uptycs

Whitepaper

The Essential CNAPP Buyer's Guide

15 questions you must ask your vendor

Contents

Introduction

The CNAPP vendor landscape has grown dramatically over the last three years. The following questions tilt towards more sophisticated capabilities. 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.

1 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

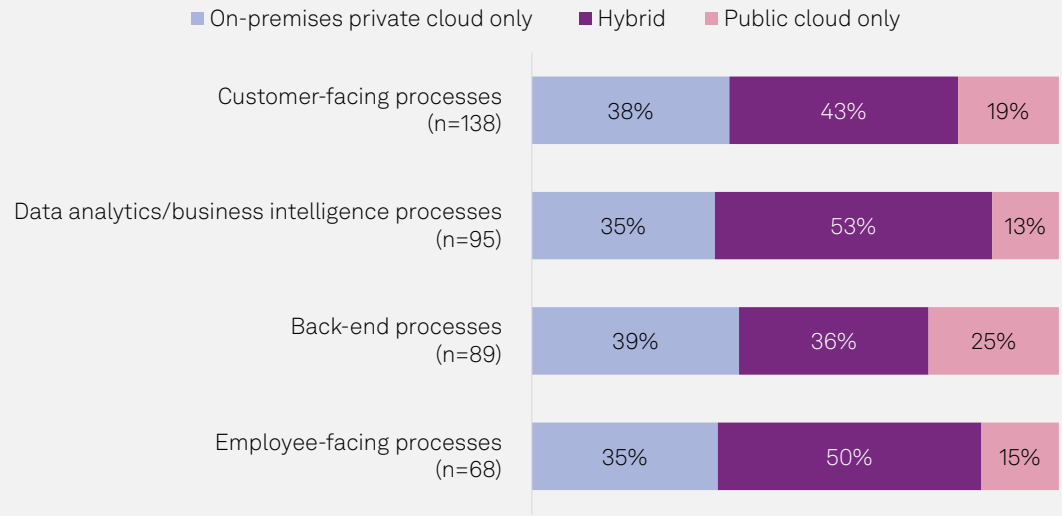

Q. Which business processes are (or were) the main focus of your organization's application modernization's efforts? (n = 400).
Q. In which environment are/will the modernized applications for your organization's [main focus of application modernization process] be running in production? (n = 388).
Base: 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 / 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 which might hold true for you. Over time, the proportion of workloads running in the cloud versus on-prem will grow (no surprise there), and the proportion of ephemeral workloads versus long-running (stateful) workloads will also grow for both public cloud and on-prem deployments.

Agentless-only CNAPPs or immature agents can't capture security telemetry from ephemeral workloads as the workload may be gone by the time an agentless scan is scheduled/initiated. 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.

Are more ephemeral workloads in your future?

