Guide

The Essential CNAPP Buyer's Guide

15 questions you must ask your vendor

uptycs

Table of Contents

Introduction

The CNAPP vendor landscape continues to evolve as cloud environments, application architectures, and security operations become more complex. The following questions focus on capabilities that can help distinguish one platform from another. Assess which ones are most relevant to your organization and environment, add them to your CNAPP RFP or evaluation process, and make a more informed decision.

01 Can you provide centralized and consistent security and compliance across my developer ecosystem, data center, private cloud, and public cloud, including ephemeral workloads?

Modern applications will live largely in hybrid environments. That's one of the conclusions of 451 Research's 2023 annual survey on how enterprise IT decision-makers are adapting their applications in the cloud era.

Figure 3: Landing zones for modernized workloads by modernization focus

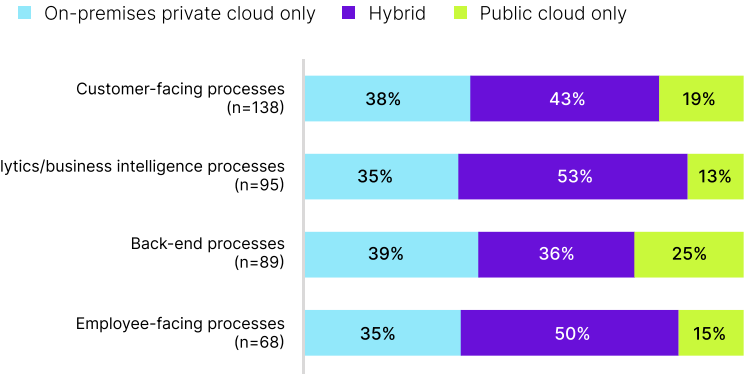

Q. Which business processes are (or were) the main focus of your organization's application modernization's efforts? (n = 400).
Q. In which environment are/will the modernized applications for your organization's [main focus of application modernization process] be running in production? (n = 388).
Base: Organizations that have modernized or will modernize applications.
Source: 451 Research's Voice of the Enterprise: Cloud, Hosting & Managed Services, Application Modernization 2023.

Having siloed security tools and separate views for private and public clouds and data centers causes blind spots. These blind spots can lead to inconsistent security policies, misconfigurations, vulnerabilities, and other costly mistakes, especially during major initiatives such as digital transformation, app modernization, AI infrastructure, cloud migration, and cloud repatriation or cloud exit.

It's also important to have deep historical visibility into ephemeral workloads. The diagram below represents a trend we're seeing across the Uptycs customer base that might hold true for you. Over time, the proportion of workloads running in the cloud versus on-prem is likely to grow, and the proportion of ephemeral workloads versus long-running, stateful workloads is also likely to increase across both public cloud and on-prem deployments.

Traditional posture management is a valuable step, but it's only the beginning. Agentless-only CNAPPs or immature agents can't capture security telemetry from ephemeral workloads if a workload is gone before a scheduled agentless scan begins. An effective CNAPP should combine broad agentless coverage with deep runtime telemetry, historical context, and real-time response, moving beyond risk identification and prioritization to active protection and prevention. For that, you'll need a CNAPP with a sophisticated agent and the ability to efficiently collect deep historical data and insights. You'll also need visibility into how the software is built, for example, across the developer laptop, GitHub, Jenkins, Artifactory, and Kubernetes cluster.

