



Sysdig 2022 Cloud-Native Security and Usage Report

Everyone is trying to shift left, but the reality is $\frac{3}{4}$ of running containers have at least one "high" or "critical" vulnerability.

Meanwhile, many companies adopt cloud for operational efficiency, but more than half of containers deployed have no limits, which could waste resources.

Contents



Executive Summary	3
Security and Compliance	5
Managing vulnerabilities	5
Runtime security threats	9
Cloud security and identity.	12
Containers and Kubernetes	15
Kubernetes capacity planning	16
Kubernetes clusters and nodes	17
Kubernetes namespaces, deployments, and pods	18
Services running	20
Container platforms deployed	23
Container, image, and service lifespans.	25
Alerts	27
Demographics and Data Sources	29
Conclusion	30

You can find
all of our past
reports [here](#).

Executive Summary

For the past four years, we've provided insights into container usage and security through real-time, real-world customer data. As our security and monitoring capabilities grow, our unique vantage point lets us explore how companies address the evolving risks in an increasingly cloud-native world. Each year, we uncover new patterns in how infrastructure, applications, and containers evolve over time. Armed with these insights, we bring you the Sysdig 2022 Cloud-Native Security and Usage Report.

Our customers tell us that security and compliance concerns surrounding their container environments are top of mind due to their ephemeral nature. These short lifespans can create unique challenges for incident response, forensics, and troubleshooting. Containers are but one complex piece of the cloud adoption story.

Securing modern cloud workloads requires controls around vulnerabilities, configurations, entitlements, and runtime threat detection. We see high numbers of vulnerable containers in production, many instances of exposed cloud storage, and many concerning threat detection events. However, we also note the tremendous growth in the adoption of the CNCF Falco project, which helps organizations detect runtime threats inside containers, hosts, and Kubernetes environments.

Getting accurate capacity information about Kubernetes deployments is a tough challenge due to the ephemeral nature of those environments. This report shows over 50% of containers have no CPU or memory limits defined which could lead to performance issues and cost overruns. With container density growing again this year, organizations are shifting toward Prometheus as the standard way to monitor their infrastructure and applications. The use of Prometheus metrics among our customers grew to 83% this year. The growth of Falco and Prometheus clearly indicates a preference for open source and open standards.

These findings are based on the data we gather from the billions of containers that our customers run over the course of a year. This allows us to report on many different aspects of actual usage of containers rather than rely on survey results. In this report, you will find details about security, compliance, services, alerting, and Kubernetes usage patterns. This information can be useful for determining the real-world state of security and usage for container environments at companies around the world, from a broad range of industries.

Key 2021 Trends

Cloud Security

- 88%** cloud roles are non-human
- 73%** cloud accounts have public S3 buckets
- 73%** YoY growth in Falco downloads

Container Security

- 75%** containers running with "high" or "critical" vulnerabilities
- 62%** detect shell in container events
- 76%** containers running as root

Container Usage

- 34%** unused CPU resources
- 51%** no memory limits
- 60%** no CPU limits
- 83%** custom metrics are Prometheus

Security and Compliance



As organizations move more container workloads to production, they are recognizing the need to integrate security and compliance into the DevOps workflow. “Shift security left” has become a mantra that refers to identifying flaws in software as early as possible to avoid deploying insecure workloads. Scanning for vulnerabilities is a critical component of this workflow, especially when images are pulled from public registries.

Our data also highlights the need for stringent runtime policies and continuous workload re-assessment to account for things like configuration drift, previously undisclosed vulnerabilities, or suboptimal configurations to reduce risk.

Managing vulnerabilities

Whether the container images originate from private or public registries, it is critical to scan them and identify known vulnerabilities prior to deploying into production. We assessed all images our customers deployed for OS and non-OS vulnerabilities. We found that OS images have significantly fewer flaws than non-OS images, likely because that they are usually supported and maintained by industry vendors.


OS vulnerability snapshot

We noticed that **4% of OS vulnerabilities are "high" or "critical"**. Although this may seem low, if an OS vulnerability is exploited, it can compromise your entire image and bring down your applications. Additionally, OS vulnerabilities can have a very large blast radius because many different workloads are affected at the same time.

