DevSecOps Disruptors and Innovators

April 27, 2023

Cybersecurity Research Analyst

Alex is a cyber security research analyst at Stratascale. His background in both research and practical security gives him a unique perspective on providing security with a risk-based approach. He focuses his expertise on emerging technologies, data-driven IT strategy, and tactical solutions to large security problems.

Submitted by Alex Banghart on 27, Apr 2023
DevSecOps Disruptors and Innovators

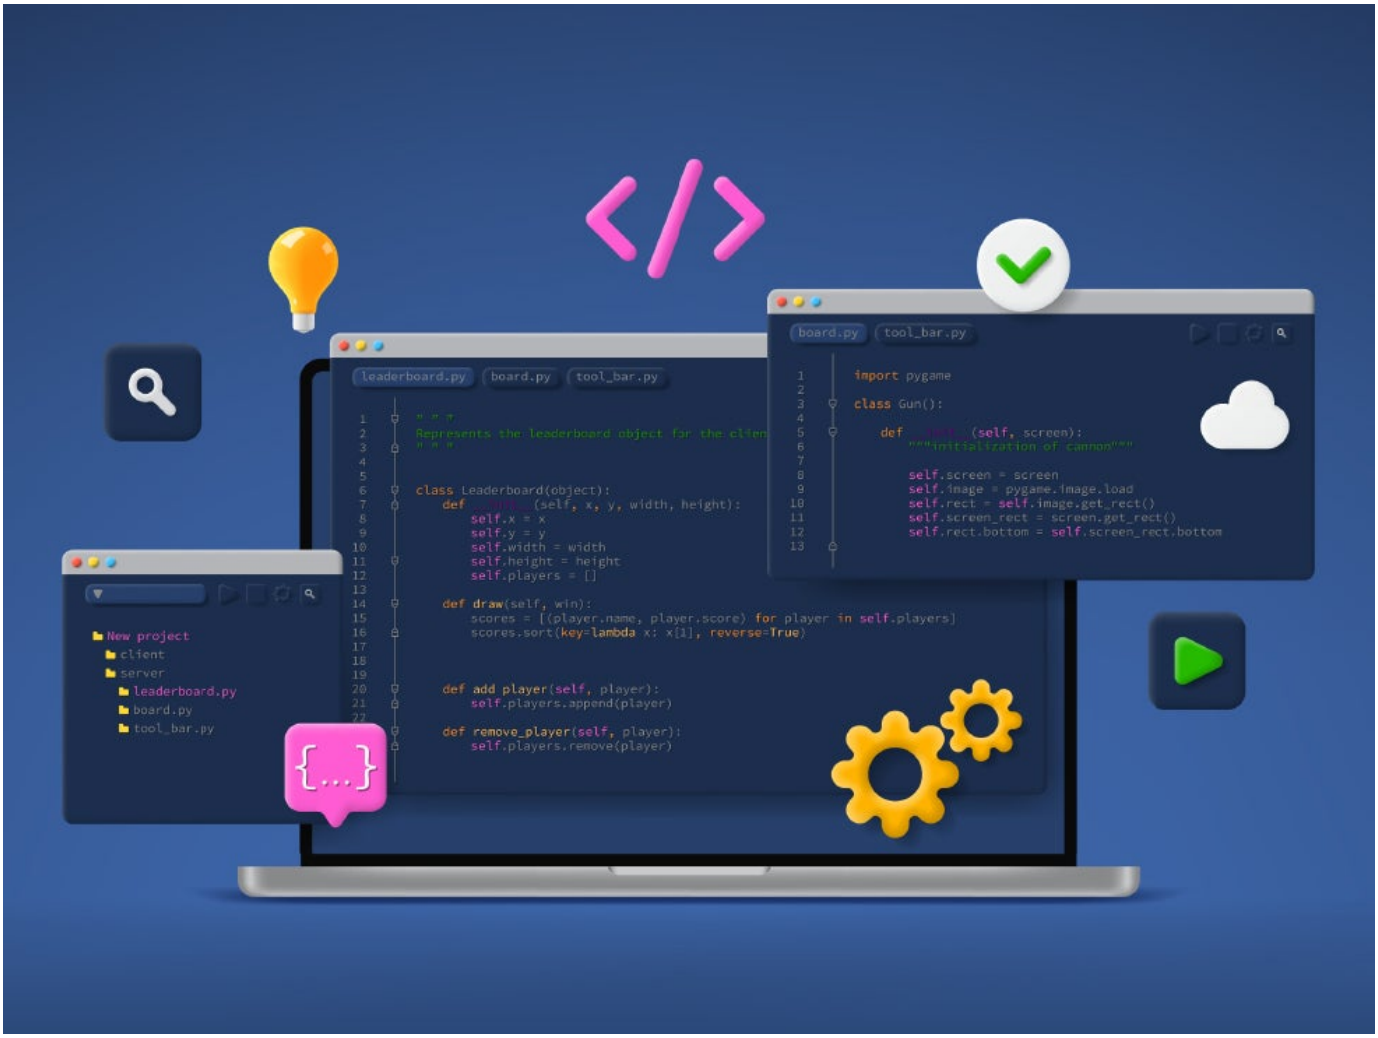

DevSecOps

INTRODUCTION, TRENDS, AND INNOVATION

What is DevSecOps?

DevSecOps means Development, Security, and Operations. The central idea of DevSecOps is that an integrated security approach should include collaboration throughout the entire development lifecycle. DevSecOps seamlessly integrates infrastructure and application security into Agile and DevOps tools and processes.

DevSecOps is best defined by IBM as “software, safer, sooner.” Fixing application and security issues is faster, cheaper, and simpler in the early stages of the development lifecycle.

Why it’s Important

GitLab recently conducted a study that shows organizations are not embracing an approach of security throughout the entire development lifecycle. Rather, they’re finding vulnerabilities extremely late in the development lifecycle. The farther in the development cycle you are, the more costly these vulnerabilities are in both remediation time and cost. According to the WhiteHat 2019 Application Security Statistics Report, an average of more than 50% of apps are vulnerable for organizations that haven’t adopted DevSecOps.

Development and security teams are still not aligned and have strenuous relations because accountability and ownership are not clearly defined. Many organizations still lack a fundamental understanding of who should be responsible for resolving security flaws.

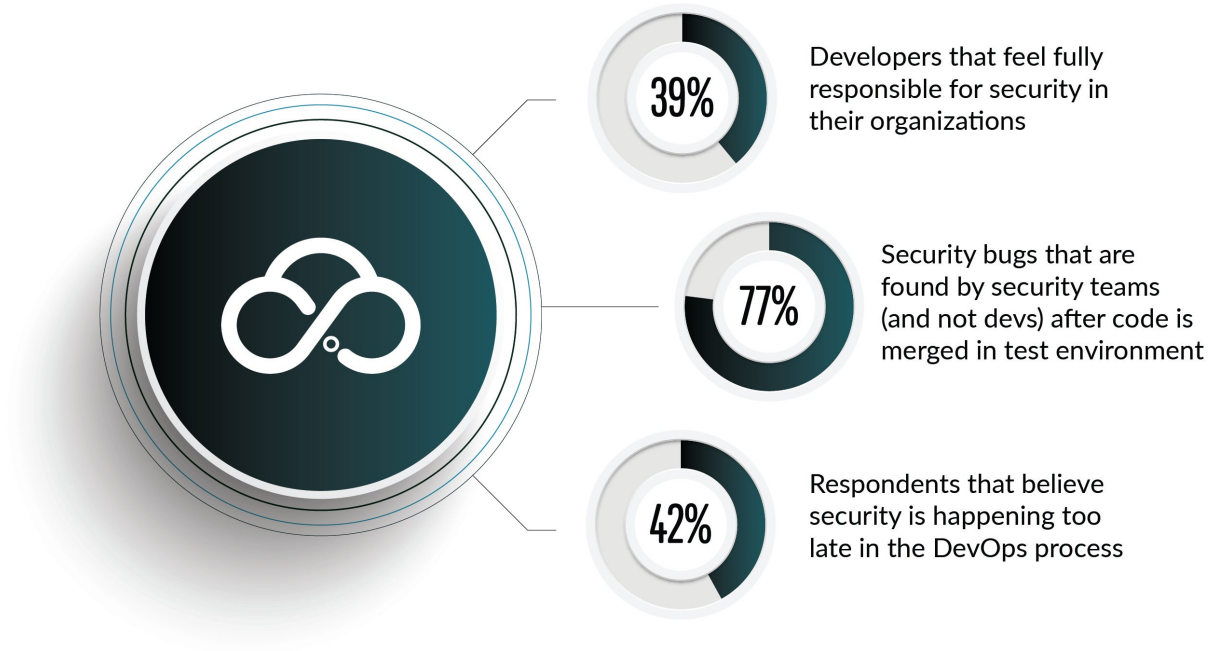

Source: “A Maturing DevSecOps Landscape,” GitLab, 2021

DevSecOps Lifecycle

A traditional DevOps lifecycle has the following steps: Plan, Develop, Build, Test, Release, Deliver, Deploy, Operate, Monitor, Feedback. This loop includes repeating areas that allow for continuous improvement within the cycle, as well as the constant agile development that most firms utilize.