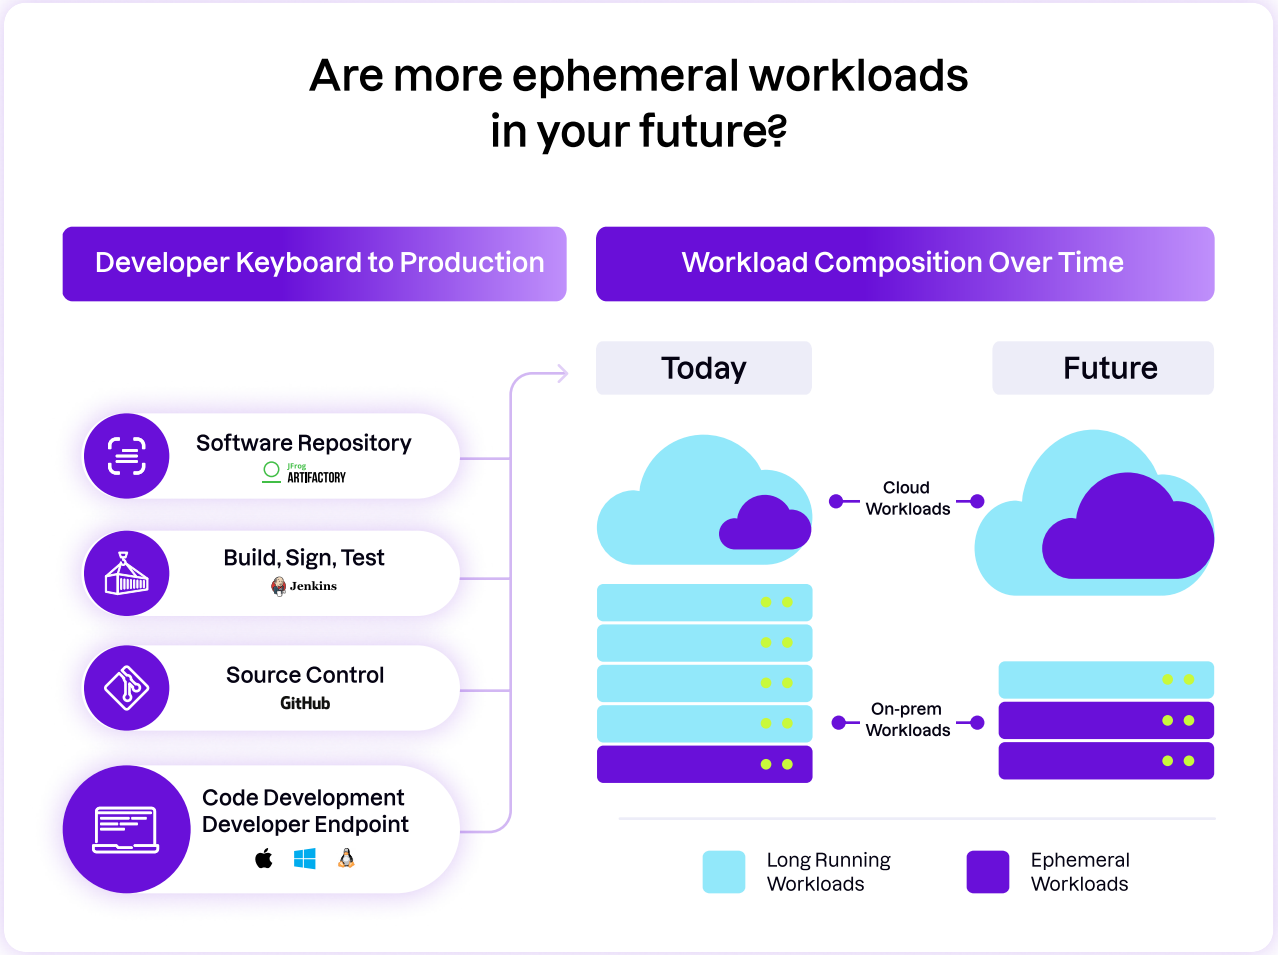

Uptycs AI-Native CNAPP empowers teams with a single data model that unifies security insights from runtime workloads, identities, cloud infrastructure, Kubernetes, and your software development pipeline. This comprehensive view enables you to prioritize risk and detect, investigate, and respond to threats across every layer of your cloud security journey, from build time to runtime.

Operationalize fixes and prevention policies at cloud speed, from build time to runtime, with integrations, exception management, and customizable rules and policies. This allows teams to consistently prioritize risk, enforce policies, and respond to threats across hybrid cloud environments.

02 Are you able to manage the quantity and complexity of data emanating from my application infrastructure?

The cybersecurity industry has a data problem. We've been on a collision course with a massive volume, variety, and velocity of data. There's an ever-growing number and diversity of infrastructure units, system changes, regulations, and people and tools touching IT infrastructure. AWS has more than 14,000 identity and access management permissions alone.

Number Infrastructure/Data Units

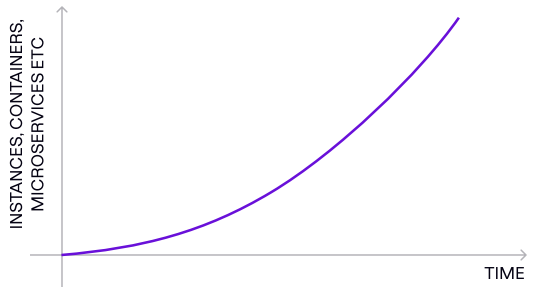

Frequency of Changes

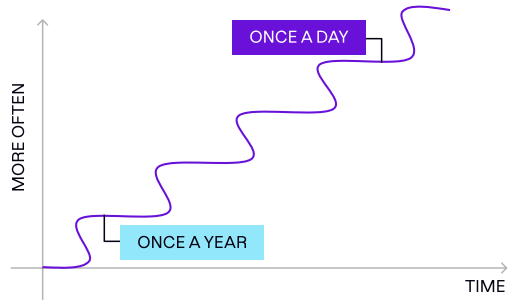

Number of Cybersecurity Regulations

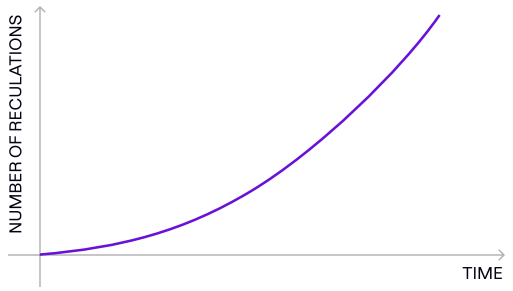

Engineers Involved

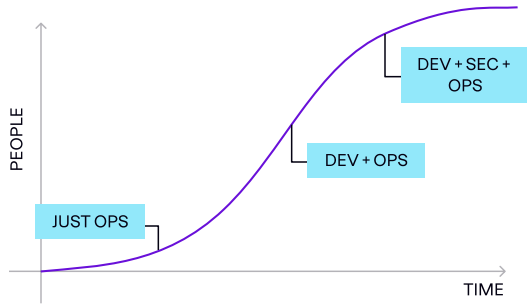

Machine Identities and Entitlements

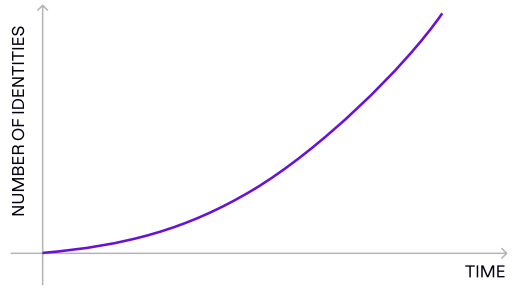

Tools and Services Used

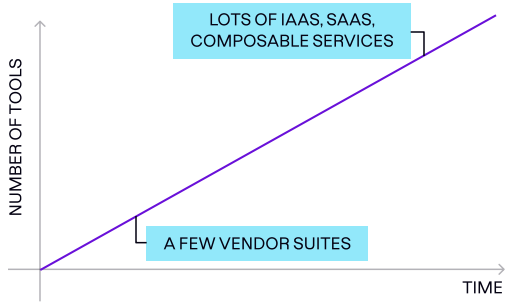

Google Ads which pioneered streaming analytics for subsecond, real-time bidding\

Akamai for horizontal scaling and high-speed data transfer (Uptycs CEO was the former chief architect at Akamai)

Salesforce which provides one data model with multiple use cases–sales, marketing, and support

SAP, arguably the leader in business process management and analytics-powered outcomes

The idea is to normalize security telemetry close to its collection point, stream it into the Uptycs security data lake, then bring everything together under a unified data model and security console. No black boxes, no ETL, and no need to put in a support ticket to get new correlated insights.