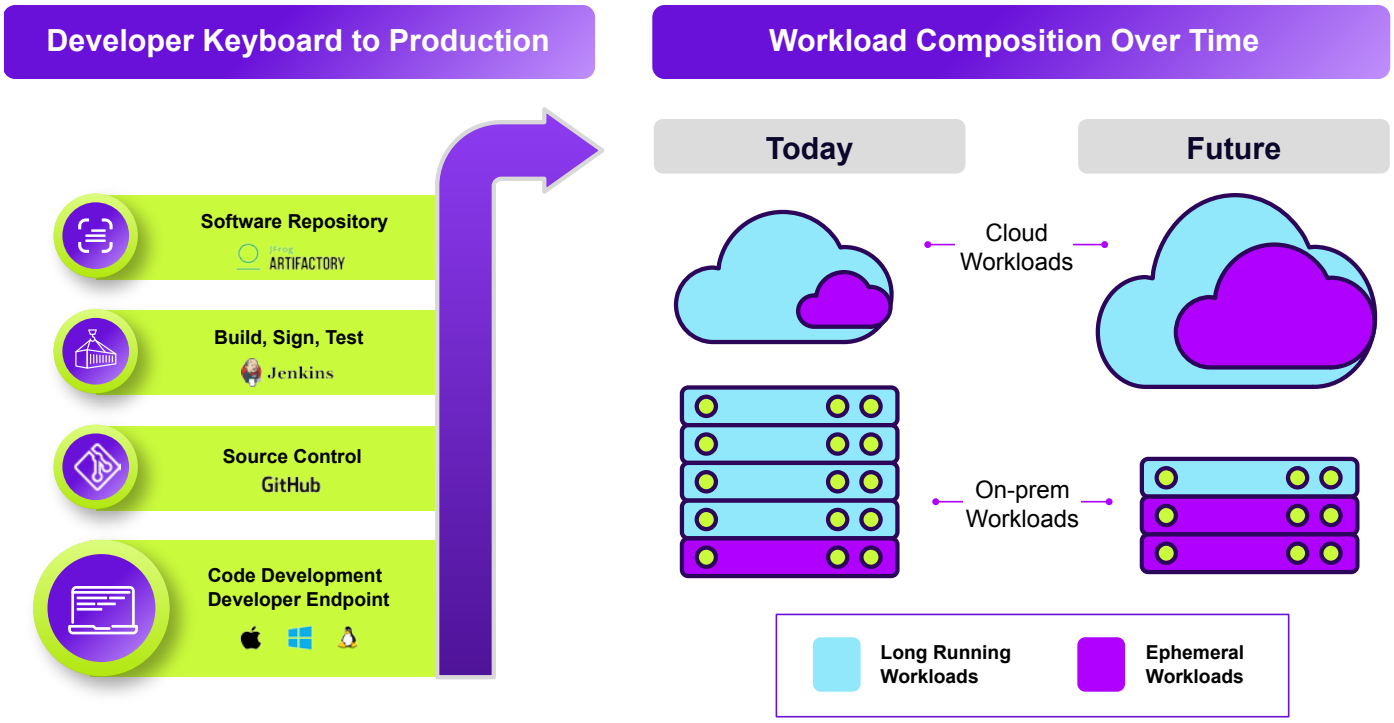

Uptycs ensures security and compliance coverage from the time your developer starts writing code to when the code is deployed and running in a production environment. Start with the coverage you need today with the confidence that you have the flexibility to move workloads to public cloud, adopt more ephemeral workloads, or move things back again. With Uptycs, you can create your enterprise platform (people, process, and technology) to securely move workloads across environments, avoid cloud vendor lock-in, and lower operational costs.

2 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 the number of people and tools that touch IT infrastructure; AWS has more than 14,000 identity and access management permissions alone