Non-OS vulnerability snapshot

What teams don't check for are vulnerabilities in third-party libraries. We found that **56% of non-OS packages have "high" or "critical" severity vulnerabilities**. Developers might be unknowingly pulling in vulnerabilities from non-OS open source packages, like Python PIP or Ruby Gem, and introducing security risk. Gaining visibility into third-party dependencies and determining whether they are truly exploitable is a challenge. Prioritization using runtime context allows teams to address the most urgent vulnerabilities first.

OS Vulnerabilities by Severity



Non-OS Vulnerabilities by Severity




Scanning in build phase vs. runtime

DevOps teams are “shifting left” and enabling testing earlier in the development lifecycle. Security is key in this process. We looked at where in the workflow images are scanned for the first time. Ideally, packages are analyzed as part of the build phase in the CI/CD pipeline and no build can proceed to production without meeting certain security thresholds. **52% of all images are scanned in runtime, and 42% are initially scanned in the CI/CD pipeline.**

An important caveat is that the vast majority of runtime scans are not of custom packages built by DevOps teams, but of images that contain third-party software downloaded from a vendor. This can

include Kubernetes components, commercial products deployed as containers, or open source software. Scanning third-party containers in runtime is prevalent because it's simpler and resembles the legacy approach to vulnerability management. However, the best practice here is still to “shift left” and scan the corresponding images as part of the infrastructure-as-code checks in the CI/CD pipeline. Sysdig ensures that all containers are continuously rescanned post-deployment to discover any newly disclosed vulnerabilities, but scanning as much as possible pre-deployment will prevent known flaws from appearing in production at all.

Where Images are Scanned



Vulnerable images in production

How often are risky deployments actually blocked? Even when images are thoroughly scanned for flaws prior to deployment, the vulnerabilities are not immediately addressed. Organizations must consider the tradeoff between delaying deployments to fix the problems or accepting security risks to release the software faster. We found that 85% of images that run in production contain at least one patchable vulnerability. Furthermore, **75% of images contain patchable vulnerabilities of "high" or "critical" severity**. However, Sysdig's most mature and meticulous customers have reduced this metric to below 5%.

Tens of thousands of vulnerabilities are discovered by Sysdig in customer environments every day. It's impossible to fix every single one, even if we limit the scope to only "high" and "critical" severity. Teams must prioritize the problems that pose the greatest risk to the organization by considering additional factors, like the workload's business criticality and which components are actually exploitable during runtime.

Patchable Vulnerabilities in Runtime



Public and hosted container registries

Container registries provide repositories for hosting and managing container images. For the first time, Quay has overtaken Docker, now accounting for 26% of customer adoption. This measure includes both

privately hosted and public repositories. Registry solutions hosted by cloud providers are also increasingly popular, with Red Hat and AWS registry usage doubling compared to last year.

Container Registries



We found that **public sources are being trusted more and more, with an increase from 47% last year to 61% this year**. Using public registries poses a risk because few are validated or checked for vulnerabilities. In some cases, the convenience of using public repositories may outweigh the risk, but the best practice is to enforce explicit policies about which registries are approved for use in the organization.

Images Pulled from Public vs. Private Registries



Runtime security threats


As container adoption soars, organizations do their best to incorporate best practices into their pre and post-production activities. Vulnerabilities are only one part of the cloud security program. Even on fully remediated workloads, insecure configurations can create serious security problems. Furthermore, once workloads are deployed into production, teams need to be able to detect anomalous behavior and trigger security alerts in real time to respond to potential threats.

"As the number of security issues continue to increase, it's promising to see Falco's adoption momentum, indicating the community's commitment to improving runtime security. As more people get involved with the project, cloud and container security is strengthened by the collective group working together against bad actors."

- Chris Aniszczuk, Chief Technology Officer, CNCF

The growth of Falco

Falco, the CNCF open source project contributed by Sysdig, is showing huge amounts of momentum and interest again this year. **The project now has over 40 million Docker Hub pulls**, which represents 370% growth since becoming an Incubating project in January 2020. Falco enables the definition of runtime policies that detect security violations and generate alerts. Falco users can leverage Sysdig Secure to automate rule creation and tuning, which leads to faster violation detection and resolution. Innovative contributions to the project this year expanded its capabilities beyond containers to broader cloud security use cases.