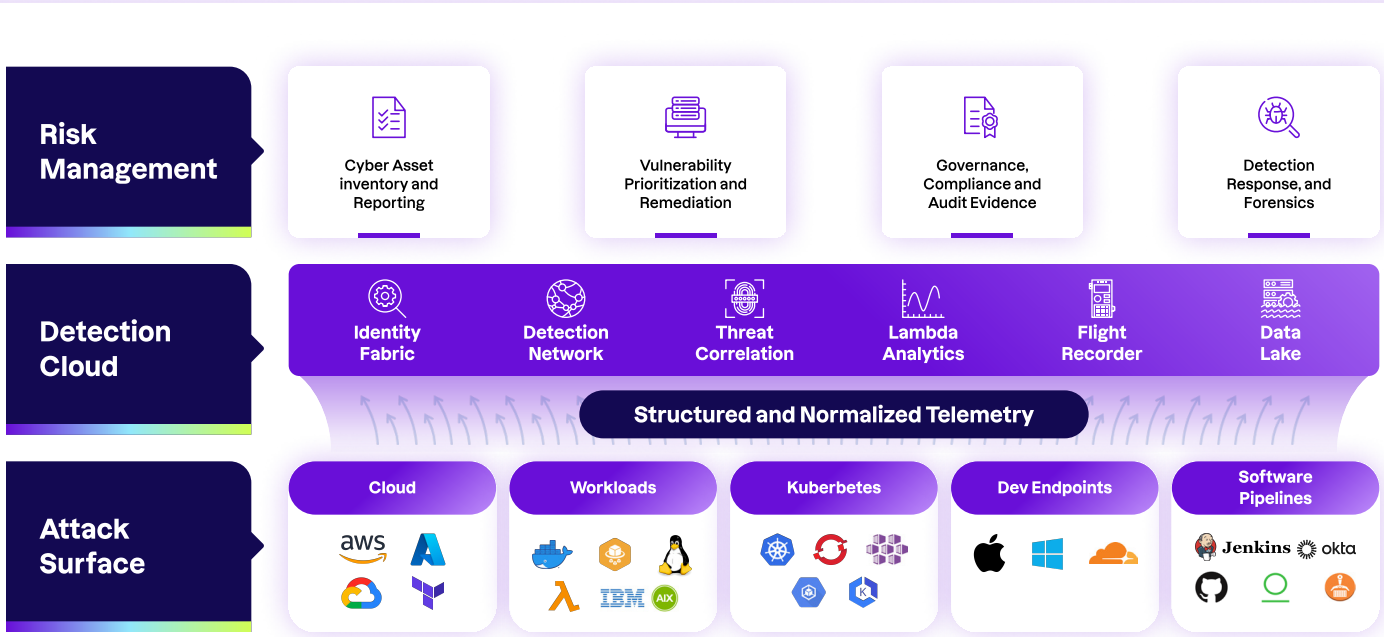

It's a much more scalable architecture for tackling the data challenge, and it's powered by a three-stage analytics pipeline.

1. Collect.

Uptycs collects real-time telemetry via eBPF from every attack surface across your hybrid cloud fleet, including process, network, and file activity, as well as software catalogs and packages. This telemetry is stored continuously for historical lookups, providing complete visibility for compliance audits, threat hunting, and continuous risk prioritization. This is all made possible by the Uptycs Sensor, a performance-optimized, osquery-based agent.

2. Aggregate.

Data from your modern attack surface is aggregated within an Extended Detection Network to facilitate high-speed cohort analysis.

3. Analyze.

Uptycs analyzes structured telemetry across cloud infrastructure, workloads, identities, Kubernetes, and development environments to correlate risks, detect malicious behavior, and support investigation and response. Preventive guardrails also continuously monitor the software pipeline for issues such as branch protection gaps, code scanning alerts, and registry misconfigurations, helping teams mitigate software supply chain risks earlier.

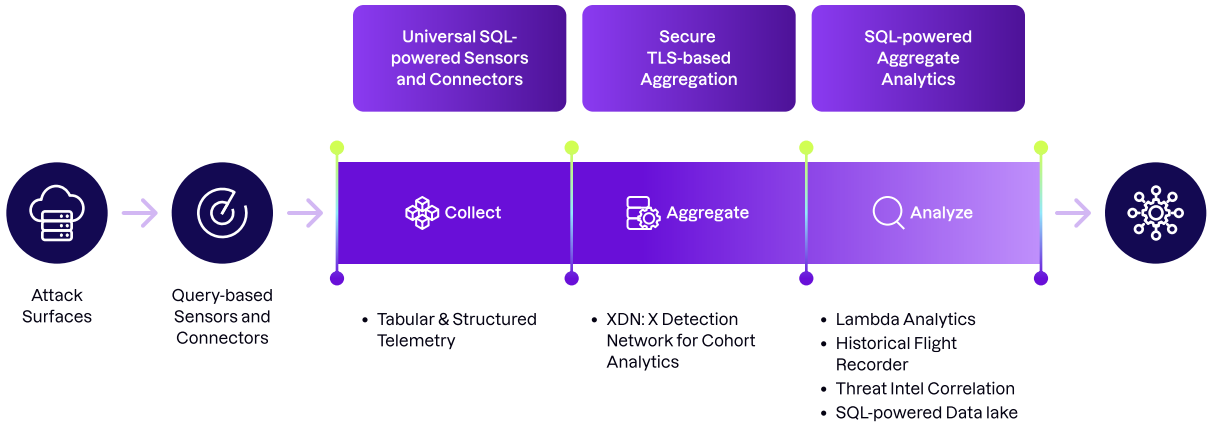

With this more modern shift-up architecture, you have full control and visibility over your security data, can get the correlated insights you care about most, and take decisive action. You can also get alerts, dashboards, and reports packaged through the lens of industry standards such as MITRE ATT&CK, CIS Benchmarks, FedRAMP Controls, SOC 2, and more.