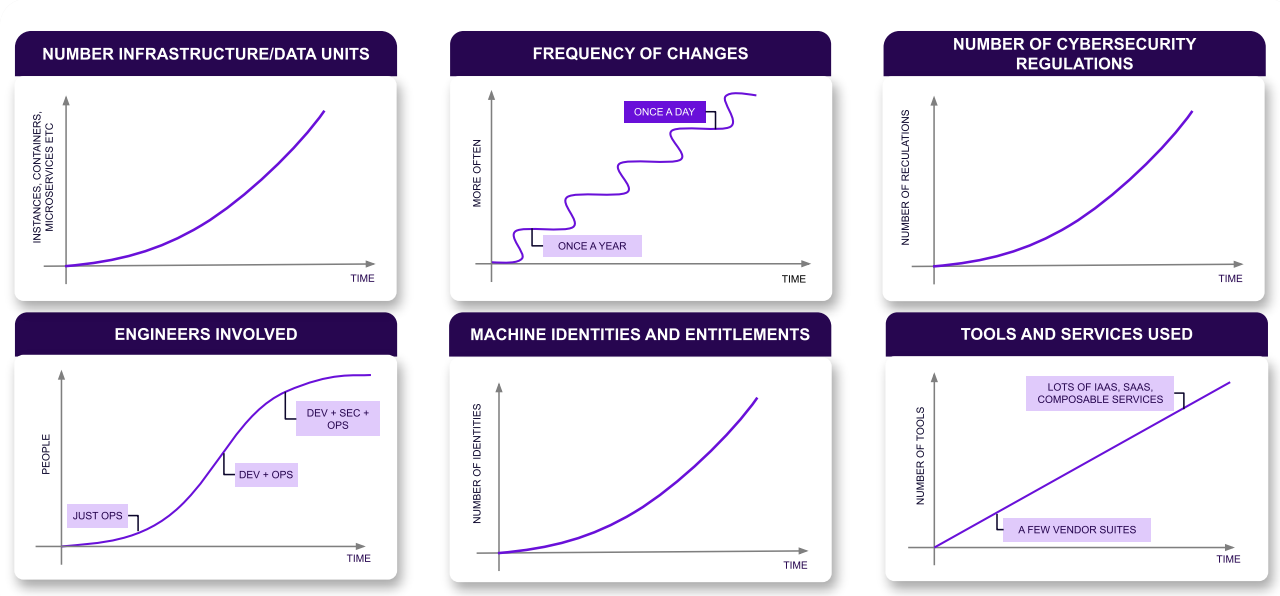

At Uptycs, we solved the cybersecurity data challenge first by leveraging a new approach called shift up security. It's an architecture inspired by industry leaders outside of cyber such as:

Google Ads which pioneered streaming analytics for sub-second, 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 up into a data lake (what we call your Uptycs Detection Cloud), 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.

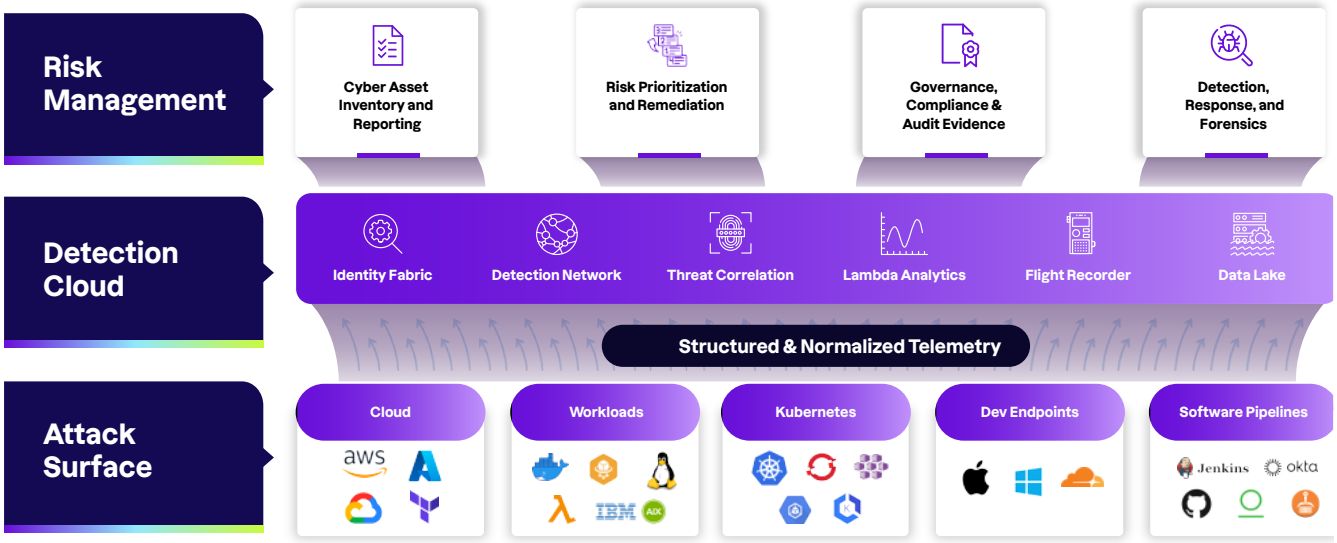

It's a much more scalable architecture for tackling the data challenge and it's powered by a three-stage analytics pipeline.

