asset that’s theoretically exploitable but has no indications of compromise.

Automated factor analysis provides better context around risk

The dashboard shows the following details for the issue 'EC2 Instance Exposed to Internet':

- Issue Factor Category:** ALL FACTORS (selected), RISK INCREASING, RISK NEUTRAL, RISK DECREASING
- Summary:**
 - Business Value:**
 - Application Value: CROWN JEWEL
 - Cost: EST. COST 10¢-\$10/HOUR
 - Functional Role: UNKNOWN
 - Account Value: UNKNOWN
 - Environment: NOT ASSESSED
 - Asset Value: NOT ASSESSED
 - Data Classification: NOT ASSESSED
- Advanced:**
 - Attack Accessibility:**
 - Internet-Facing: 1+ PORTS AND 1+ IPS
 - Authorized Entity Count: 3+ PRINCIPALS ACCESS
 - Asset State: ASSET STATE IS ON
 - Data Exposure:**
 - Encryption: DATA AT REST, NO VOLUME, ENCRYPTION
 - Data Integrity:**
 - IAC DRIFT, NOT TEMPLATE BASED
 - Service Disruption:**
 - Restorability: BACKUP IS 7+ DAYS OLD
 - Role Permissions Level: NOT ASSESSED
 - Compliance:**
 - Visibility: AGENT, NO SSM AGENT
 - Visibility: LOGGING, 1 LOG SOURCE, ENABLED
 - Visibility: TAGGING, HAS TAGS
 - Data Compliance: NOT ASSESSED

Based on the results of the factor analysis, Longbow quantifies risk for each issue, allowing analysts to instantly identify their top remediation priorities.

Grouping and Filtering

The tool also allows analysts to get the exact view of risk they're looking for by grouping and filtering issues.

For instance, they could aggregate all the issues by asset to see which assets have the highest risk score. They could also group by application or business unit to show which areas would benefit the most from targeted remediation efforts.

Analysts might also want to look at vulnerabilities first and then move on to misconfigurations. It's all about giving them the freedom to focus their efforts while aligning with their unique goals and priorities.

Transparent, Explainable Security

Longbow is not a black-box model. It doesn't just give analysts recommendations and expect them to trust those recommendations. Instead, it shows them why the tool provides the results it does.

All factors included in the analysis are identified and summarized so users can see for themselves which ones are driving the risk score up or down. The tool also provides detailed descriptions of those factors. In some cases, this could even include linking users to the original source that Longbow pulled the factor details from.

The screenshot shows a detailed view of a security issue titled "EC2 Instance Exposed to Internet". The interface is dark-themed and includes a sidebar with navigation options like "SUMMARY", "SOLUTIONS", "CONTEXT", "FACTORS", "FINDINGS", "Timeline", and "ADVANCED". The main content area is divided into sections: "DESCRIPTION" (EC2 Instance is accessible via public internet), "ASSET" (AWS EC2 Instance: arncaws:ec2:us-west-2:221433242586:instance/i-Ob...), and "FINDINGS" (17 hours ago, Longbow: Internet-Facing, 1+ Ports and 1+ IPs). A prominent "100 URGENT" risk score is displayed, along with a "50 Severity" indicator. The "URGENCY FACTORS" section lists 19 items, including "Internet-Facing" (1+ Ports and 1+ IPs), "Authorized Entity Count" (3+ Principals Access), "Application Value" (Crown Jewel), "Visibility: Agent" (No SSM Agent), "Encryption: Data at Rest" (No Volume Encryption), "Role Permissions Level" (Medium Risk Permissions), "Visibility: Logging" (2 Log Sources Enabled), and "Functional Role" (Server). The "SEVERITY FACTORS" section lists "Visibility: Tagging" (Has Tags) and "Restorability" (Backup Is 7+ Days Old). Navigation buttons for "CREATE TICKET" and "ACT" are visible in the top right corner.