03 Have you ever taken a customer's production environment offline or had to move your software into reduced functionality mode?

Cloud security platforms must process large volumes of telemetry without introducing operational risk to production environments. Buyers should understand how a vendor manages scalability, agent upgrades, resource controls, and continuity of service.

Uptycs has significantly optimized its osquery-based agent for stability and performance, minimizing the memory, CPU, and disk I/O footprint. On Linux, the Uptycs Sensor uses eBPF to noninvasively collect system-level telemetry with very low CPU overhead. Average CPU utilization is approximately 1%, and you can set a ceiling for CPU consumption.

Uptycs will automatically adjust the workload based on your ceiling.

Reliable performance on Linux servers helps avoid disruptions to production applications, including HPC (high-performance computing) and HPCaaS environments. Uptycs runs across a wide variety of Linux distributions, including IBM Linux on Z, LinuxONE, and AIX.

That's why enterprises like PayPal, Comcast, and Nutanix rely on Uptycs to secure the development ecosystems they use to build their applications and run their workloads.

04 Can you detect the top 500 toolkits used by threat actors, either using an agentless or agent-based approach? If so, what's your method of detection?

Rarely do organizations deploy full-blown antivirus software on production Linux workloads due to resource requirements that may affect workload availability and performance.

Threat actors are prone to deploy malicious toolkits for cryptomining, reverse shells, and ransomware. The entry point for malicious code can be an existing vulnerability or a compromise introduced through the supply chain, such as stolen developer credentials.

To address these risks without impacting workload performance, Uptycs lets your team take advantage of industry-standard YARA rules to identify malware in your environment. YARA gives security teams a flexible way to identify and classify malware and suspicious patterns across files and process memory.

Uptycs simplifies the process of creating, running, and saving custom YARA rules, offering a user-friendly approach to managing and updating your threat detection arsenal. With features like eBPF-based auditing for Linux workloads, containers, and cloud-specific threats, Uptycs extends the reach of your YARA rules, enabling the monitoring of network, socket, and process activities.

Out of the box, Uptycs maintains hundreds of YARA rules to detect 500+ APT toolkits. You can also create and deploy custom YARA rules to scan process binaries and process memory. Monitored files are scanned using several hundred YARA rules, with events raised immediately following a match. In addition, you can scan any file or process ad hoc in real time.

YARA helps cybersecurity professionals identify and classify malware by crafting rules that pinpoint specific patterns within files. Developed by Victor M. Alvarez, known as @plusvic on X, and a dedicated software engineer at VirusTotal, YARA empowers malware research and digital forensics. Its ability to dissect and analyze files for threats has rendered it an indispensable tool for security analysts worldwide.

> "An analogy I often use is: YARA is to files what Snort is to network traffic."
>
> **Victor Alvarez**
> PenTest Magazine, February 2016

Essential YARA Resources:

- VirusTotal YARA page
- InQuest / awesome-yara on GitHub
- "Resource Efficient Malware Scans with YARA + osquery," Uptycs
- "Threat detection and malware hunting with Uptycs"
- "Classify Malware with YARA," John Hammond
- Deploying YARA Scanning at Scale for Advanced Attacker Detection

05 Can you use runtime data to prioritize vulnerabilities and block vulnerable processes from running?

In addressing a widespread vulnerability like Log4j, the key is not just to detect but also to prioritize effectively for rapid response. Unlike traditional methods, which may take a full day to indiscriminately flag thousands of VMs, Uptycs leverages runtime data to precisely identify and prioritize vulnerabilities based on their actual use in your environment.

If you have a large environment with 100,000 VMs, the number of flagged vulnerabilities can be tens of thousands. However, most of these VMs will have the vulnerable .jar file on disk but will not be running it. In many cases, the vulnerable Log4j .jar was never even used. It came with installed Java software, but no applications use the vulnerable class.

Where do you start when you're under threat from such an easily exploitable and widespread critical vulnerability? Do you apply a sledgehammer? For example, in the wake of Log4j, many companies took production applications offline by shutting off internet access while they figured out what to do.

This is where Uptycs' deeper runtime scanning comes in. Uptycs correlates runtime activity, vulnerability context, and behavioral detections mapped to MITRE ATT&CK to help teams understand which risks require immediate attention. It also enables teams to attribute and establish provenance for code-level activity, supporting faster root cause analysis.

Uptycs will identify both the hosts that have a Java process running and those that have the vulnerable .jar open. When this check is performed, the scope of investigation within a large environment drops from tens of thousands to hundreds of VMs.

Yet, knowing where to focus is only half the battle. You also need to fix the problem. With Uptycs, you can customize and provision rules to automatically remediate. For example, you can create a rule to immediately terminate a process when the binary is coming from a vulnerable package or .jar file, or when a process launches with a shared library open. Once the vulnerable processes are stopped, you can quickly recover by updating the packages and .jar files and return to business as usual.

06 Can you instantly kill a reverse shell, cryptomining application, or ransomware activity?

What do threat actors, including nation-states, do when they find a vulnerability? They leverage the vulnerability, perform reconnaissance, and persist. Then, at the opportune time, they create reverse shells to communicate with command-and-control systems so attackers can execute actions designed to steal intellectual property, PII, and financial data. Even more common is for threat actors to steal credentials and pretend to be a genuine user.

That's why you need the capability to instantly kill a reverse shell as soon as it's launched. Stopping this activity at launch can prevent the attacker from progressing to the next stage of the attack chain or exfiltrating data.

Uptycs also provides out-of-the-box remediation capabilities, such as killing a container process, quarantining a host, or pausing a process to take a deeper look at the system's file contents and processes, all from a single click and in real time.