1. **Collect**. Uptycs uses open SQL to collect tabular structured telemetry from your modern attack surface such as workloads, clouds, containers, Kubernetes, and software pipelines. This is all made possible by the Uptycs Sensor—a performance-optimized osquery-based agent.
2. **Aggregation.** Data from your modern attack surface is aggregated within an Extended Detection Network to facilitate high-speed cohort analysis.
3. **Analyze**. Uptycs provides best-in-class analytics capabilities via a Lambda Architecture for streaming, a flight recorder with 13-month data retention, a SQL-powered data lake for ad hoc investigations, and real-time correlation with threat intel and reputation databases.

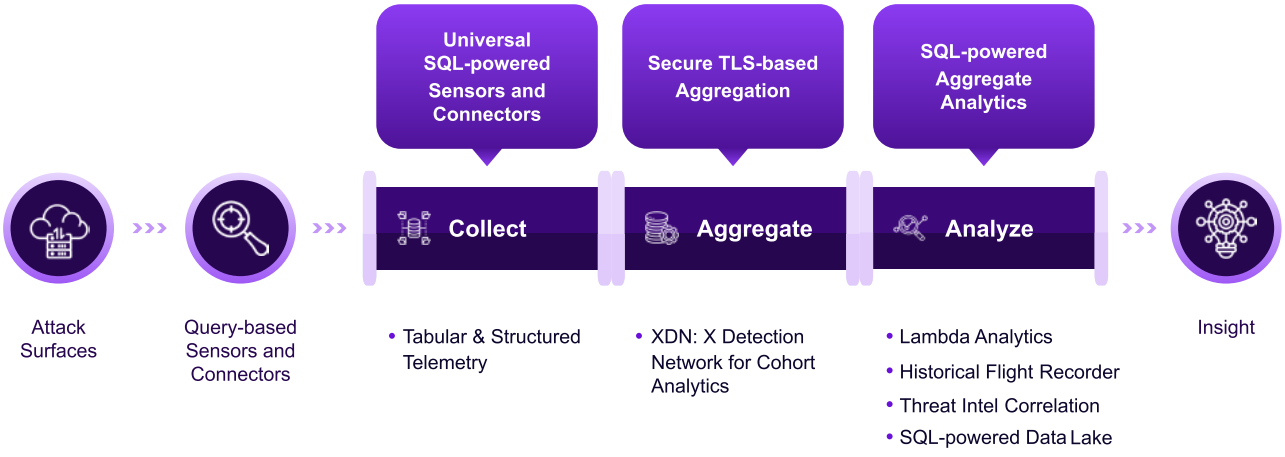

With this more modern shift up architecture, you have full control and visibility of your security data, can get the correlated insights you care about most, and take decisive action. You can also get alerts, dashboards, and reports packaged through the lens of industry standards such as MITRE ATT&CK,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

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

Competing cloud security solutions are crashing because they can't manage the data needed to effectively contextualize and prioritize risk. In other cases, security vendors have had to set their platforms to reduced functionality mode to manage agent upgrades.

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 ~1% and you can set a ceiling for CPU consumption and Uptycs will automatically adjust the workload based on your ceiling.

Reliable performance on Linux servers avoids issues for production applications, including HPC (high-performance computing) and HPCaaS environments. Uptycs runs across a wide variety of Linux distros, 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.

4 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 customers deploy full-blown antivirus software on production Linux workloads due to their unpredictable resource requirements which can even take workloads down.

Threat actors are prone to drop malicious toolkits for crypto mining, reverse shells, and ransomware. The entry point for malicious code can be an existing vulnerability or introduced into the supply chain (e.g. by stealing 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 is more effective and performant than the signature-based approach used by antivirus tools.

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-

yara 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

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. Also, monitored files are scanned with several hundred YARA rules with events being raised immediately following a match. In addition, you can scan any file or process ad hoc in real time.

5 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 not running. 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 will identify both the hosts that have a Java process running and have the vulnerable .jar open. When this check is performed, the area on where to focus 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 where a shared library is opened. Once the vulnerable processes are stopped, you can then quickly recover by updating the packages and .jar files, and return to business as usual.

6 Can you instantly kill a reverse shell, crypto mining app, 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 to command and control centers so the C&C can execute actions 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. By doing this you have thwarted the bad actor and prevented them from moving to the next stage in the kill chain and exfiltrating data.