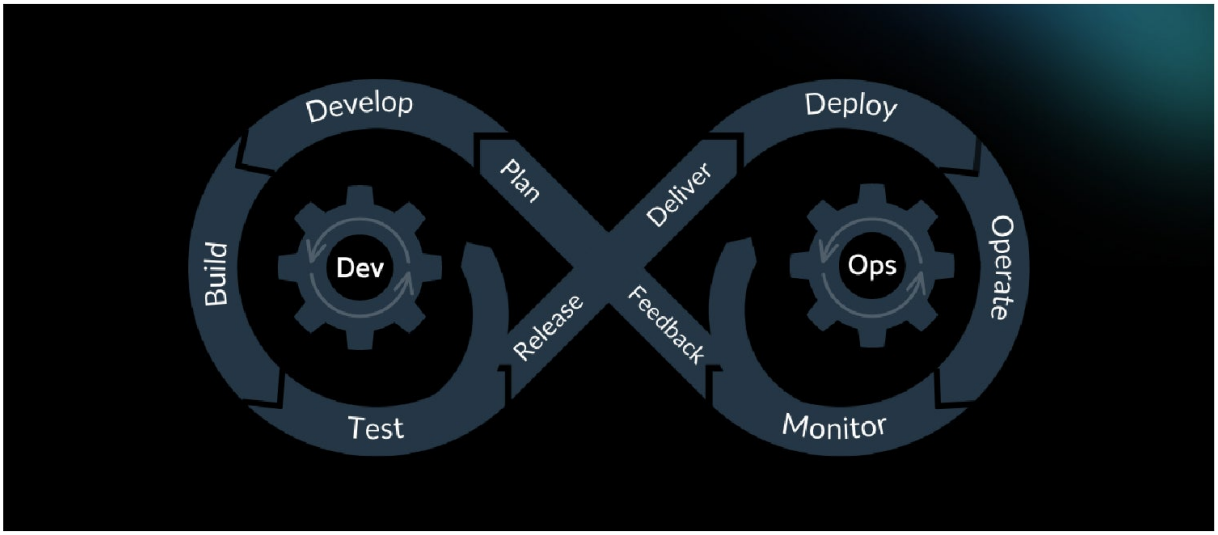

Source: DoD Enterprise DevSecOps v2.0

A shift to DevSecOps involves adding security measures at each of the steps. This ensures that vulnerabilities are found sooner and can be remediated more easily. The red shields in the diagram below illustrate the suggested security controls and where they fit into the lifecycle:

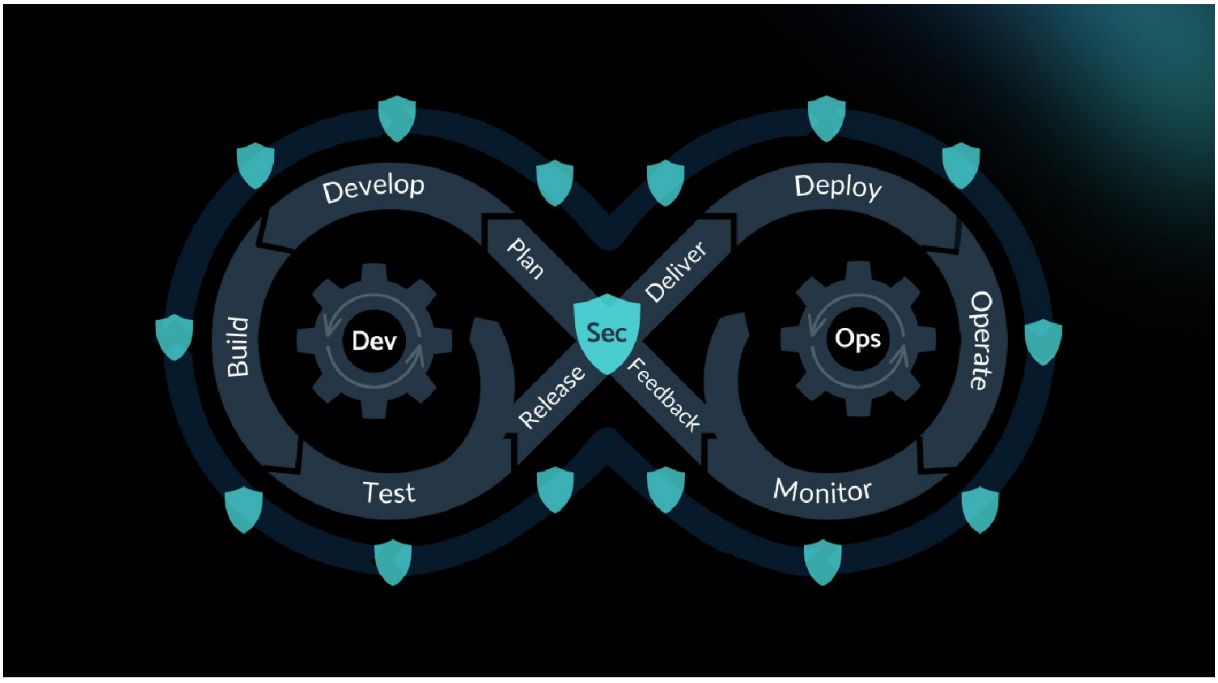

Source: DoD Enterprise DevSecOps v2.0

There are numerous benefits to implementing DevSecOps. After implementing a comprehensive DevSecOps program, the United States Department of the Air Force noted the following improvements:

- 106x faster lead time from development to deployment.
- 208x more frequent code deployments.
- 7x lower change failure rate.
- 22% less time on unplanned work/reworks.
- 50% less time remediating security issues.
- 2,604x MTTR (Mean Time to Recover).
- 40% reduced development costs.

- 44% more time focused on new capabilities vs. maintaining legacy code.

Source: DoD Enterprise DevSecOps Initiative & Platform One DAU Presentation

It’s also important to note that the costs to remediate fall substantially when vulnerabilities are found earlier in the lifecycle. By contrast, defects found after release cost 640x those remediated during coding.

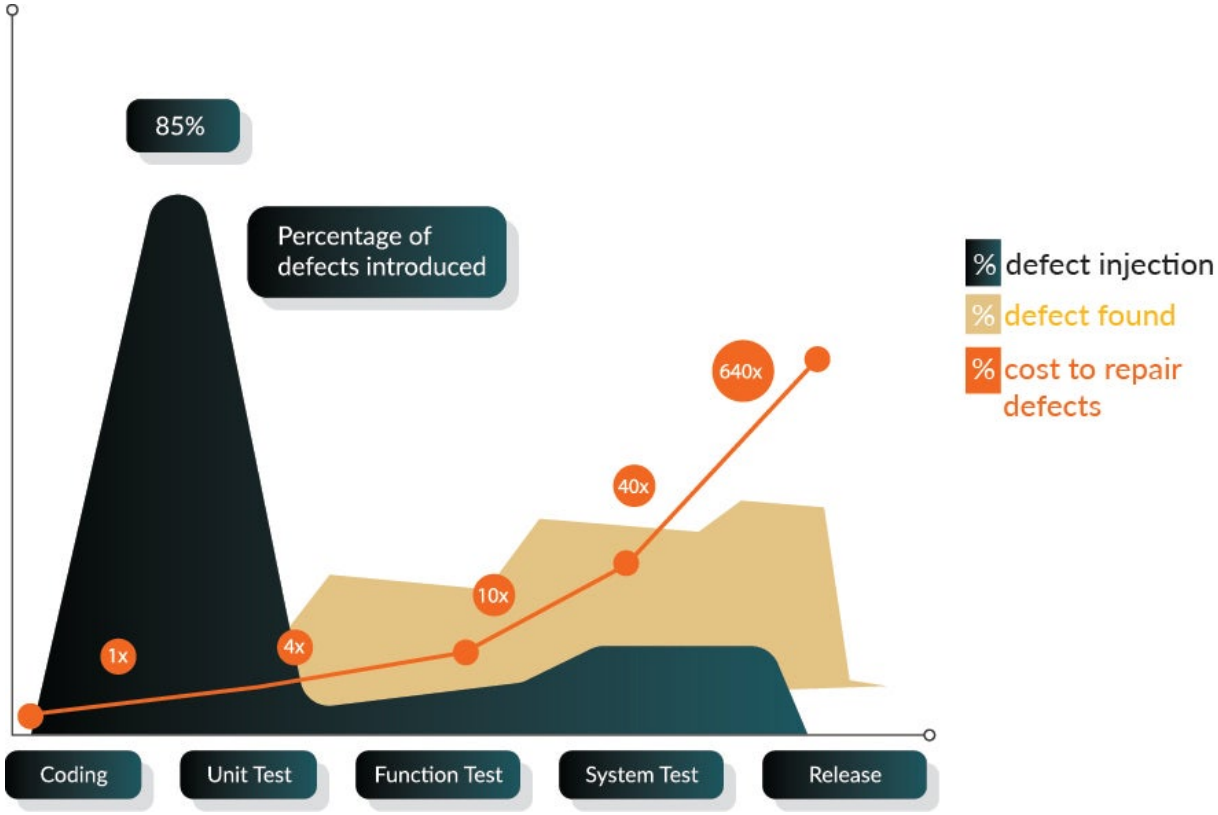

Source: Applied Software Measurement: Global Analysis of Productivity and Quality

Key Trends in DevSecOps

Infrastructure as Code (IaC) adoption

IaC means managing IT infrastructure in software instead of in traditional hardware. According to Gartner: “By 2023, 60% of organizations will use infrastructure automation tools as part of their DevOps toolchains, improving application deployment efficiency by 25%.”