07 How do you identify anomalous behavior? What types of historical data are you feeding into your AI models?

Many attacks involve a threat actor using stolen credentials to log in and execute malicious activities. To standard detection tools, these activities may look like they are being performed by a legitimate user. Detecting suspicious activity often requires behavioral analysis across large volumes of identity, process, network, and cloud data. Uptycs captures detailed behavioral data, including:

- Processes launched
- Active times of day and geographies
- Logins and logouts
- Commands and processes executed
- Command lines used
- Folders they operate in
- IP addresses reached
- HTTP requests and API calls made
- Parameters used in the API
- And more

Effective security AI depends on complete, structured, and relevant telemetry. Uptycs combines behavioral, identity,

cloud, workload, and runtime context through a unified data model. Juno AI Analyst can use this context to support investigations and provide evidence-backed answers that analysts can verify.

08 Can you blend runtime insights with historical data? Can you show me IoCs (indicators of compromise) based on runtime insights that were run in the past? What about IoCs on ephemeral hosts?

As assets become distributed across cloud and physical infrastructure, preventing malicious activity is an increasingly complex battle on multiple fronts. Not only is it harder to reliably detect an attack in real time, but it's also common for attackers to cover their tracks. Virtual machines and other ephemeral assets can be created and destroyed at any time, meaning evidence of an attack may disappear before it is detected. This is especially true for zero-day attacks and vulnerabilities because IoCs may only be published when the zero-day is announced. An environment may therefore have been exposed before public indicators or detection content became available.

Historical data lets you uncover malicious activities that may have been completely missed using runtime insights alone. Historical data allows you to create a clearer picture of how an attack occurred. Where real-time threat intelligence may reveal only fragments of an attack, historical data enables you to reconstruct what happened and when.

However, the greatest challenge of scanning historical data is scale. A large hybrid cloud estate's historical dataset is enormous. Handling it efficiently requires both extremely robust infrastructure for storing data and a unique system architecture capable of querying vast datasets in real time, which can span 100,000 machines or more. As shown below, the ability to blend historical data with real-time insights can also help busy incident response teams focus on what truly matters.

The table below is illustrative of a vulnerability assessment across 100,000 workloads for a widespread vulnerability like Log4j:

Uptycs Assessment Type	Severity	Business Impact	Runtime Insights Running now		Runtime Insights Ran in the past but not running now		Historical Runtime Insights Ephemeral hosts		Uptycs Assessment Type
Description	High (based on agentless or agent-based periodic scan)	Sensitive data on hosts? Lateral movement risk? Production Environment?	Where are these vulnerable processes running right now?		Any IoCs on those hosts two weeks ago?		Any IoCs on ephemeral hosts? Were there consequences? Any evidence of data exfiltration?		
Number of vulnerabilities	100,000	10,000	400	+	400	+	900	=	1,700

Of course, for any vulnerability, knowing whether a fix is available and whether the vulnerability is exploitable helps prioritize remediation efforts.

Many CNAPPs can identify whether a vulnerable workload is exposed to the internet, and some can show vulnerable packages running now. Buyers should also ask whether a platform can determine if those packages ran previously, whether related indicators were observed, and whether a breach occurred.

Uptycs can, and once detected, you can quickly remediate with automated workflows and curated step-by-step guidance.

The same goes for Kubernetes and containers. Uptycs' breadth and depth of telemetry cover the Kubernetes control plane down to an individual container process, so you can prioritize vulnerabilities loaded in memory across your container fleet.

09 Can you show end-to-end image security visibility from the code repository to runtime?

Without a secure DevOps pipeline, organizations must be more reactive in remediating vulnerabilities. Images with known vulnerabilities, exposed credentials, or malware should never be promoted to a stage where they could accidentally be deployed to production systems. A shift-left principle that allows security teams to prevent vulnerabilities from making their way to runtime is key. In addition, security teams shouldn't have to use separate tooling to identify vulnerabilities in CI and registries, which can create blind spots and increase the time spent evaluating results instead of taking action.

Uptycs scans images across CI and registries for vulnerabilities, malware, and secrets. With CI scanning, Uptycs shows security results and can pass or fail image builds, allowing teams to correlate findings with the SDLC. The same visibility extends to registry scanning, so teams can filter by a specific CVE and identify where a vulnerability is present across CI, registries, and runtime.

Uptycs gives you a unified view of your image security across the SDLC to answer key questions such as:

- Did the developer image go through the security controls across the SDLC? If not, what was missed? Who was the developer responsible for this image?
- Was GitHub configured according to best practices and CIS guidelines?
- Was the newly introduced code properly reviewed according to the organization's established practices?
- Did the build machine itself, for example, the Jenkins worker node, have any vulnerabilities?
- Was the image properly signed and attested by your security team?
- Did the image get deployed from a trusted registry? Images should only be pulled from trusted registries that are scanned for vulnerabilities.