Uptycs also provides out of the box remediation capabilities such as killing a container process, quarantining a host, or simply 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.

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

Most attacks begin when a threat actor uses stolen credentials to log in and execute malicious activities. In the eyes of standard detection tools the threat actor's activities look like they're being performed by a real user. The only way to detect if activity is malicious is to use ML and AI to analyze large volumes of behavioral data.

Uptycs captures very fine behavioral data including:

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

Fact is, you can have the most beautiful AI model in the world, but it won't deliver the right outcomes without complete data sets. Uptycs has the data.

8 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 hybridized 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. Heightening this security problem, virtual machines can be created and destroyed at any time which means the origin of an attack can be erased from existence before it's ever detected. This is especially true for zero-day attacks and vulnerabilities because the IoCs (indicators of compromise) are only published when the zero-day is announced. So, there's a good chance your infrastructure was infected two to four weeks prior.

Historical data lets you uncover malicious activities that may have been completely missed using run-time insights alone. Historical data allows you to create a clearer image of exactly how an attack occurred. Where real-time threat intelligence can only reveal fragments of an attack, historical data permits you to playback exactly 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:

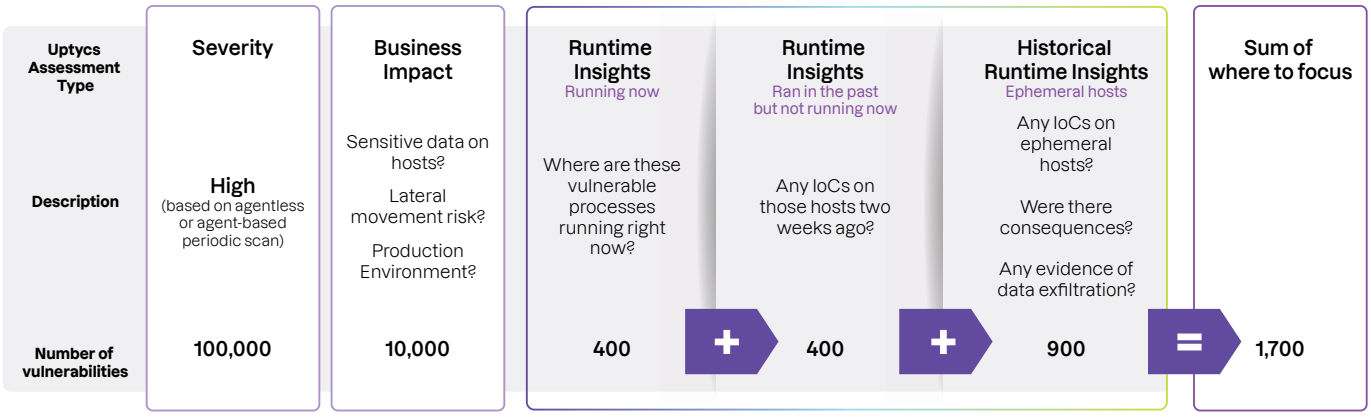

Of course, for any vulnerability, knowing whether a fix is available or not and whether the vulnerability is exploitable or not, helps prioritize remediation efforts.

The fact is, most CNAPPs can tell you if a vulnerable workload is exposed to the internet, and some can show you vulnerable packages running now, but can they tell you if those packages were running three weeks ago, and if and how you were breached? 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.

9 Can you show end-to-end image security visibility from code repo to runtime?

Without a secure DevOps pipeline, customers must be more reactive in remediating vulnerabilities. Images with known vulnerabilities, leaking credentials, or containing malware should never be promoted into a position 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 catch vulnerabilities in their CI versus their registry, which can lead to blind spots and more time spent evaluating results versus taking action.

Uptycs uniquely scans images across CI and Registry for vulnerabilities, malware, and secrets. When you start with CI Scanning, Uptycs shows you the security results and pass/fail image builds so you can correlate your security findings to your SDLC. You also get the same view with registry scanning, so you can filter by a given CVE vulnerability and check on what parts of their SDLC, CI registry, and runtime where a vulnerability is present.