IaC and codified infrastructure accelerates the ability to build a foundation for automation and testing by creating repeated and automatic software-driven processes. IaC allows for processes to be enshrined in the code, which reduces human error and ensures that these processes are repeatable, secure, automated, and efficient.

Defending against attacks originating from third-party code

Organizations find themselves increasingly exposed to attack vectors originating from vulnerable third-party code libraries. Log4J is a good example of why a robust DevSecOps program is so important for monitoring development and production code.

As Nicolas Fränkel points out: “Wise developers don’t reinvent the wheel: they use existing libraries and/or frameworks. From a security point of view, it means users of such third-party code should carefully audit it. We should look for flaws: both bugs and vulnerabilities.”

Root-cause analysis with AIOps

As the complexity of the cloud increases, it becomes impossible to manage these environments manually. It’s becoming increasingly important for DevSecOps teams to use automation to collect observable data and telemetry.

Broader business implementation comes from increasingly sophisticated AI algorithms and the increased speed at which AI can discover new data relationships. It’s important for the DevSecOps team to be able to identify the root cause of IT problems in real time and, in some cases, provide automatic remediation. This real-time analysis is important because the team integrates security scans to test code under development and to continually identify new security vulnerabilities in production environments.

GitOps framework: the new normal

GitOps is a framework of practices for managing infrastructure and application configuration using Git, an open-source version control system. This makes Git the only reliable source and control mechanism for dynamically creating, updating, and deleting system architectures.

GitOps uses pull requests to validate changes and automatically deploy to the system infrastructure. By concentrating these configurations in one place as much as possible, your team will have more control.

As more organizations move to continuous integration and continuous delivery (CI/CD), there will be more opportunities to implement GitOps. This approach allows teams to apply automation to testing, deployment, delivery, and governance.

Expansion of serverless architecture

Serverless computing is a cloud-based model for building applications and hosting that enables organizations to leverage resources as needed. Serverless architectures will interest teams who want to build, manage, and scale their applications without having to manage all the underlying infrastructure. The serverless model provides tools for cloud providers to manage their infrastructure and build applications in a modular fashion.

Enterprises can grow dynamically by delegating infrastructure management tasks to cloud providers. While serverless computing can be more costly than managing your on-premises infrastructure, organizations only pay for the resources they use. Serverless computing also improves disaster recovery and IT system resiliency, as cloud providers host the infrastructure.

Microservices expanding in place of the monolithic applications

Microservices are closely related to serverless computing. Rather than building monolithic applications that are time consuming and costly to develop and test, teams can add flexibility by splitting into independent units. This allows teams to remove the limitations of traditional application development.

By modularizing services, organizations can benefit from more versatile, step-by-step development that meets their core business needs. And when problems arise, microservices allow developers to tackle problems in a cohesive way rather than break the entire application. This style of modular application development helps the DevSecOps team maintain agility and flexibility while paying attention to code quality and security.

DevSecOps Disruptors and Innovators Coverage Overview

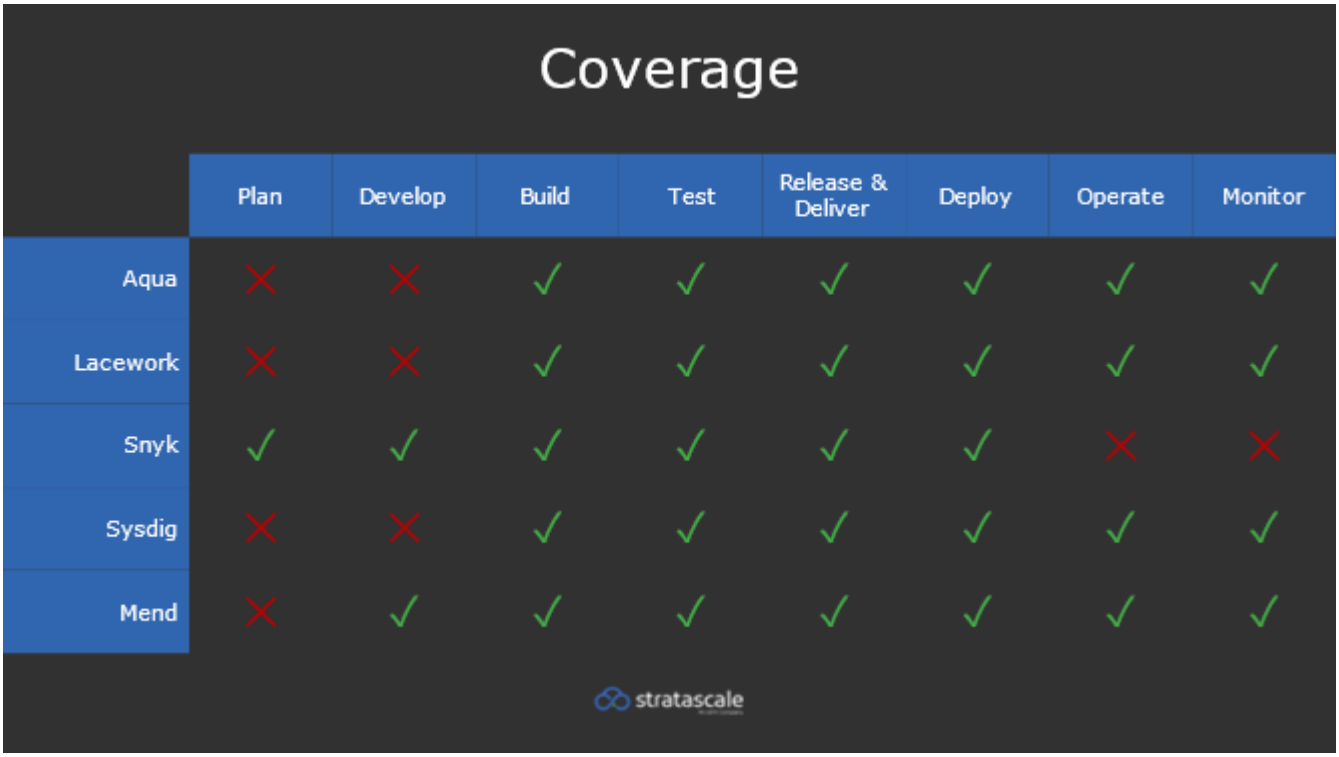

In-Depth Dives

Overview:

Aqua Security was founded in 2015, as containers and serverless technologies were just emerging. Aqua recognized that the dramatic change in application development and architecture required an equally dramatic shift in security.

Funding:

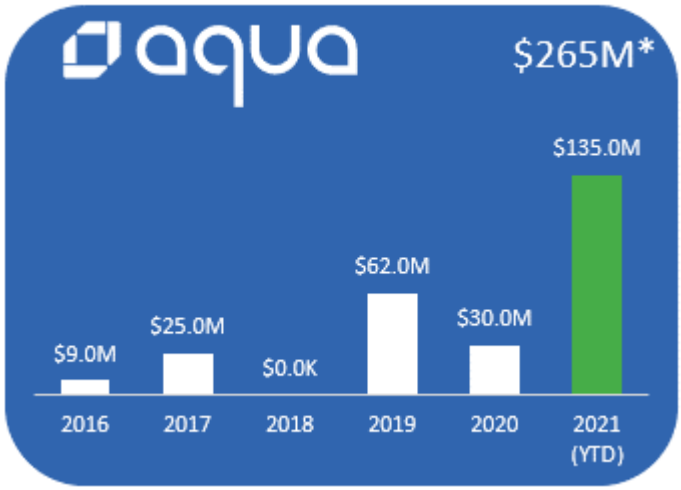

Capabilities:

Automate DevSecOps

- **Empower DevOps to fail fast and fix early:** Accelerate deployment and save costs by detecting and fixing security and compliance issues early.

- **Reduce the application attack surface:** Improve security by preventing vulnerable, untrusted code from being deployed, dramatically reducing the attack surface.

- **Regain control of your pipeline:** Going fast doesn’t mean running blind. Get full visibility into the security posture of your dev pipeline and avoid last-minute surprises.

Aqua enables organizations to automate secure development and deployment of applications in their DevOps pipelines by embedding comprehensive security testing and powerful policy-driven controls early on — all fully automated.

- **Image and function scanning:** Comprehensively scans container images and serverless functions for known vulnerabilities, embedded secrets, OSS licensing issues, malware, and configuration issues.