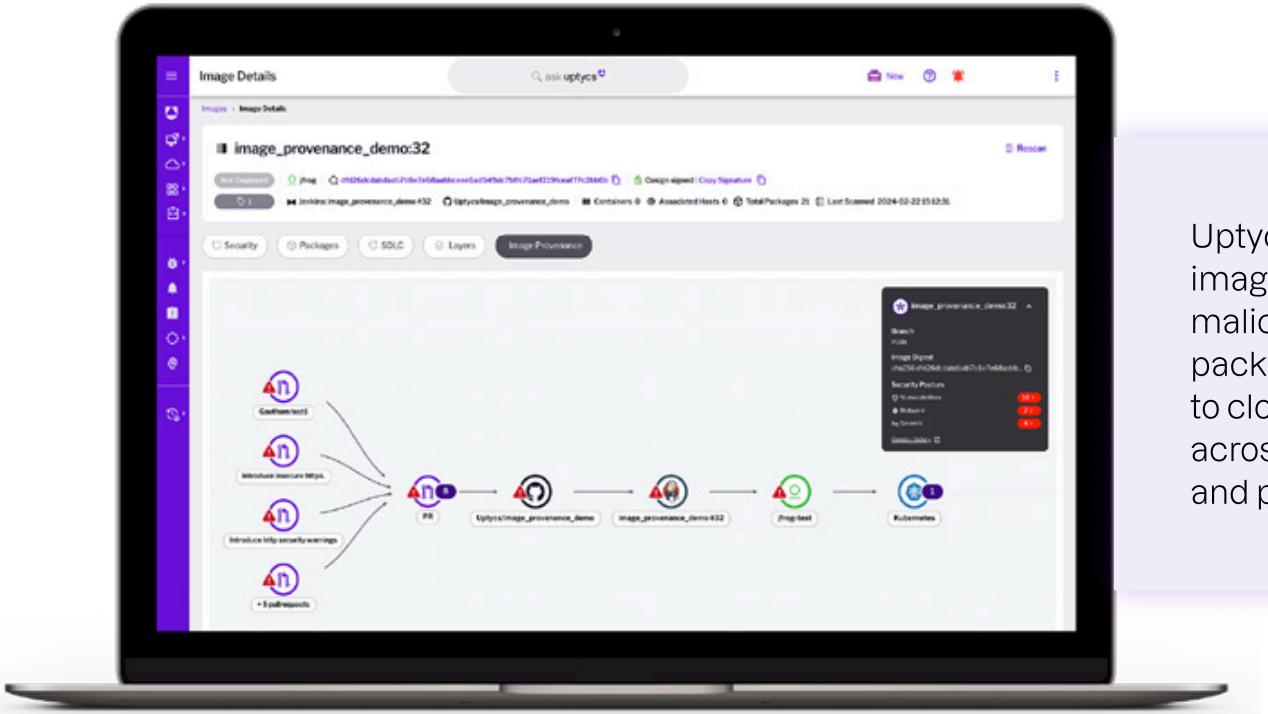

Uptycs goes beyond image scanning for malicious or vulnerable packages to give you code to cloud traceability across the supply chain and provenance.

These insights can be key to ensuring your development teams and business units follow the right processes while maintaining end-to-end security controls across the software pipeline.

For more information, check out our Mastering Kubernetes Security e-Book or the Uptycs for Kubernetes and Container Security Solution Brief.

10 Can you help reduce friction between developers and security teams?

Disconnected tools and limited visibility can force security teams to treat every risk the same way, such as applying a hard block and leaving developers to determine how to resolve it.

With Uptycs, security teams can define flexible posture controls and visibility across the software development lifecycle, from build time to cloud runtime, enabling developers to focus on the key security risks that matter most. For example, for a developer pipeline in GitHub Actions that builds a container image, SecOps can set pipeline policies to fail or ignore a build based on its context. The policy could be set to fail if any malware is found. Conversely, it could be set to ignore vulnerabilities from base-layer images that are not owned by the developer or vulnerabilities that have no available fix.

This enables better collaboration between security and development teams by giving developers clearer ownership and context while allowing security teams to enforce controls based on risk and remediation availability.

11 Can you help me securely scale the software development teams connected to my critical infrastructure?

As we've seen with LastPass and CircleCI, major breaches can begin at the developer's endpoint. That's why it's critical to detect malware and vulnerabilities on developers' laptops and identify suspicious behavior as they move code in and out of repositories and into production.

In the following example, the attacker has taken over a developer's laptop, tried to log in to Okta but failed because of MFA, used credentials found in a .txt file to access GitHub and copy a code repository, and then used secrets found in code to move into AWS. There, the attacker attempts to elevate privileges and exfiltrate a highly sensitive customer database.

Developer endpoints are often outside the scope of traditional CNAPP deployments, which can leave gaps between development activity and hybrid cloud or Kubernetes infrastructure.

Uptycs provides a unified view that can connect suspicious developer endpoint activity with identity, repository, cloud, Kubernetes, and workload context. This creates an early-warning system that can identify and stop threat actors before they access critical data and cloud services.

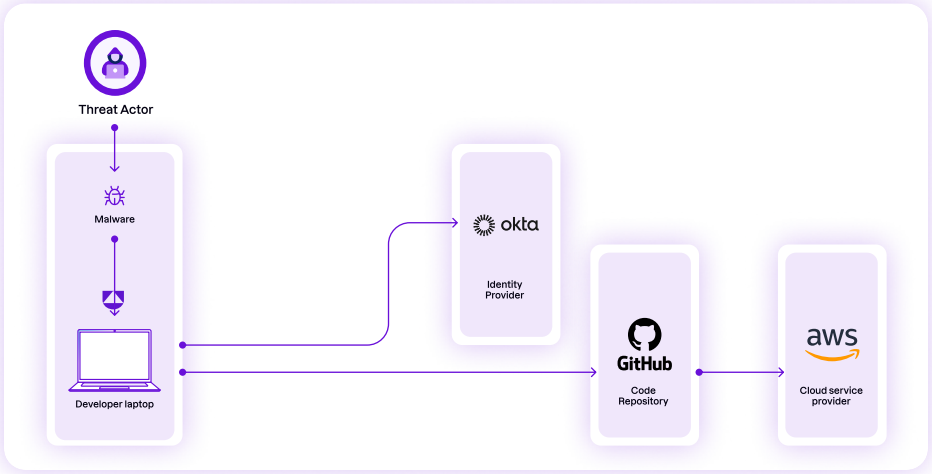

With the trend toward remote and hybrid global workforces, and 64% of firms outsourcing some portion of application development, you may also want to consider an additional visibility and access layer in the form of a Zero Trust Score and integration with a Zero Trust Network Access provider such as Cloudflare.