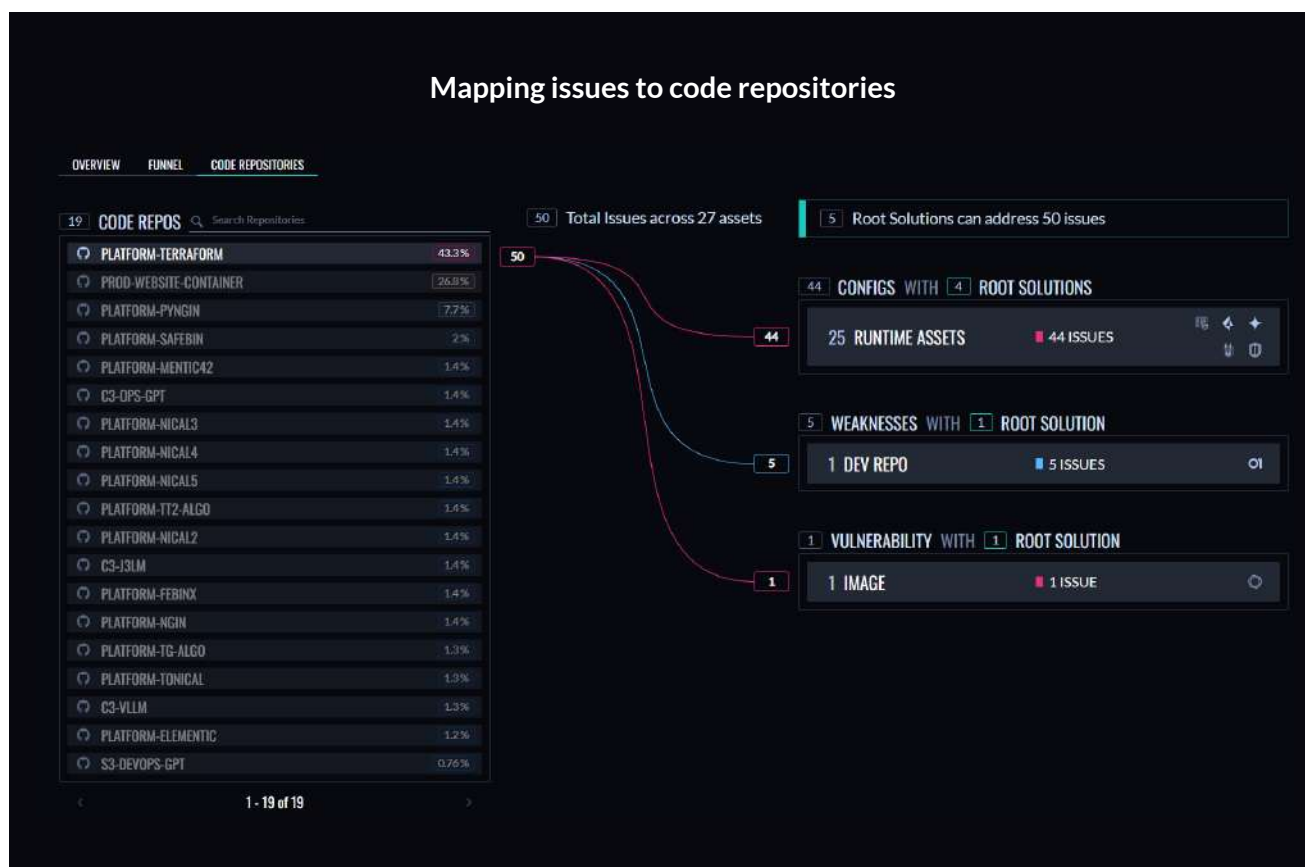
Next-Level Capability 2: Identify Root Issues

Finding issues in production is important, but the ultimate goal of any security team is to remediate risk, not just identify it. Remediating issues in production requires holistic visibility into where those issues originated.

Analysts can't identify root issues using detection tools alone, because CNAPP tools don't look back and AST tools don't look forward. Instead, they either have to do time-consuming manual analysis or just leave developers to figure it out for themselves. This makes life harder for both security and developers, and it keeps the remediation rate low.

Longbow automatically maps runtime issues back to the codebase. It can do this because it runs at a level above the detection tools, with visibility and analytics capabilities those tools can't match.

The analysis includes identifying vulnerabilities and misconfigurations in IaC code repositories and container images. The tool visualizes the output of this mapping, helping analysts see which areas of the codebase introduce the most risk.