- **Continuous image assurance:** Prevents unapproved images and functions from being deployed in your environment, preempting operational errors, image sprawl, and rogue deployments — based on your policies.

- **CI/CD integration:** Natively integrates with Jenkins, Azure DevOps, Bamboo, GitLab, TeamCity, and more to scan images as they are built, providing actionable feedback to developers within their IDE.

- **Collaborative remediation:** Provides real-time actionable information on vulnerability and configuration remediation — feeds back to developers within their CI/CD tools, sent via Slack, or as a ticket in Jira.

- **Centralized secrets management:** Leverages your existing secrets vaults to securely deliver, update, and rotate secrets in containers with no container restart and no exposure outside the container, both in transit and at rest.

Modernize Security

For cloud-native apps, ensure the flow of “good code” with application-aware controls while automatically detecting and preventing suspicious activity.

- **Secure once, run anywhere:** Aqua’s security controls are consistently enforced across orchestrators, whether on-premises or in cloud environments.

- **Mitigate zero day and insider threats:** Least-privilege whitelisting detects and prevents any anomalous behavior, privilege escalation, or code injection.

- **Ensure business continuity:** Granular response ensures that only suspicious activities are blocked without disrupting application uptime.

Aqua replaces outdated signature-based approaches with modern controls that leverage the cloud-native principles of immutability, microservices, and portability. Using machine-learned behavioral whitelisting, integrity controls, and nano-

segmentation, Aqua makes applications more secure than ever possible before.

- **Continuous image assurance:** Prevents unvetted or unapproved images from running anywhere in your environment — based on known vulnerabilities, embedded secrets, OSS licensing, malware detection, and secure image configuration.

- **Image-to-container drift prevention:** Enforces container immutability and detects any unapproved changes to running containers by continuously comparing them to their originating images, including executables, privilege elevation, and image parameters.

- **Enforcing least privileges:** Uses machine learning to automatically profile container behavior, whitelisting runtime parameters such as system calls, file access, network access, and executables, improving isolation and preventing privilege escalation.

- **Granular monitoring and logging:** Monitors container, pod, node, and cluster activity to detect and report on all policy violations, run/stop events, and login events — all of which can be sent to your choice of SIEM (such as Splunk, ArcSight, and more).

- **Container-level application firewall:** Segments workloads by automatically creating dynamic firewall rules between container services, ensuring that only whitelisted connections are allowed and alerting on or blocking network traversal attempts.

Enforce regulatory compliance controls for PCI, HIPAA, GDPR, and beyond, across the cloud-native application lifecycle.

- **Manage, maintain, and prove compliance:** Collect detailed, continuous data on images, containers, orchestrators, and hosts, providing granular data streams sent to your choice of SIEM/GRC tool.

- **Enforce compliance across the stack:** Gain real-time visibility and control over the compliance posture of images, containers, hosts, and clusters, from development to production.

- **Easily implement regulatory requirements:** Use custom compliance checks in your pipeline and out-of-the-box compliance runtime templates for PCI-DSS, HIPAA, NIST, and more.

As organizations deploy applications using containers, serverless functions, and other cloud-native technologies, they're stepping into uncharted compliance territory. Aqua can help — with purpose-built cloud-native compliance controls, full visibility and auditing, and "compliant by default" templates to facilitate compliance and with less hassle.

- **Event logging and reporting:** Granular audit trails of access activity, scan events and coverage, Docker commands, container activity, secrets activity, and system events.

- **CIS-certified benchmark checks:** Automatically assess node configurations against Docker and K8s CIS benchmarks using Aqua OSS tools or scheduled testing and reporting.

- **Global compliance templates:** Pre-defined compliance runtime policies mapped to specific security standards, such as NIST, CIS, PCI, and HIPAA.

- **Full user accountability:** Enforce granular user accountability and controlled super-user permissions.

- **"Thin OS" host compliance:** Scan and monitor hosts for vulnerabilities, malware, and login activity, as well as discover and scan images stored on hosts.

- **Compliance enforcement controls:** Ensure only images and workloads that pass compliance checks are allowed to run in your environment.

Serverless Containers and Functions

Protect workloads in serverless containers and functions: Extend security across the cloud-native spectrum, enabling elastic, efficient deployment with security and compliance for services such as AWS (Amazon Web Services) Fargate and Lambda.

- **Modernize applications with confidence:** Reap the benefits of scalable, efficient microservices while controlling risk and maintaining regulatory compliance.

- **Centralize control over code deployment:** Manage and enforce security policies across cloud-native architectures and diverse deployment models.

- **Maintain and prove compliance:** Track, monitor, and audit vulnerabilities, sensitive data, and privileges in containers and serverless functions.

Aqua provides extensive lifecycle and runtime controls to secure serverless container services (AWS Fargate, Azure Container Instances) and serverless functions (such as AWS Lambda), allowing organizations to take a preventative approach to serverless computing, reducing the attack surface and reining in code sprawl.

- **Risk assessment and mitigation for serverless:** Comprehensively scan container images and serverless functions for known vulnerabilities, embedded secrets, OSS licensing, malware, configuration, and permissions.

- **Unified security management for CaaS and FaaS:** Manage hybrid cloud-native applications that use containers and functions from a single console, providing flexibility and control.

- **Embedding serverless security into the CI/CD pipeline:** Automatically test code as it is built into functions and container images, alerting on failing builds that don't conform to policy.

- **Injecting security into serverless containers:** Aqua’s MicroEnforcer injects security controls into the container, either during build or as part of a deployment task, to create self-protecting containers.

- **Protect AWS Lambda functions:** Assess the risk of AWS Lambda functions by discovering over-provisioned permissions and roles, embedded access credentials and keys, as well as code vulnerabilities. Prevent execution of functions that violate your organization’s security policy.

Hybrid Cloud and Multi-cloud

Security built in: Make security an enabler of cloud migration, hybrid-cloud, and multi-cloud deployments, with persistent controls that follow your workloads wherever they run.

- **Secure once, run anywhere:** Run and move applications across Azure, AWS, Google Cloud, IBM Cloud and more, with unified, consistent security.

- **Scale multi-tenant environments as needed:** Monitor, manage, and enforce security policies on multiple clusters, for multiple tenants, either on-premises or in the cloud.

- **Future-proof your investment:** Avoid cloud provider lock-in and ensure the flexibility of future deployment with no need to reconfigure your security controls.

Aqua provides full lifecycle security controls for containers and serverless functions, coupled with compatibility and integrations with all orchestrators and cloud providers. Aqua enables organizations to keep their cloud strategies flexible while

ensuring uniform security and compliance enforcement across environments.

- **Integrated with your cloud:** Aqua has partnered with AWS, Azure, GCP (Google Cloud Platform), IBM Cloud, and VMware to ensure that its software not only runs smoothly on every environment but leverages native image registries, orchestration, IAM, monitoring, and log collection.

- **Multi-cluster segmentation:** Aqua's application-contextual container firewall allows you to segment workloads within the same environment or across clusters and clouds, preventing the spread of attacks while enabling you to deploy with maximum efficiency.

- **Built for multi-tenant environments:** Manage multiple team deployments or multiple customer tenancies from a central console. Maintain separation of data and access, ensuring complete isolation between tenants.

- **Scans images and functions:** Comprehensively scan container images and serverless functions for known vulnerabilities, embedded secrets, OSS licensing issues, malware, and configuration issues — before they are deployed.

- **Protecting serverless:** Aqua protects containers both on VMs (Aqua Enforcer) or in serverless environments such as AWS Fargate and Azure Container Instance (Aqua MicroEnforcer) — controlled from a single console with consistent policy enforcement.

Overview:

Lacework is a software company with a specialized unified cloud security platform. As described by CEO Dan Hubbard: "We started Lacework with a group of the most

insightful thinkers in cloud security. We developed a concept that cloud security could be provided as an end-to-end experience with better context, more intelligence, and more sophisticated threat detection. Our founding team had a vision for applying machine learning to the behaviors of resources, users, and all activity in the cloud, which would provide a more accurate picture of where actual threats existed. This approach provides customers with more contextual, accurate data and alerts.”

Funding:

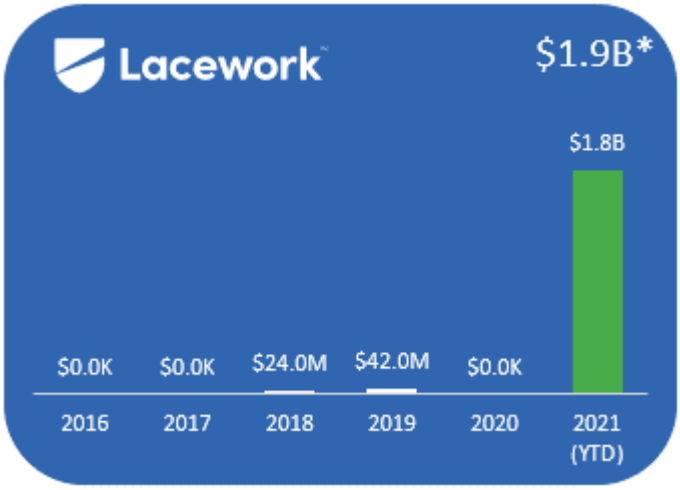

Capabilities:

Visibility and Analysis for Compliance in Multi-cloud Environments

Complexity is an enemy of security; a unified view is essential to simplify the complexity of having multiple configurations. Lacework does this across AWS, GCP, and Azure by bringing multiple clouds into one portal. This means no logging into different disparate tools to evaluate your stance. It’s a single pane of glass to audit all your cloud platform configurations.