Containers running as root

Although teams understand the need to scan for vulnerabilities, they may not be scanning for common configuration mistakes. What we see is that **76% of images ultimately run as root**, allowing for privileged containers that can be compromised. From talking to our customers, in practice, even if risky configurations are detected at runtime, teams don't stop containers in order to continue deploying quickly. Instead, they run within a grace period and then decide on the remediation step. Additionally, such high numbers of workloads not abiding by best practices may be indicative of the increasingly broad adoption of container technologies by organizations that have not yet evolved their DevSecOps processes to accommodate the new operational model.




Container security alerts generated

Security monitoring of containerized workloads aims to discover and impede malicious actors, as well as to identify risky exposures and reduce the attack surface. Sysdig customers regularly receive Falco security alerts indicative of poor security hygiene in these environments.

- 62% detected a terminal shell in a container, indicating that these workloads are not being treated as immutable and increasing the risk of tampering.
- 38% detected a container launched with sensitive mount points, meaning that the container is able to alter important files on the host system.
- 28% detected the launch of privileged containers, which means the container has root capabilities of the host machine.

Container Runtime Security Alerts



Although containers are a great fit for immutable microservices, some organizations deploy them for other use cases. For example, an application that was containerized for cloud migration but not yet refactored will behave more like a traditional virtual machine. Privileged status, sensitive mount points, and terminal shell activity in these types of workloads are traditionally more acceptable. As organizations mature their cloud-native capabilities, they should strive to eliminate or reduce the insecure behaviors that lead to these alerts.

"When the news on log4j came out, we received calls from our customers asking what the impact was. Using Sysdig Secure, we were able to find out in less than 5 minutes what the potential risk would be, which was zero. Also, we were able to get communications out in under an hour and proactively reduced our customers' anxiety about the potential issue."

- Sam Brown, Product Security Manager, Expel

Cloud security and identity


Containerized workloads are right at home in cloud environments because the same governing philosophies of immutability, scalability, and everything-as-code apply. However, there is a lot more to worry about when securing the cloud environment than the containers themselves.

Identity and access

Most security incidents in the cloud occur due to misconfigurations. These can include granting excessive privileges, unintentionally exposing assets to the public, or neglecting to change weak default configurations.

The root user is the most privileged user account. Cloud security best practices and the CIS Benchmark for AWS indicate that organizations should avoid using the root user for administrative and daily tasks, yet 27% of teams continue to do so. Creating dedicated roles with appropriate permissions for performing administrative tasks is much less risky than leveraging the root user account. Furthermore, **48% of organizations don't have multi-factor authentication (MFA) enabled** on this highly privileged account, which makes it easier for attackers to compromise the organization if the account credentials are ever leaked or stolen.

Cloud Identity and Access



Users and roles


Identity and access management is a key control in cloud environments. Sometimes, it is the only control standing between an attacker and your critical assets and sensitive data. The sheer number of permissions that need to be managed increases significantly in modern environments. In addition to human users, applications, cloud services, commercial tools, and many other entities require access. Understanding and harnessing this complexity requires new skills and tools.

"Moving to the cloud is an opportunity to do away with some technical debt and push for stronger efforts at best practice across the board, but the learning curve can be quite steep. Our focus has been on improved security, better resiliency, and reduced planned downtime, and we are seeing considerable success in those areas."

- Charles Jones, Director of Information Security at NCSOFT

One interesting element of cloud access management is that there are many different types of identities. **Only 12% of roles in organizations' cloud environments are assigned to human users.** The non-human roles may be assumed by users to perform certain specific tasks, or they could be used by applications, service providers, or other third parties. A best practice is to follow the principle of least privilege and explicitly assign the minimum necessary permissions to each role. Unfortunately, most users and roles sacrifice security by granting excessive permissions because it's easier and faster to operate this way.

Cloud Users and Roles



Configuration and compliance

Compliance is a security driver for most organizations, especially those in highly regulated industries. However, teams are often unaware of configuration mistakes that lead to their cloud environments falling out of compliance. Security best practices, like CIS benchmarks, or specific compliance standards, like PCI DSS, include managing configurations.

Risky configurations are common in the cloud, partially due to the frequent changes that occur in these highly dynamic environments. We found that **73% of cloud accounts contained exposed S3 buckets** and **that 36% of all existing S3 buckets are open to public access**.