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 as per best practices and CIS guidelines?
- Was the newly introduced code properly reviewed based on the correct organizational 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 ideally being scanned for vulnerabilities.

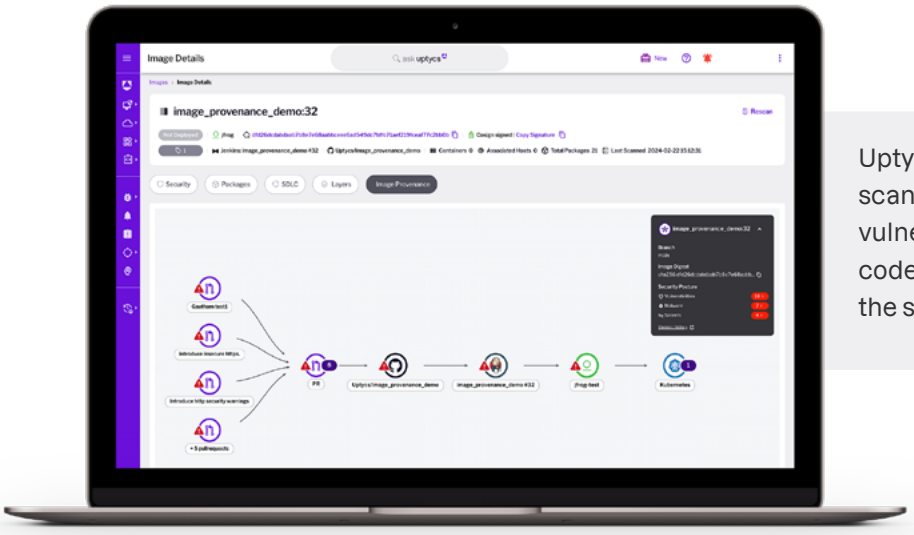

Uptycs goes beyond image scanning for malicious or vulnerable packages to give you code to cloud traceability across the supply chain and provenance.

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

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

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

With disparate tooling and lack of visibility, many security teams are forced to treat all security risks in the same way, for example with a hard block, and then leaving developers on their own to figure out how to get their work unstuck.

With Uptycs, security teams can define flexible posture controls and visibility across all parts of the software development life cycle from build to cloud runtime enabling developers to focus on the key security risks that matter most. For example, for a given developer pipeline in GitHub Actions that builds a container image, SecOps can set pipeline policies to fail or ignore depending on the context. For example, the policy could be set to fail if any malware is found. Conversely, the policy could be set to ignore vulnerabilities from base layer images, which are not owned by the developer, or ignore vulnerabilities that have no fix. This enables better collaboration between security and development teams allowing developers to focus on the most important security issues which reduces inefficiency throughout the organization.

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

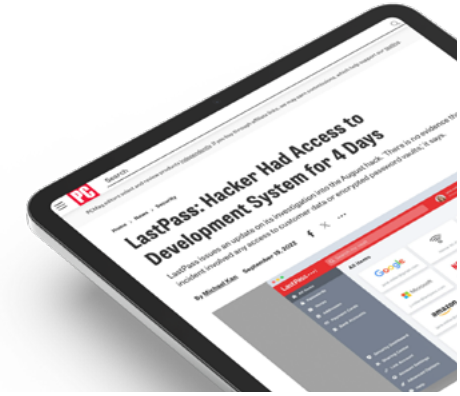

As we've seen with LastPass and CircleCI, major breaches often start with the developer's endpoint. That's why it's critical to be able to detect malware and vulnerabilities on your developers' laptops and reveal any suspicious behavior as they move code in and out of repositories and into production.

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 where they will give themselves elevated privileges, and exfiltrate a highly sensitive customer database.