As configurations change, Lacework will monitor and alert any time a configuration falls out of compliance. This ensures that security and compliance teams immediately become aware of issues, so they can be fixed before data and cloud resources are compromised.

Lacework delivers deep visibility for configurations across all of an enterprise’s cloud accounts and workloads so organizations can ensure compliance with industry, governmental, and institutional standards. Operating on multiple cloud platforms can increase the threat vector of the overall infrastructure and add complexity to an already challenging task. Lacework operates as a comprehensive, centralized solution to identify, analyze, and alert on configuration issues.

Identify Configuration Issues:

- Find Identity and Access Management (IAM) vulnerabilities, including root account, password requirements, and usage of MFA.
- Check for logging best practices — enable log files across regions, and enable that log files are validated and encrypted.
- Monitor critical account activity such as unauthorized API calls and use of the management console for unauthorized purposes.
- Confirm secure network configurations, including limiting access to vulnerable ports, enforcing “least access” privileges, and checking for the use of flow logging.

Track Configuration Changes Continuously:

- Daily re-audit to maintain compliance and protection.
- Monitor account activity for abnormal activity, even when that activity is technical authorized.
- Receive customizable alerts when items change from compliant to non-compliant.

Ongoing Monitoring of Activity:

- Detection and alerting of activity on all cloud platform resources, such as new activity in a region, activation of new services, or changes to access control lists.

- Changes to users, roles, or access policies.

- Access or customer master key tampering.

- Reduce alert fatigue with customizable alerts and reports that eliminate repetitive or irrelevant results.

Configuration Compliance Management

Lacework checks across the industry-accepted CIS Benchmark for secure configurations for cloud accounts and workloads. Additionally, Lacework includes supplemental checks for common compliance frameworks like PCI-DSS, SOC 2, HIPAA, and others.

Using Lacework, compliance and security teams have continuous analysis and historical reporting available so they can understand what is being checked, where problems exist, an analysis of the problem, and the steps needed to remediate the misconfiguration. The product supplies link directly to the resources in question to reduce the time to remediate.

The Lacework configuration compliance solution is built to detect behavioral anomalies, so even if configurations do meet required standards, unauthorized use or abnormal activity is detected and alerted on. This ensures that organizations are aware of issues that might go undetected by solutions that only identify non-conforming compliance rules.

The Power of Polygraph for Configuration Compliance

Lacework’s foundation is Polygraph, a deep temporal baseline built from collecting high fidelity machine/process/users interactions over a period of time. The Polygraph is used to detect anomalies, generate appropriate alerts, and provide a tool for users to investigate and triage issues.

Fundamentally, Lacework's Polygraph technology dynamically develops a behavioral and communication model of your services and infrastructure. The model understands natural hierarchies (processes, containers, pods, machines, etc.) and aggregates them to develop behavioral models. A behavioral model is, in some sense, the essence of how a customer's infrastructure operates. With this model, Polygraph monitors your infrastructure for activities that fall outside the model. In addition, Polygraph continually updates its models as your data center behavior changes.

Spot IaaS account configurations that violate compliance & security best practices the could put your company at risk with Lacework's comprehensive configuration compliance monitoring tools for enterprise DevOps teams.

Cloud Workload Protection

Visibility and Analysis for Container and Cloud Workloads

Lacework's cloud workload security platform provides visibility to all processes and applications within an organization's cloud workloads and container environments. The breadth and depth of workload visibility provided by Lacework help security teams detect vulnerabilities and then utilize its machine learning analysis to identify anomalous behavior that poses threats.

Traditional security solutions rely on network logs and the firewall rules to identify potential risks, but those approaches required a manual effort and could not keep pace with the speed of modern cloud deployment methodologies. Lacework was built specifically to deliver contextual data about cloud events — every update, configuration change, access point, and a million other activities that might represent potential threats.

Lacework tracks all machine/process communications, the users associated with those processes, and the amount of data that was transferred between processes during a given time. This deep level of detail allows teams to save time on their investigations because all the relevant information is on one workload security platform.

Automated workload intrusion detection: Lacework's cloud workload protection security platform is fully automated with no rule-writing required. Using

sophisticated machine learning, Lacework's workload security platform learns what constitutes normal behavior versus those that indicate potentially malicious activity.

Examples of such anomalous activities are when a user launches a new unknown application, when an application connects to a suspicious endpoint, or when privileges are unexpectedly escalated. When Lacework identifies a potential threat, a contextual alert is generated with relevant data to allow users to investigate and triage the issue within your cloud workload environment.

Lacework's Automated Cloud Workload Security Approach Provides:

- No missed events: Lacework will always alert you on new activity so that you're given a chance to investigate any behavior within your workload environment that could potentially be malicious.

- Low alert noise: Lacework will only alert you on what is new or anomalous, preventing alert fatigue within your organization.

- Simple operations and maintenance: Automated workload detection means no writing and maintaining error-prone rules. With Lacework, you will not need to constantly maintain rules, allowing you to focus on securing your environment.

Cloud workload protection at scale and the speed of business: The modern cloud infrastructure allows organizations to deploy, scale, and configure their infrastructure faster than ever. This ability to automate and operate at DevOps speed poses a challenge to traditional security approaches.

Lacework's approach is to automate workload security with detection of threats and anomalies and to provide investigative insights that humans can understand. Lacework's cloud workload security solution supports public clouds AWS, GCP, Azure, and computer hosts and containers.

Multi-Cloud Security

Visibility and Detection of Misconfigurations and Account Vulnerabilities

Lacework provides comprehensive cloud account security for AWS, Azure, and GCP accounts, including insights about configuration changes that could lead to threats. At the console level of a cloud environment, an organization can inadvertently apply misconfigurations that could leak data or open up an easy attack surface to a hacker. With continuous updates and broad administrative access within cloud environments, account changes are normal. Yet, with increased activity comes increased vulnerability.

Through API integration between accounts, Lacework looks at all the security-relevant configurations and identifies where the organization is passing or failing certain account security best practices for these particular configurations. These checks are run continuously, and security teams receive automated alerts about any configuration changes that violate security compliance. Among the myriad issues that Lacework is looking for, it is able to identify things such as:

- S3 buckets in AWS that are misconfigured and left publicly open.
- Security groups allowing unrestricted access to SSH.
- IAM users that don’t have MFA enabled.
- Security groups that are misconfigured.
- New regions spun up specifically for Bitcoin mining.

Data from the cloud accounts are ingested, and Lacework applies machine learning to logs to generate high-fidelity alerts on any behaviors or events that could be an indicator of compromise at the account resource level. Lacework also proactively alerts on any account security misconfigurations at the time they occur.

- **Identify Configuration Issues:**

- Find Identity and Access Management (IAM) vulnerabilities, including root account, password requirements, and usage of MFA.
- Check for logging best practices enable log files across regions, and enable that log files are validated and encrypted.
- Monitor critical account activity such as unauthorized API calls and use of the management console for unauthorized purposes.
- Confirm secure network configurations, including limiting access to vulnerable ports, enforcing “least access” privileges, and checking for the use of flow logging.

- **Track Configuration Changes Continuously:**
  - Daily re-audit to maintain compliance and protection.
  - Monitor account activity for abnormal activity, even when that activity is technically authorized.
  - Receive customizable alerts when items change from compliant to non-compliant.

- **Ongoing Monitoring of Activity:**
  - Detection and alerting of activity on all cloud platform resources, such as new activity in a region, activation of new services, or changes to access control lists.
  - Changes to cloud account users, roles, or access policies.

- Access or customer master key tampering.

- Reduce alert fatigue with customizable alerts and reports that eliminate repetitive or irrelevant results.

Container Security

Container Monitoring, Visibility, and Visualization of Deployed Containers

Lacework delivers native container security support, reduces the attack surfaces, and effectively detects threats in a containerized environment. Their cloud container security monitoring platform automatically discovers every container across a user's environment and clusters them based on different behaviors. Lacework visualizes your containerized applications in real time, providing a clear understanding of communications, launches, and other cloud runtime behaviors.

Host security: Containers can be thought of as lightweight virtual machines with much leaner system requirements. Virtualization emulates the guest system, translating every instruction between the guest and host. Containers, on the other hand, share the kernel and execute instructions on the host directly. This implies that the main attack surface is still the host, as it is shared across containers and any compromise at the host level can compromise all containers. The other challenge is that not all services are run in the container, as there is a long list of OS level and management services which run outside containers and are part of the attack surface.

Cloud container security using behavioral patterns: Lacework's container security platform discovers every container and uses machine learning to establish each container's normal behavioral patterns. Lacework then places containers with similar behaviors into a single, logical cluster — called a "Polygraph" — each with a baseline of expected characteristics and behaviors.