In AWS, buckets are created with public access disabled by default, but it's often convenient to change this setting, whether for temporary testing or long-term ease of use. The amount of risk associated with an open bucket varies according to what kind of data is stored there. Organizations should manage access based on data sensitivity and specific use case, while abiding by the principle of least privilege.



36%

of all S3 buckets are open to the public

73%

of cloud accounts contain publicly exposed S3 buckets

"We use Sysdig to identify and alert us to suspicious activity and misconfigurations, and more generally, workloads that may cause security risk."

- Security Engineer, BlaBlaCar

Containers and Kubernetes


Metrics, usage, and adoption

Each year, we take a look at details specific to the count and activity around containers and Kubernetes, including density and lifespans. This provides insight into the rate of adoption, but also illustrates the scale and efficiencies being achieved. In this section, we also answer questions like: How many clusters are organizations operating? How many Pods run per node? How much capacity does a cluster use? We look at a range of details about what teams are doing with Kubernetes. Because Sysdig automatically collects Kubernetes labels and metadata, we're able to provide cloud-native context for all of the data insights we discover, from performance metrics and alerts to security events. This same capability enables us to capture each of the following usage metrics from the cluster all the way to Pods and containers, all with a simple query.

Kubernetes capacity planning

In an ephemeral, dynamic environment like Kubernetes, capacity management and planning are inherently difficult. Limits on how many resources a container can use often go undefined. In addition, environments where developers are allowed to choose their own capacity needs can lead to over allocation and these are rarely rightsized. In looking across the customers in our largest region, we found that **60% of containers had no CPU limits defined and 51% had no memory limits**

defined. Of those clusters that did have CPU limits, an average of **34% of CPU cores were unused.** Without knowing the utilization of their clusters, organizations could be wasting money due to overallocation or causing performance issues by running out of resources. Given the average cost of Amazon Web Services CPU pricing, an organization with 20 Kubernetes clusters could be spending up to \$400,000 per year more than they need to due to underutilized CPU resources.




Kubernetes clusters and nodes


Some organizations maintain a few large clusters while others have many clusters of varying sizes. The charts in this section show a distribution of cluster count and nodes per cluster for users of the Sysdig platform. The large number of single clusters per customer, and relatively small number of nodes, is an indication that many enterprises are still early

in their use of Kubernetes. The use of managed Kubernetes services in public clouds is another factor that impacts these data points. **This year, we observed a shift towards more clusters overall and more nodes per cluster.** Such shifts may indicate that cloud-native deployments are starting to mature by utilizing more resources.

Number of Clusters



Number of Nodes per Cluster




Kubernetes namespaces, deployments, and pods

We saw a **slight shift this year toward more namespaces per cluster** and more Deployments per namespace. Again, this may indicate a maturing of these cloud-native environments.

Namespaces

Kubernetes namespaces provide logical isolation to help organize cluster resources between multiple users, teams, or applications. Kubernetes starts with three initial namespaces: default, kube-system, and kubepublic. How namespaces are used varies across organizations, but it is common for cloud teams to use a unique namespace per application.


Namespaces per Cluster



Deployments per namespace

Deployments describe the desired state for Pods and ReplicaSets and help ensure that one or more instances of your application are available to serve user requests. Deployments represent a set of multiple, identical Pods with no unique identities, such as deployments of NGINX, Redis, or Tomcat. The number of Deployments per namespace provides an idea of how many services compose our users' microservices applications.


Deployments per Namespace



This year we saw a significant increase in the number of Pods per cluster, with **54% of organizations running more than 100 Pods per cluster** compared to only 19% last year. However, the average number of Pods per node fell, indicating that teams are deploying more smaller nodes to handle their workloads.


Pods per cluster

Pods are the smallest deployable object in Kubernetes. They contain one or more containers with shared storage and network, as well as a specification for how to run the containers.



Pods per node


A Pod remains on a node until its process is complete, the Pod is deleted, the Pod is evicted from the node due to lack of resources, or the node fails.



What services are customers running?