Efficient Solutions Based on Root Cause

Suppose an analyst identifies an issue in a particular container. Longbow could show that the container came from a particular image and that the image came from a particular code repository.

Then, the tool can show the analyst the individual files within that repository and what dependencies those files have. Based on this, the analyst might identify that a single line of code in a Docker file is the root cause of 20 different issues across 20 production assets.

Without Longbow, they'd have no way to easily discover that. They'd be stuck treating them as 20 separate issues, instead of 20 instances of the same issue. They might end up remediating them one by one—assuming they get remediated at all. With Longbow, they can simply fix the Docker file once to remediate all the issues.

Analysts can also filter by origin to see the most common sources of risk in their codebase. This helps them work more efficiently, since they can ensure the remediation steps they're taking are based on actual prioritized risk, not guesswork.



Next-Level Capability 3: Remediate Risk with Targeted Solutions

Getting holistic visibility into risk and analyzing the root causes of that risk are both important aspects of application security. But neither of them matters unless they drive remediation. Longbow not only provides targeted solutions based on root cause; it identifies Best Next Actions, helping analysts see how their work could have the biggest impact.

Group Solutions for Efficient Remediation

When users group issues in Longbow, their settings carry over into the solution stage. If they've grouped issues to determine which applications are introducing the most risk, the tool could provide targeted solutions based on application risk.

If an analyst urgently needs to remediate risk for a particular application, they can see all the solutions that apply to that application. Then, they can assign those solutions to developers who know that application the best. Those developers will be well-positioned to remediate risk quickly while minimizing disruption.

Longbow groups solutions by where the remediation takes place. If a developer is remediating an urgent issue in a particular file, the tool can help them identify other, less-urgent issues in the same file. Since they'll have the file open anyway—and they'll need to QA the updated file whether they remediate one issue or many—why wouldn't they work on the other issues as well?

Even if the issues aren't urgent, enough of these "small wins" could add up to significant risk remediation. The alternative is for developers to bounce around randomly in the codebase, tackling individual issues as they arise. Remediating issues one by one is inefficient and doesn't make developers feel like a valued part of the security process.

Choose From Multiple Solutions

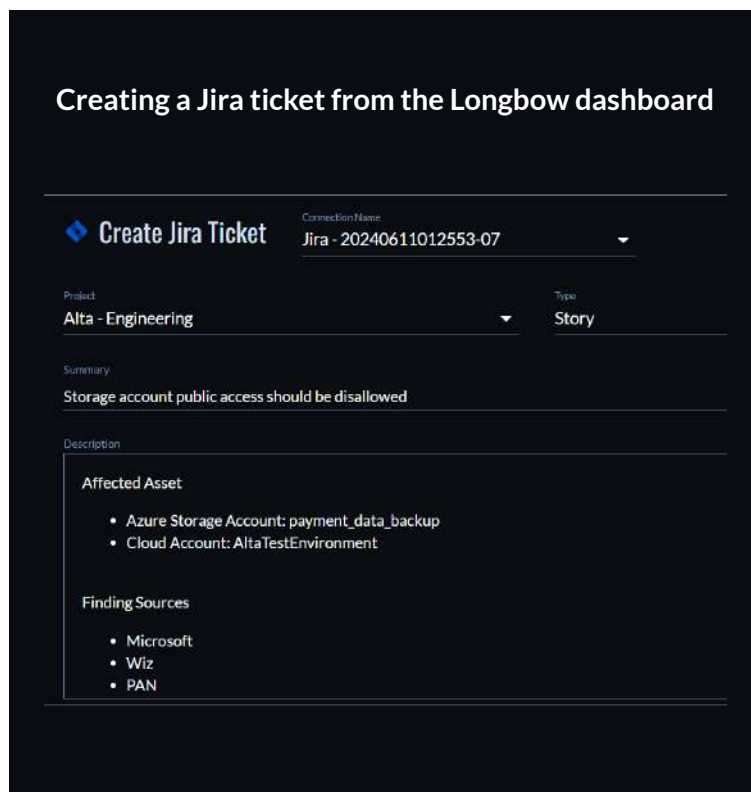
There's always more than one way to address a security issue, and Longbow was built with that understanding in mind. It provides multiple solutions for each issue and lets analysts pick the one that's best suited for the situation.