Clustering containers based on behavior dramatically simplifies the visualization of a containerized cloud in a Lacework Polygraph by representing dozens or even hundreds of similar containers as a single item. This means new containers or configuration changes do not generate alerts as long as behaviors stay within the

expected baseline. This also reduces notification clutter by delivering high-precision alerts only once per container cluster. Lacework's container security platform creates multiple types of Polygraphs based on different behavioral categories:

- The **communication Polygraph** baselines the communication pattern between different container clusters.
- The **launch Polygraph** baselines the launch behavior of the container clusters.
- The **privilege change Polygraph** baselines the user privilege changes within the containers.
- The **user activity Polygraph** baselines user behavior.

Continuous container security monitoring for compliance: Lacework's cloud container security monitoring platform brings multi-cloud checks into one dashboard by continuously monitoring configuration changes and API activity for containers across AWS, Azure, and GCP platforms. CIS benchmark scans are performed during container image development and container deployments.

Lacework's security platform also includes supplemental checks based on industry best practices and common compliance frameworks like PCI-DSS, SOC 2, HIPAA, NIST, and more. Unlike most other container security solutions that only identify non-conforming compliance rules, Lacework goes a step further and alerts your team about any behavioral anomalies — even when the associated configurations meet the required standards.

From automated threat detection to compliance, Lacework's offers a comprehensive approach to container security that ensures nothing is left unprotected, which point solutions can't guarantee.

The power of Polygraph: Lacework's foundation for securing containers is Polygraph, where a deep temporal baseline is built from collecting high fidelity machine, processes, and user interactions over a period of time. The Polygraph is used to detect anomalies, generate appropriate alerts, and provide a tool for users

to investigate and triage issues in their cloud container environments.

Kubernetes Security

Application, Visibility, Threat Detection, and Forensics for Kubernetes

Lacework's Kubernetes security solution provides comprehensive threat detection for dashboards, pods, management nodes, and clusters, in addition to end-to-end security for their public cloud infrastructure workloads, accounts, and containers.

With the rapid adoption of Kubernetes for application and infrastructure orchestration, there's a corresponding increase in the risk associated with data exposure and vulnerabilities throughout the application lifecycle. Without proper detection of threats, organizations could unwittingly be granting unauthorized access to Kubernetes clusters, applications, and customer data. Lacework's Kubernetes security platform identifies the risks and threats for Kubernetes-deployed infrastructures, including publicly exposed and unsecured API servers and management consoles.

Lacework was among the first cloud security vendors to highlight the need for rigorous container security. The company's original research was published earlier this year in a report titled "Containers at Risk: A Review of 21,000 Cloud Environments."

Application visibility: Lacework provides deep visibility into your Kubernetes deployment. This includes high-level dashboards of your clusters, pods, nodes, and namespaces combined with application-level communication between all these at the application, process, and network layers.

Threat detection for Kubernetes: Backed by the power of Lacework's Polygraph technology, this security solution for Kubernetes includes detection of both risks and threats that may be specifically designed to breach a vulnerability within Kubernetes, a possible misconfiguration, or a threat that can affect your infrastructure by installing malicious code onto one of your containers.

The Lacework Polygraph is designed to detect both known and unknown threats that affect Kubernetes environments through the detection of IOC's and Lacework's behavioral analysis and machine learning classification. Risks and threats are visible within the Lacework dashboard, are ranked by risk severity, and can be delivered

through the most common modern methods such as a Slack channel or a Jira ticket.

Forensics for Kubernetes: Whether you're triaging an alert or digging into deep details around the cause and effect of a change, Lacework's security platform for Kubernetes has all the information.

This SaaS service allows you to go back in time and look at all related events across your Kubernetes infrastructure that may have caused a breach or exposed you to an unknown risk. Detailed information about your containers, your applications, and your infrastructure are all available and include information related to Kubernetes such as pods, nodes, labels, namespaces, and all network information. All this information is available both within the UI and from its API.

Lastly, Lacework's Kubernetes security solution creates hourly Polygraphs that can demonstrate the change of relationships and events over time. This is a critical tool for understanding and triaging your events.

Overview:

Snyk helps you use open source and stay secure. Continuously find and fix vulnerabilities for npm, Maven, NuGet, RubyGems, PyPI, and much more.

Open source is awesome for boosting your productivity. However, taking code written by others, often with little to no vetting of its security pedigree, can put your

application at risk. Snyk enables you to find, and more importantly fix, known vulnerabilities in your open source. And it’s built by the best developers and security researchers in the space.

Funding:

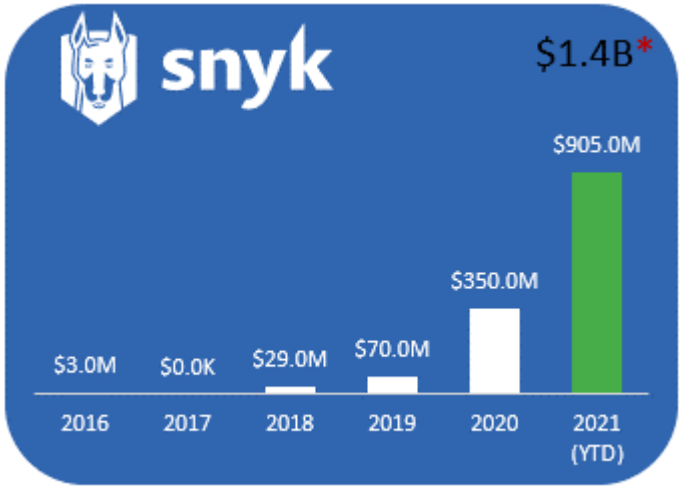

Offerings:

Developers and DevOps

Snyk Open Source: Automatically detect vulnerabilities and accelerate fixing throughout your development process.

Find Vulnerabilities in Your Open-source Dependencies Early and Across the SDLC:

- **Coding and CLI:** Detect vulnerable dependencies as you code in your IDE or CLI to avoid future fixing efforts and save development time.
- **Code management:** Scan pull requests before merging. Test your projects directly from the repository and monitor them daily for new vulnerabilities.
- **CI/CD:** Prevent new vulnerabilities from passing through the build process by adding an automated Snyk test to your CI/CD.

- **Production environment:** Test your running environment to verify there is no exposure to existing vulnerabilities and monitor for newly disclosed vulnerabilities.

Analyze easily and make data-driven security decisions with Dependency tree view: Accelerate your triaging process with Snyk’s dependency path analysis, which allows you to understand the dependency path through which transitive vulnerabilities were introduced.

- **Dependency health:** Broaden your security coverage by identifying whether or not there is a risk associated with dependencies within your open-source libraries.

- **Runtime prioritization:** Prioritize your fixes based on an analysis of the vulnerabilities that are called at application runtime and bear a higher risk.

- **Exploitability data:** Use exploitability indicators to identify those that are easy for attackers to weaponize.

- **Accuracy control for minimizing false positives:** Receive high-accuracy alerts that are verified and qualified by Snyk’s dedicated security research team.

Fix Quickly to Reduce Exposure:

- **Minimal fix required:** Snyk identifies the minimal upgrade required in order to clear a vulnerability and notifies when there is a risk of breaking the code.

- **Transitive dependency fix:** Accelerate triaging of transitive vulnerabilities with Snyk’s fix suggestions for the direct dependency.

- **Fix pull request:** Automate fixing with a one-click fix pull request populated with the required upgrades and patches.

- **Precision patches:** When upgrading is too disruptive (or not available), fix quickly and precisely with Snyk's proprietary patches (developed in collaboration with the maintainer).

Monitor Continuously to Maintain Your Code Security Level:

- **Newly disclosed vulnerabilities:** Automatically monitor your projects and deployed code and get notifications whenever new vulnerabilities are disclosed.

- **Gating new dependencies:** Prevent new vulnerabilities from passing through any stage of the development process.

- **Reporting:** Understand the state of all your security vulnerabilities and license issues in one place. Monitor how your team addresses issues with an auditable inventory of dependencies used in your projects.

Alerts and notifications: Get updates on newly identified vulnerabilities through preferred channels including Slack, Jira, email, and more.

Automate Open-Source Security Management and Governance, at Scale:

- **Security policies:** Automatically prioritize and deprioritize vulnerabilities using fully customizable security rules.

- **License policies:** Create, customize, and manage license compliance policies across your organization with Snyk License Compliance Management.

- **Project tags and attributes:** Manage and control your projects more easily by assigning them with built-in attributes or your own customized tags.

- **API:** Tune security automation to fit into your existing development workflows and ensure both developer experience and consistent platform governance.

Vulnerability Database

Snyk Intel Vulnerability Database: Empower agile development teams with trusted data and insights to rapidly secure open-source code. Snyk Intel Vulnerability Database is the most advanced and accurate open-source vulnerability database in the industry. Continuously curated by an experienced Security Research Team, Snyk's Open Source Vulnerability Database maintains its high standards, which enable your teams to be optimally efficient at containing open-source security issues while maintaining their focus on development.