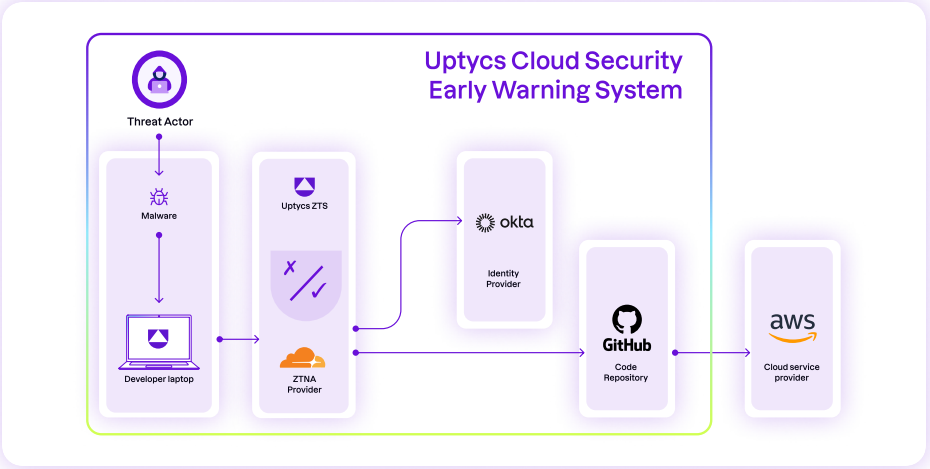

You can use Uptycs as a full-fledged XDR, deploy Uptycs exclusively across outsourced development teams, or add it as a premium visibility and security layer on top of solutions like CrowdStrike and SentinelOne, as SEI and Lumin Digital have done. Let the endpoint protection team cover these assets as they see fit, but since developer endpoints are part of critical cloud infrastructure, they deserve to be treated as such.

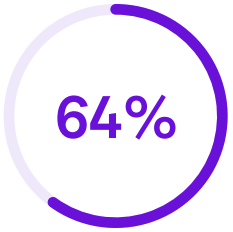

of organizations turn to outsourcing for app development

Source: Avasant Research 2023 State of Global IT Outsourcing

12 How can I customize your solution to fit my needs?

Lack of customization in cloud security tooling can hinder SecOps efficiency and slow down threat detection, investigation, and response times. An Uptycs customer who switched from Lacework and evaluated Wiz puts it best: “Most CNAPP tools are all-or-nothing, but we can customize Uptycs to do what we need. It helps us focus on what matters.”

Many aspects of Uptycs are customizable, including dashboards and reporting, compliance frameworks, and detection rules. For example, custom detection rules and intelligence can be added to Uptycs in just a few clicks. With broad hybrid and multicloud support, Uptycs allows teams to adapt dashboards, detections, policies, reports, and compliance controls to their environments and operating models.

Uptycs has thousands of custom detections, including over 2,300 behavioral detections covering MITRE ATT&CK. Since Uptycs is an open platform, you can see how detections work, how they are scored, and the TTPs behind them. You can then use this information to customize existing detections and create new ones that make the most sense for your organization.

With Uptycs, you can help reduce alert fatigue because you control what is turned into a detection or an alert.

Lookout

“With Uptycs, the same alerts can be categorized as either normal on test environments, for example, or extremely concerning if they occur on production systems. We’re able to create matching asset groups in Uptycs, to which we can apply our backend and all the things we want to do based on our nomenclature. This offers a whole new range of capabilities and security operational efficiencies for us."

Grant Kahn
Director of Security Engineering, Lookout

13 Can your solution help with forensic investigations and remediations? If so, how?

Some CNAPP platforms provide a predefined view of risk but offer limited flexibility for deeper investigation, custom queries, or remediation. Buyers should ask whether they can investigate the questions that matter to their environment and take action directly from the platform.

Uptycs can significantly aid forensic investigations by leveraging its capability to store and analyze up to a year of historical data. This historical depth, combined with the Uptycs Sensor for detailed runtime insights and agentless scanning for broader coverage, gives teams the context needed to conduct thorough investigations.

Advanced container workload protection capabilities further strengthen Uptycs' forensic capabilities, enabling deep analysis of ephemeral container processes alongside developer endpoints and workload activities.

Security teams can also leverage YARA and Sigma rule scanning to uncover malicious activity that may be masquerading as a legitimate process, such as a port scan hidden behind a different process name.

Juno AI Analyst can accelerate investigations by translating analysts' questions into queries, correlating relevant evidence, and showing the data and reasoning behind its conclusions.

Uptycs combines deep telemetry, the ability to query infrastructure like a database, and customizable detection logic to support forensic investigations across cloud infrastructure, workloads, containers, identities, and developer endpoints.

14 Can you run your platform for us in a Managed Detection and Response model?

Some CNAPP vendors provide platform support but do not offer the operational resources to monitor, investigate, and respond on a customer's behalf.

Uptycs offers Managed Detection and Response (MDR) services at three levels: Managed Onboarding, Managed Monitoring, and Managed Protect. Each is designed to augment your security team's capabilities.

Managed Onboarding supports deployment and integration within your environment. Managed Monitoring provides ongoing oversight, identifying and notifying teams about potential threats. Managed Protect provides more comprehensive management, including continuous monitoring and active threat containment.

This tiered approach allows organizations to select the level of service that best fits their needs, enhancing their cybersecurity posture without having to expand their internal workforce.

Additionally, Uptycs offers Professional Services to further support customers, including tailored onboarding, custom dashboard development, specialized configuration, and integration with SIEM and SOAR tools, enhancing security posture and operational efficiency.

Uptycs Managed Services Coverage

	Managed Onboarding	Managed Monitoring	Managed Protect
Prepare	✓	✓	✓
Identify	✗	✓	✓
Contain	✗	✗	✓
Eradicate	✗	✗	✓
Recovery	✗	✗	✗

15 How can you help me improve outcomes for my SIEM or reduce my bill?

For many enterprises, SIEM remains a critical security tool, but organizations must also manage rising data-storage costs and an increasing volume of detections.