The top open source solutions running in containers

Open source software has changed the face of enterprise computing. It powers innovation across not just infrastructure, but especially application development. Sysdig has the ability to auto-discover the processes inside containers gives us instant insight into the solutions that make up the cloud-native services that teams run in production. Below are the **top 12 open source technologies deployed** by organizations:



The 2021 list includes a wide range of services, each critical to the function of modern applications in its own way:


- HTTP server and reverse proxy solutions — NGINX
- NoSQL, relational, and in-memory database solutions — MongoDB, Postgres, and Redis
- Logging and data analytics — Elasticsearch
- Programming languages and frameworks — node.js, Go, and Java/JVMs
- Message broker software — RabbitMQ

Given the wide range of options available in the open source community, it's surprising that the services in our list have remained fairly consistent over the past three years. We purposely omitted Kubernetes components like etcd and fluentd, as well as Falco. Because these are deployed by default, they end up at the top of the list for every Kubernetes user. This year we saw that NGINX and Go (aka golang) continue to be among the top open source projects used when developing cloud-native applications. The top 12 solutions above are widely deployed and trusted services. If you're in the market for similar services, you can't go wrong with taking advantage of what these open source solutions offer. There is, however, a long tail of software solutions available.

Custom metrics

Custom metric solutions give developers and DevOps teams a way to instrument code to collect unique metrics. This approach has become a popular way to monitor applications in production cloud environments along with tracing and log analysis. Of the three mainstay solutions, JMX, StatsD, and Prometheus, it was Prometheus that gained for the third year in a row. Year-over-year, **Prometheus metric use increased to 83% compared to 62% last year**. As the use of new programming

frameworks expands, alternatives like JMX metrics (for Java apps) and StatsD continue to decline, with JMX dropping precipitously to only 4% this year compared to 19% last year. It is clear that with the strong connection between Prometheus and Kubernetes, more organizations are adopting Prometheus metrics as they move toward cloud-native architectures.





Top Prometheus exporters

One of the most successful open source projects to emerge from the CNCF, Prometheus has become synonymous with cloud-native monitoring. It is now widely adopted as a metric standard in projects like Kubernetes, OpenShift, and Istio. In addition, an increasing number of “exporters” are available to provide metric output for a wide range of third-party solutions. We expect the popularity of Prometheus to continue its growth within our customer base, particularly as Sysdig now offers full Prometheus compatibility for large-scale environments.

For this ranking, we looked at each github project listed on prometheus.io, measured the number of issues, stars, and forks for each, and correlated the results against the number of Dockerhub or other repository pulls.

Top 10 Prometheus Exporters


Name	Maintainer
node_exporter	prometheus
blackbox_exporter	prometheus
jmx_exporter	prometheus
redis_exporter	oliver006
windows_exporter	prometheus-community
postgres_exporter	wrouesnel
elasticsearch_exporter	justwatchcom
mysqld_exporter	prometheus
kafka_exporter	danielqsj
snmp_exporter	prometheus

What container platforms are being deployed?


Runtimes and orchestration

The breakdown in runtime usage did not change significantly this year, with Docker still being the most used container runtime among organizations. However, it dipped below 50% for the first time in the last five years that this report has been published. At the same time, container orchestrators have almost entirely consolidated down to one. If you include orchestrators based on Kubernetes like Red Hat OpenShift, **Kubernetes is used 96% of the time.**

Runtimes




Orchestration




Containers-per-organization

To get a sense of the scale at which enterprises are currently operating, we looked at the number of containers each one runs across their infrastructure. Over half of the organizations run 250 or fewer containers. At the high end, only 6% are managing more than 5,000 containers. It is common for adoption to begin at a small scale, sometimes born from developers who push for containerization as a means to accelerate software delivery. DevOps and cloud teams report that once the benefits are proven, adoption accelerates as more business units look to onboard to the new platform. However, this year showed movement toward an overall increase in the number of running containers, which may indicate that more and more workloads are moving to containers and away from traditional architectures.




Base image OS

Most people use a base image because it's easier than creating your own. We can see that **RHEL, which includes the UBI (Universal Base Image), is by far the most popular at 36% of base images used.** This may be because RHEL has a long history of usage in the enterprise and would be an easy choice as organizations move to cloud-native workloads. Interestingly, **only 25% use Alpine.** By using slimmed-down base images like Alpine, organizations can debloat their container environment and reduce their attack surface.



Container density

Over the past five years, the median number of containers per host increased in every report. However, this year showed only a 12% increase year-over-year compared with the 33% increase of last year. It is possible that the number will continue to increase slightly in the future, but that density will probably come at the cost of overall image size. While the primary goal of containers is to speed development and deployment, many organizations are benefiting from increased utilization of hardware resources thanks to container efficiencies.