Comprehensive Security Coverage:

- **Best coverage:** The Snyk Intel Vulnerability Database goes far beyond CVE vulnerabilities and other public databases, including many additional non-CVE vulnerabilities derived from numerous sources More vulnerabilities covered than the next largest publicly available commercial vulnerability database.

- **Know sooner:** Snyk exposes many vulnerabilities before they are added to public databases. of the JavaScript vulnerabilities in NVD were added first to the Snyk open-source Vulnerability Database.

- **Detect Faster:** Because Snyk exposes many vulnerabilities before other sources, you can detect and correct issues faster.

Vulnerability Database Methodologies:

- **Enriched data from numerous vulnerability databases:** These include CVE, NVD, and more. Data derived from these resources are analyzed, tested, and enriched before being included in the database.

- **Dedicated proprietary research for new vulnerabilities:** Snyk's dedicated security team is focused on uncovering severe vulnerabilities in key components. A recent disclosure by its team is Zip-Slip. See more examples in the footnote below.

- **Threat intelligence systems:** Listen to chatter on security bulletins, Jira boards, Github commits, and more to automatically identify vulnerabilities not yet reported. Previously surfaced vulnerabilities from this source include Apache Airflow and Marked.

- **Community relationship:** Snyk collaborates with the community and operates bug bounties for new disclosures. This activity results in hundreds of community disclosures, such as f2e-server.

- **Collaboration with academia:** The team partners with PhD academia labs such as Berkeley, Virginia Tech, and Waterloo to exchange tools, methods, and data. Snyk then exclusively discloses findings.

Team of security experts: Snyk's security database is managed by a team of experts, researchers, and analysts, ensuring the database maintains a high level of accuracy with a low false-positive rate. Snyk database authority was validated by the leading security institutes. Snyk was granted a CVE numbering authority, is a member of the Node foundation security membership group, and is a contributor member of OWASP. The team is headed by Snyk's co-founder, Danny Grander, a veteran security researcher. Previously, Danny built cyber solutions for government agencies, led vulnerabilities research, and managed research and development teams. Danny is a competitor and frequent winner of CTF at DefCon, CCC CTF, and Google CTF.

Curated, Enriched and Actionable Content:

Hand-curated Content and Enriched Metadata:

- Vulnerability description: hand-curated content and summaries, including code snippets where applicable.
- All items in the database are analyzed and tested for their accuracy (version ranges, vulnerable method, and more).
- CVSS score and vector assigned to 100% of vulnerabilities.

Triage Support:

- **Vulnerable functions called in runtime:** For issue prioritization, Snyk is able to alert when a vulnerable function is being called during the runtime of the application.
- **Exploitability:** Snyk indicates when a vulnerability has a published proof of concept of how it can be exploited. Published exploit code serves as a good indicator of exploitability because it enables attackers to easily weaponize a vulnerability.

Enterprise Features

Snyk for Enterprise Security: Regain visibility into open-source risk and empower your developers to address it.

Find Vulnerabilities:

- Map the full application dependency tree.
- Find vulnerabilities in all open-source dependencies.
- Use CLI, integrations, or API to add projects to be tested.

- Continuously test for newly disclosed vulnerabilities.
- Dependencies are tested against Snyk’s comprehensive vulnerability database.

Reports:

- **Visibility:** View the state of all your security vulnerabilities and license issues in one place, with an overview optimized for displaying on a big screen.
- **Accountability:** See how quickly your team addresses issues.
- **Auditable:** Get an inventory of all the dependencies used in your projects that you can export as a CSV.

Licenses:

- **Review compliance:** Get an inventory of the licenses used in your projects and their dependencies.
- **Stay compliant:** Prevent problematic licenses from being introduced when a GitHub pull request is raised.
- **Custom Policy:** Create a bespoke license policy for your organization. Set the severity level of specific licenses and get alerted when a project includes a problematic license.

Groups:

- **Team flexibility:** Set up areas for your teams to focus on the projects relevant to them.

- **Superpowered reports:** Get an overview of your vulnerability status across all your organizations.

- **Quick filters:** Save filters in your reports so you can quickly access the data you care about.

Issue Tracking:

- **Issue lifecycle management:** Integrate Snyk with your issue management tool of choice, and manage Snyk issues with your standard process.

- **Jira integration:** Triage security vulnerabilities and license issues in Jira with your team.

- **Webhook integration:** Connect Snyk to any notification system using a generic webhook integration.

Test a local project: Snyk's CLI helps you find and fix known vulnerabilities in your dependencies, both ad hoc and as part of your CI (Build) system. The Snyk CLI requires you to authenticate with your account before using it. It supports Node.js, Ruby, Python, Java, Scala, Go, and .NET.

Overview:

Sysdig enables companies to confidently run cloud-native workloads in production. With the Sysdig Secure DevOps Platform, cloud teams embed security, maximize availability, and validate compliance.

Funding:

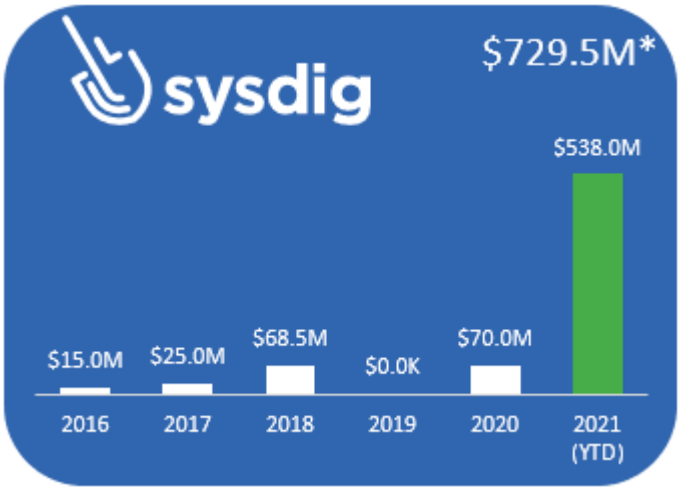

Capabilities:

Embed Security, Compliance, and Performance into DevOps Workflows

The Sysdig Secure DevOps Platform converges security and compliance with performance and capacity monitoring to create a secure DevOps workflow. It uses the same data to monitor and secure, so you can correlate system activity with Kubernetes services. This enables you to identify where a problem occurred and why it happened — and you can use this single source of truth to investigate and troubleshoot performance and security issues.

Two Core Products, One Secure DevOps Workflow:

- **Sysdig Secure:** Efficiently resolve vulnerabilities, block threats at runtime, and respond to incidents quickly — and be confident about your Kubernetes security.

- **Sysdig Monitor:** Deliver the performance and availability your users expect via context-based monitoring — and manage the risk, health, and performance of your microservices.

You can’t secure what you cannot see: Existing DevOps tools and point solutions simply don’t provide adequate security or visibility. Secure DevOps for cloud-native requires purpose-built tools.

- **Embed security:** Detect vulnerabilities and misconfigurations with a single workflow. Block threats without impacting performance. Conduct forensics after the container is gone.

- **Maximize availability:** Prevent issues by monitoring performance and capacity. Accelerate troubleshooting with a single source of truth. Scale Prometheus monitoring across clusters and cloud policies.

- **Validate compliance:** Verify that configurations meet CIS best practices. Ensure application compliance with NIST, PCI. Accelerate audits by correlating Kubernetes activity.

Critical ecosystem component: The Sysdig Secure DevOps Platform provides cloud-native security and monitoring, delivering visibility and control for secure operations. Open by design, it supports and adds value to a wide range of leading tools, solutions, and cloud services from multiple providers.

Secure DevOps Across the Cloud-Native Lifecycle:

- **Build:** Scan and block vulnerable images and enforce best practices pre-production.

- **Run:** Block threats, enforce compliance, and monitor application performance.

- **Respond:** Proactively alert on incidents, reduce MTTR with forensics, and capture detailed audit records.

Manage Cloud Security Risk

With daily updates in a cloud-native environment, it is easy for new vulnerabilities to be introduced and applications to fall out of compliance. It can take days or weeks to detect and respond to container specific attacks, leaving your company open to data breaches, reputational damage, and compliance fines. Ultimately these issues can

distract DevOps teams and slow down releases.

Built for Kubernetes and container security: To make your teams as productive as possible, you need to automate and merge security and compliance into the DevOps workflow. Your tool of choice should address security requirements across all stages of the Kubernetes lifecycle and integrate into the DevOps ecosystem.

- **Deploy securely:** Use a single workflow for detecting vulnerabilities and misconfigurations in containers. Verify configuration meets CIS benchmarks and application compliance with NIST and PCI.

- **Block threats at runtime:** Prevent threats without impacting performance using Kubernetes-native controls. Strengthen Kubernetes security using automated policies.

- **Respond quickly:** Automatically remediate by triggering response actions and notifications. Conduct forensics after the container is gone. Enable audit by correlating Kubernetes activity.

Sysdig secure: Sysdig Secure embeds Kubernetes security and compliance into the build, run, and respond stages of the application lifecycle. Now you can identify vulnerabilities, check compliance, block threats, and respond faster. Powered by the open-source cloud-native runtime security project called Falco. Read more about how Sysdig Secure extends Falco.