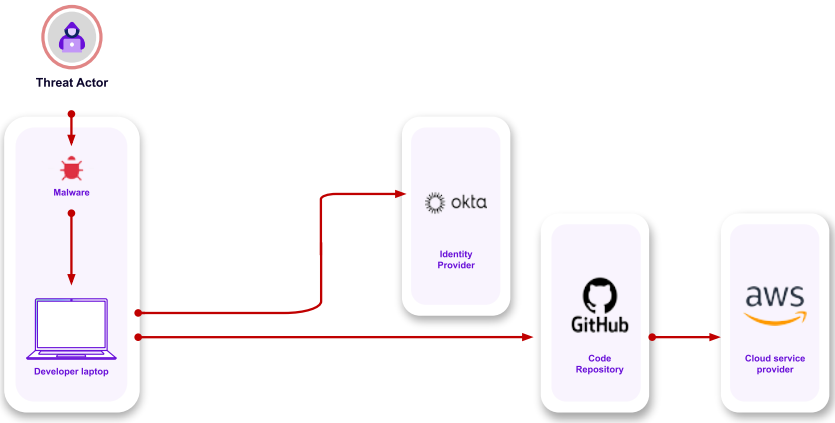

With the trend towards remote and hybrid global workforces (64% of firms outsource some portion of application development), you may also want to consider an additional visibility and access layer in the form of a Zero Trust Score and integration with a Zero Trust Network Access provider such as Cloudflare.

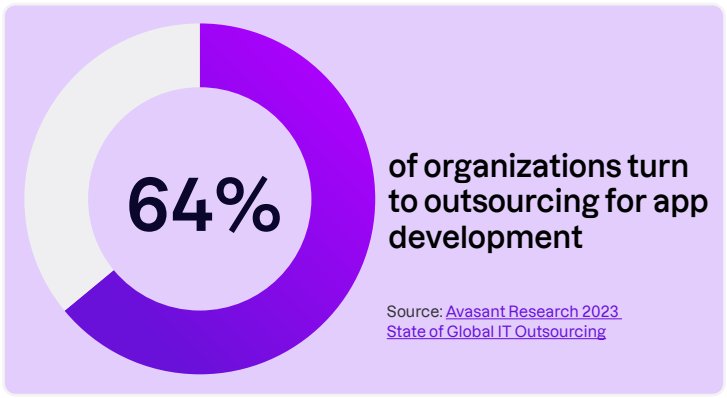

The fact is, competing CNAPPs struggle when it comes to developer endpoint security and their blast radius across your hybrid cloud and Kubernetes infrastructure.

Only Uptycs can give you the type of cloud security early warning system shown below that identifies and stops threat actors before they can access crown jewel data and services in the cloud.