Container, image, and service lifespans

The measure of how long containers, container images, and services live was one of the most popular data points from our 2021 report. It reflects just how dynamic modern applications are from both a development and a runtime perspective.

The short life of containers

Comparing container lifespans year-over-year, we see a similar pattern where the vast majority of containers are alive for less than a week.

About 44% live less than five minutes!




Many containers only need to live long enough to execute a function and then terminate when it's complete. Seconds may seem short, but for some processes, it's all that is required. The ephemeral nature of containers remains one of the technology's unique advantages, in that things are designed to change as needed. However, it also presents new issues to consider for monitoring, security, and compliance because many monitoring and security tools can't report on entities that no longer exist.

Continuous development and image lifespans

Containers are a perfect companion to the agile movement, accelerating the development and release of code, often as containerized microservices. Our image lifespan data reflects the shift in the time between code releases and the reality that CI/CD pipelines are helping developer teams deliver software updates at a faster cadence than ever before. The data shows that about **half of container images get replaced — also known as churn — in a week or less**. For most, if not all, of today's businesses, speed to market matters and makes all the difference in maintaining competitiveness. Code is being deployed more frequently, which creates new container images. Containers give businesses what they need to turn great ideas into reality, fast.


Container Image Lifespans



Service lifespan

Services, the functional software components of our applications like database software, load balancers, and custom code — are continuously being improved. However, at the same time, it's important to keep services up and running around the clock to be able to meet customer expectations. The data show that service lifespans have remained relatively consistent compared to last year.

Service Lifespans




Alerts

Analysis of trends with the types of alerts set by teams helps us understand the kind of conditions that our users identify as having the most potential for disruption to their container operations.

The top 10 alert conditions

There are more than **800 unique alert conditions** being used across our customers today. The graphic below represents the most commonly used alert conditions, along with the percentage of customers using each. Kubernetes.node.ready continues to be the most used, along with important resource and uptime metrics.


Top 10 Alerts



Alert channels

We looked at the communication channels users have configured to receive alerts. **Slack took the top position, greater than purpose-built incident response platforms and even email.** We find the results interesting because unlike PagerDuty and Opsgenie, for instance, Slack is not considered an incident response platform. It's likely that Slack is being used for non-critical alerts handled during normal work hours, while solutions like PagerDuty are being used for "waking people from bed."


There are a number of alerts that do not have a notification channel configured, but this isn't necessarily a bad thing. This could be because the alert is for informational purposes only, or because the Sysdig platform itself provides enough information to satisfy the demands of the alert in question.



Demographics and Data Sources

The data in this report is derived from an analysis of more than 3 million containers that our customers are running at a given moment. We also pulled from public data sources like GitHub, Docker Hub, and the

CNCF. The data originates from container deployments across a wide range of industries and regions, with organizations ranging in size from mid-market to large enterprise.



Conclusion

Cloud technologies continue to expand their role in transforming how organizations deliver applications. With security becoming a growing concern among DevOps teams, it is good to see that teams are implementing security during the build process. However, more work is needed to secure both containers and cloud services to prevent possible vulnerabilities from entering production. Runtime threat detection will continue to be critical to securing the cloud, as even the most robust programs will not address all software vulnerabilities and misconfigurations. The key trends from our fifth annual report highlight the continued growth in container environments, and the growing dependency on open source-based solutions to run them:

- Broader adoption of cloud and containers results in many insecure behaviors like containers running as root and cloud accounts having excessive privileges. However, continued growth of Falco indicates that organizations are leveraging cloud security tools to try to reduce risk.
- Capacity planning is difficult in Kubernetes environments. Only with proper container resource limits and continuous monitoring can organizations ensure they are not overspending or risking performance issues for their applications.
- Kubernetes and container environments continue to grow as organizations move their workloads to the cloud. Using open standards like Prometheus and adhering to security best-practices are critical behaviors that will increase visibility and reduce risk for cloud-native applications.

Thank you for reading the Sysdig 2022 Cloud-Native Security and Usage Report. We look forward to following and documenting the evolution of the container market in the coming year. See you then!



Sign up for a free 30-day trial and get visibility for security, compliance, and monitoring in minutes.

LEARN MORE

www.sysdig.com