- **Image scanning:** Scan container images in the CI/CD pipeline and block vulnerabilities before they reach production.

- **Validate compliance:** Validate compliance across the lifecycle of containers, Kubernetes, and cloud-native workloads. Identify violations.

- **Runtime security:** Detect and block attacks, combining deep visibility into system calls with Kubernetes metadata, labels, and audit events.

- **Forensics and audit:** Record a snapshot of pre- and post-attack activity through system calls. Incident response and post-mortem analysis.

Sysdig Monitor

Monitoring and Troubleshooting for Containers

Sysdig Monitor allows you to maximize the performance and availability of your cloud infrastructure, services, and applications. Cloud monitoring at scale, with full Prometheus compatibility, provides deep visibility into rapidly changing container environments.

Full Prometheus compatibility: Sysdig Monitor is the first commercially available platform that is fully compatible with Prometheus. Give developers their preferred monitoring approach without the management headache. Sysdig scales to millions of metrics with long-term retention and a single backend.

Sysdig Monitor features: Sysdig Monitor is the tool you need to gain visibility into the infrastructure, applications, and services that drive your business:

- Full-stack monitoring.
- Prometheus compatibility.
- Topology maps.
- Dashboards.
- Adaptive alerts.
- Kubernetes troubleshooting.
- Sysdig teams.

Overview:

Mend is the only all-in-one licensing, security, and reporting solution for managing open-source components. It is also the only one that operates in real time by automatically and continuously scanning dozens of open-source repositories and cross-referencing this data directly against the open-source components in your build.

Mend helps you find optimal components, automatically alerts you about known security vulnerabilities, bugs, new versions, patches, and fixes in the components you're using. It automates the creation and enforcement of your company's licensing policies and centralizes inter-departmental communications and approval processes. It keeps detailed inventories and due diligence reports.

It's compatible with pretty much all programming languages, build tools, and development environments. And possibly the best thing about it: You just plug in and forget about it — unless there's a problem. It's everything software development teams need to get the maximum out of using open-source components without the headache, so they can focus on what they should be doing — making beautifully constructed software.

Funding:

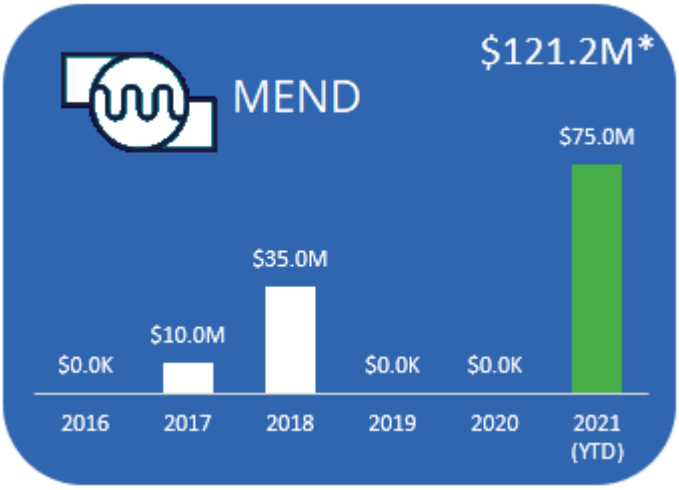

Capabilities:

Mend Core

Mend is with you in every step of the software development lifecycle and keeps monitoring your open-source components, even after you release, based on the last build inventory report. From the moment you plug it in, it works continuously and automatically from behind the scenes, keeping close tabs on your open-source components.

With You Throughout the SDLC

- **Selection:** Mend's browser plugin will save your developers time and help them choose the best open-source components for your product. The selection tool will show your developers all the required information on each component, including whether it meets your company policy while they are browsing open-source repositories.

- **Maintenance:** Mend continues to cross-reference your last open-source inventory against its continuously updated database and notifies you of any issues that come up. Automatically maintaining your products means maintaining long-term relationships with your customers.

- **Detection:** Every time a new component is added to the build, Mend's plugins automatically calculate and assign a digital signature then cross-reference against its database, which is continuously updated. This is used to determine if it's an open-source component.

- **Release documentation:** Automated licensing applications, copyright creation (EULA), and certificates of originality make it easy to get your product out there.

- **Alerting:** Once an open-source component is detected, Mend pulls all relevant information from its massive database and send you alerts within minutes after

the build is complete:

  - **Policy alerts:** Mend cross-references all information about the open-source components in your build with your company policy and notifies you about problematic licenses or any other policies you have set up.

  - **Security alerts:** When there’s an open CVE on one of your components, along with their level of severity, relevant CVE info, and if there's a fix.

  - **Bugs alerts:** Including their level of severity and if there are any patches.

  - **Version alerts:** For a component’s more updated version.

- **Reporting:** Get detailed, real time, inventory, licenses, risk, due diligence reports, and more at any given time according to the last time you ran your build. Effortlessly provide comprehensive updated reports with one click.

- **Policy:** Set up automated policies by defining your acceptance, rejection, and internal approval process protocols per open-source licenses type, security vulnerability severity, software bugs severity, library age, and more. As soon as a developer attempts to add an open-source component that is not automatically accepted per your policies, you’ll get an alert.

Mend for Developers

Work Smarter, Not Harder:

- **Choose right from the first step:** Get all the information needed while browsing for open-source components.

- **Integrate with your IDEs and Repos:** Detect problematic components at early stages in the SDLC.

- **Automate remediation process:** Speed up remediation by automatically updating vulnerable and outdated components.

Code Securely Without Slowing Down Development:

- **Mend Remediate:** Identifies vulnerable and outdated open-source components in your repos and automatically generates Pull Requests (PRs) with a suggested fix. Automated workflows based on vulnerability severity, CVSS score, or a new version release can be defined.

- **Repo integration:** Detects all open-source components in the repos UI, enforces policies automatically, and generates inventory, security, and compliance reports. It also alerts on vulnerabilities and provides detailed information, including a suggested fix.

- **IDE integration:** Provides developers with real-time information about open-source vulnerabilities in their IDE UI with practical remediation guidance, so they don’t need to switch between applications or wait until committing the code.

- **Browser integration:** Allows developers to view a snapshot of a component’s details while browsing on web pages such as StackOverflow and Maven Central before downloading the component.

Mend Prioritize

The sharp rise of reported open-source vulnerabilities in recent years presents software development and security teams with new challenges. Teams can no longer fix all bugs and remain on schedule, and prioritization is quickly becoming a necessity in order to focus limited remediation resources on the most critical issues.

Effective usage analysis technology helps teams to do just that — prioritize. It scans open-source components with known vulnerabilities to assess their security impact on your software. Prioritization is based on whether your proprietary code is making

calls to the vulnerable method, making it effective.

Prioritize Based on Effectiveness:

- Vulnerabilities effectiveness level is displayed with shield icons.
- The summary pane displays the number of libraries analyzed, their severity, and how many are effective.
- The Analysis Statistics section at the bottom displays the percentage of libraries analyzed, and the number of effective and non-effective security alerts.

Optimize Remediation Processes:

- When an effective vulnerability is identified, a detailed call graph presents the complete paths from the proprietary code to the vulnerable functionality.
- This pinpoints the exact location of the vulnerable functionality and the path that leads to it.
- The call graph shows developers where a reference occurs, including filename, class name, and line in the code.
- These details considerably shorten review and remediation time, saving precious resources and helping organizations fix their critical vulnerabilities faster.

Simplify Vulnerabilities Remediation with Effective Usage Analysis:

- **Reduce security alerts by 70%-85%:** Significantly reduce the number of vulnerabilities by focusing on the effective ones first.

- **Speed up remediation processes:** Detailed call graphs pinpointing the path to the vulnerability will speed up remediation time.

- **Improve collaboration between teams:** Use effectiveness as an objective indicator that determines the impact of a security vulnerability to minimize friction between security and developers.

Mend for Containers:

- Mend for Containers continuously detects vulnerabilities and manages licenses from early development all the way to production.
    - **Development**
    - **Build**
    - **Container Registry**
    - **Deployment**

- **Native support for container registries:** Mend for Containers offers native integrations for Docker Hub, Amazon ECR, Azure Container Registry, Google Cloud Registry, and JFrog Artifactory.

- **Risk-free deployment with Kubernetes:** The Mend Kubernetes Controller is a designated lightweight pod located inside your Kubernetes cluster. It detects all open-source components in your cluster and alerts on issues per your organizational policies. It also supports all Managed Service Providers (AKS, EKS, and GKE).

- **Management that’s built for containers:** Mend’s management infrastructure has been designed to meet the specific needs of container

hierarchies, making operating in their environments a seamless experience.

- **The simplest way to secure open-source components in your containers:**
    - **Continuous integration:** Native integrations with all container registries and Kubernetes-managed service providers.
    - **Automated policy enforcement:** Block vulnerable components from entering your containers to ensure your container security.
    - **Real-time policy violation alerts:** Get security alerts when vulnerable components are added to your container or when new issues are reported.