#of Workloads	Daily Telemetry Per Workload	Daily Volume Sent to SIEM	Storage: 30 Days	Storage: 60 Days	Storage: 90 Days
200,000	200MB	40TB	1,200TB/ 1.2PB	2,400TB/ 2.4PB	3,600TB/ 3.6PB
100,000	200MB	20TB	600TB	1,200TB/ 1.2PB	1,800TB/ 1.8PB
50,000	200MB	10TB	300TB	600TB	900TB
10,000	200MB	2TB	60TB	120TB	180TB
5,000	200MB	1TB	30TB	60TB	90TB

Uptycs structures security data at its source, thereby avoiding expensive ETL (Extract, Transform, and Load) processes and reducing the costs associated with indexing unstructured data. Its analytics data pipeline is specialized to store and transform high-volume, real-time telemetry from your attack surface into contextualized, actionable alerts.

Uptycs can reduce your SIEM bill because it's capable of storing data with up to 100x compression, along with offering efficient retrieval using sophisticated time-domain partitioning. These capabilities make it a strong complement to a general-purpose SIEM.

Uptycs not only reduces the volume of data sent to SIEM systems but also decreases false alerts, paving the way for a more streamlined and cost-effective security operation. Furthermore, organizations may use SOAR to pull additional data from Uptycs when more context is required for a specific investigation.

By structuring and correlating high-volume security telemetry before forwarding selected findings and context to a SIEM, Uptycs helps organizations balance comprehensive security visibility with operational and data-management costs.

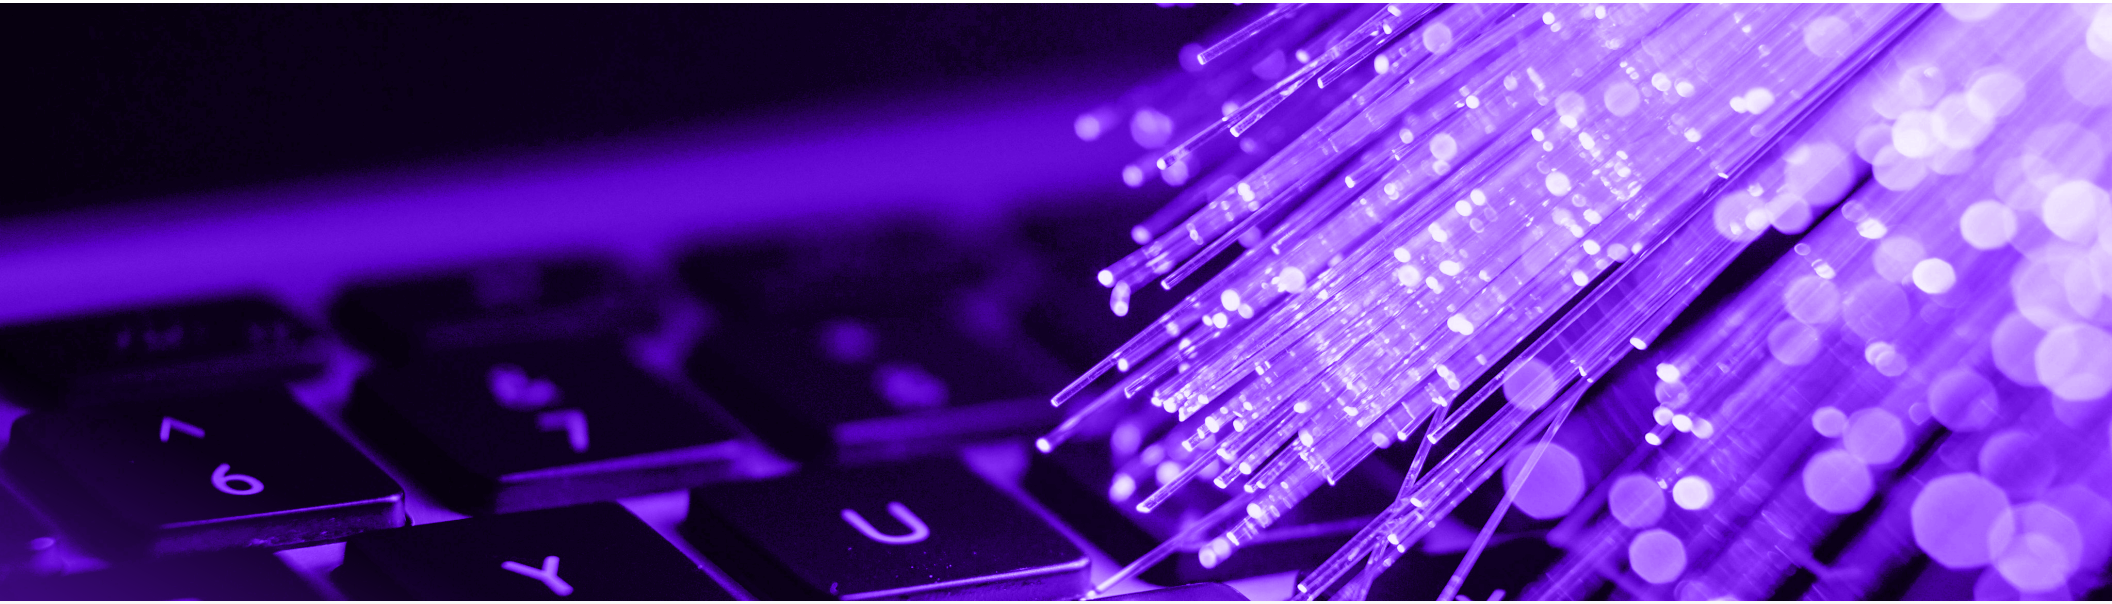

uptycs

Uptycs delivers the industry's first **Unified CNAPP** with integrated XDR, powered by Juno - the only verifiable AI Security Analyst.

Unlike traditional CNAPPs that rely on static snapshots, Uptycs unifies cloud configuration data with deep runtime telemetry across cloud, containers, and endpoints into a single Unified Ontology. Trusted by the Fortune 500, Uptycs powers Continuous Threat Exposure Management (CTEM) by normalizing telemetry into a single data lake, enabling the Juno AI Analyst to provide **"Glass Box" transparency**, turning hours of investigation into minutes of verified answers.

Learn more at uptycs.com/juno-ai