While full remediation will always be the default option, there may be times when an analyst would choose to isolate the risk instead. They may identify a vulnerability in an asset that's too important to take offline altogether. Instead, they can quarantine the server to help limit the risk facing that asset, while keeping it running.

Comprehensive Tickets on Demand

Longbow can integrate risk remediation into existing workflows. From an issue or solution in the tool, analysts can click a button to create a remediation ticket in Jira or ServiceNow. The ticket will automatically populate with all the information the developer needs to execute the selected solution.

When developers receive complete, targeted remediation tickets, they can spend less time filling in the blanks. This means they can process remediation tickets faster and get back to their day jobs sooner. They're also motivated by seeing the quantifiable impact their work has on risk remediation.



Case Study: Financial Company Revolutionizes its Approach to Risk

A global financial services provider adopted Longbow when its analysts were struggling to get a holistic view of risk. Like many large companies, it had a backlog of millions of issues it couldn't begin to remediate, with more added every day.

The company spent millions annually on detection tools, but analysts weren't getting value from those tools. The problem was not with the tools themselves; the analysts didn't have the necessary context to act on the findings the tools provided. After selecting Longbow, the utilization rate for detection tools skyrocketed, helping the company get more value from its existing security investments.

Benefits

Automated analysis: The customer went from not knowing root issues to automatically accessing root solutions at scale.

Optimized security portfolio: In addition to driving greater utilization of some tools, Longbow helped the customer identify other tools it could safely stop using. This helped cut costs significantly.

Improved collaboration: Internal teams all had their own view of risk. Longbow gave them a focal point to rally around for collaboration and communication, tearing down silos and increasing trust.

Results

10x increase in issues remediated per day

5.7 hours saved of remediation time per resource, per day

50% reduction in overall risk score



Summary

With the right ASPM tool, application security shifts from an overwhelming burden to a source of business acceleration. It's entirely possible for organizations to protect their cloud-native applications while also empowering analysts and developers to do their jobs better and faster.

By bringing together risk visibility, root-issue analysis and targeted remediation capabilities in the same place, Longbow cures alert fatigue and lays out the Best Next Actions required to secure software assets. Remediating risk consistently and proactively throughout the application pipeline helps businesses deliver secure applications at speed, giving them a key source of business agility and competitive advantage.

Learn More and Get Started Today



Are you tired of security platforms that don't look outside their own environment?



Do you need something better than tools that treat one fragment of the security landscape like the entire picture?



Could a vendor-agnostic platform help you make the most of the security tools you've already invested in?



Do you want that platform backed by experts that understand your challenges and act as an extension of your team to address them?

If you answered "yes" to these questions, then Longbow might be right for you. Our experts can demonstrate how Longbow enables an efficient, proactive approach to risk remediation throughout your application pipeline. Contact us today to schedule a briefing.

Veracode is a global leader in Application Risk Management for the AI era. Powered by trillions of lines of code scans and a proprietary AI-assisted remediation engine, the Veracode platform is trusted by organizations worldwide to build and maintain secure software from code creation to cloud deployment. Thousands of the world's leading development and security teams use Veracode every second of every day to get accurate, actionable visibility of exploitable risk, achieve real-time vulnerability remediation, and reduce their security debt at scale. Veracode is a multi-award-winning company offering capabilities to secure the entire software development life cycle, including Veracode Fix, Static Analysis, Dynamic Analysis, Software Composition Analysis, Container Security, Application Security Posture Management, and Penetration Testing.

Learn more at www.veracode.com, on the [Veracode blog](#), and on [LinkedIn](#) and [Twitter](#).

Copyright © 2024 Veracode, Inc. All rights reserved. Veracode is a registered trademark of Veracode, Inc. in the United States and may be registered in certain other jurisdictions. All other product names, brands or logos belong to their respective holders. All other trademarks cited herein are property of their respective owners.

VERACODE

Veracode Headquarters
65 Blue Sky Drive
Burlington, MA 01803

Phone 339.674.2500
Email hq@veracode.com

EMEA Headquarters
36 Queen Street
London, EC4R 1BN, United Kingdom

Phone +44 (0)20 3761 5501
Email emea@veracode.com