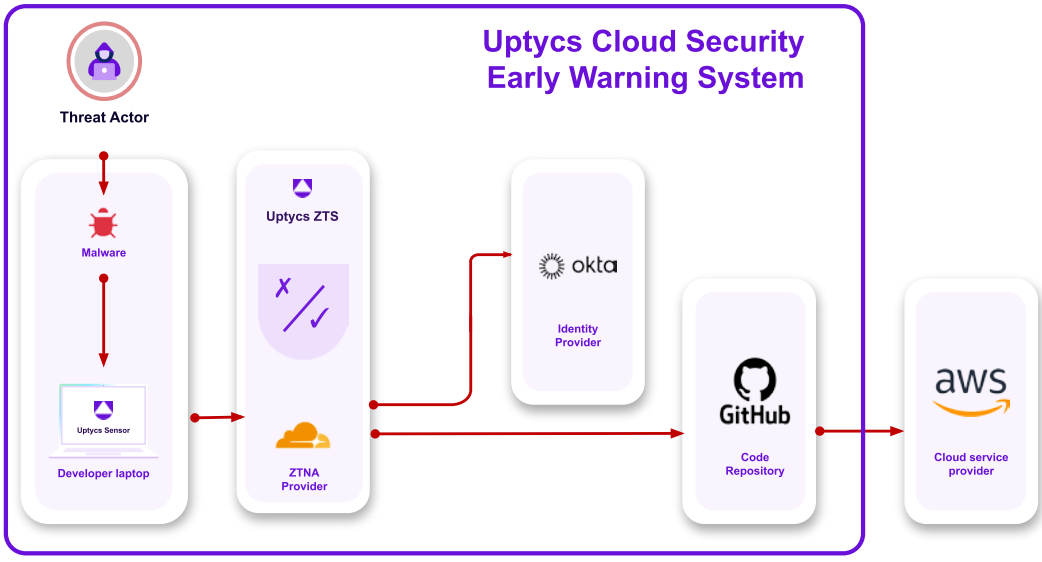

You can use Uptycs as a full-fledged XDR, have Uptycs deployed exclusively across outsourced development teams, or add it as a premium visibility and security layer on top of solutions like Crowdstrike and SentinelOne just like SEI and Lumin Digital have done. Let the endpoint protection team cover these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

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.

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 both customize existing detections and create new ones that make the most sense for your organization.

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

> "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?

Most CNAPP vendors stop at telling you what they think they want you to know. Then once they've identified an issue, you usually have limited remediation options. They are highly opinionated, which is fine, but you might also want to ask and get answers to the things you want to know.

Uptycs can significantly aid forensic investigations by leveraging its unique capability to store and analyze up to a year of historical data. This historical depth, combined with the innovative use of both the Uptycs Sensor for detailed insights, and agentless scanning for broader coverage, ensures that no stone is left unturned in investigative processes.

Advanced container workload protection capabilities further strengthen Uptycs' forensic prowess, 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 the masquerading of malicious events, such as a port scan hidden behind a different process name.

In head-to-head comparisons with competing CNAPPs such as Wiz, Uptycs' depth of telemetry, ability to query your infrastructure like database, and customizable detection logic, clearly demonstrate its superior capability in enabling robust forensic investigations.

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

Most CNAPP vendors have basic support, but won't run their platform on your behalf, or provide managed detection and response.

Uptycs offers Managed Detection and Response (MDR) services at three levels—Managed Onboarding, Managed Monitoring, and Managed Protect—each designed to augment your security team's capabilities. Starting with Managed Onboarding, Uptycs ensures a seamless integration into your environment. Managed Monitoring provides vigilant oversight, identifying and notifying about potential threats. The pinnacle, Managed Protect, delivers full-scale management, from continuous surveillance to active threat containment, ensuring peace of mind through expert monitoring and response capabilities. This tiered approach allows organizations to select the level of service that best fits their needs, enhancing their cybersecurity posture without expanding the workforce.

Additionally, Uptycs offers Professional Services to further support customers, including tailored onboarding, custom dashboard development, specialized configuration, and integration with SIEM/SOAR tools, enhancing the overall 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 there are rising data storage costs and an unwieldy number of detections to manage.

# 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

SIEMs ingest a considerable amount of data even when the total number of workloads is reduced—and that's just for workloads.

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 funneling into SIEM systems, but also decreases false alerts, paving the way for a more streamlined and cost-effective security operation. Furthermore, organizations may opt to use SOAR to pull additional data from Uptycs when more context is required for a specific investigation.

By providing a dynamic and flexible approach to security data, Uptycs ensures that security operations can be tailored to meet the specific needs of an organization, striking the right balance between maintaining comprehensive security measures and managing operational costs efficiently.

About Uptycs

Uptycs is the leading cloud security platform for large hybrid cloud environments. We protect workloads wherever they run while extending security visibility from development to runtime. That's why enterprises like PayPal, Comcast, and Nutanix rely on Uptycs to secure their mission-critical workloads.

First-generation CNAPP solutions don't have the data needed to manage and prioritize risk. With Uptycs, data is power! We have no trouble giving you deeper context so you can focus on what truly matters. Most CNAPPs can tell you if a vulnerable workload is exposed to the internet, but can they show you vulnerable packages running now, or three weeks ago, and if and how you were breached? Uptycs can.

Uptycs brings teams together to optimize security operations, ensure compliance, and accelerate remediation across cloud workloads, containers, Kubernetes, and software pipelines – all from a single security console, policy framework, and data lake.

Shift up your cybersecurity. Learn more at Uptycs.com

uptycs