



Five container development security risks & how to prevent them



Five container development security risks & how to prevent them

The digital landscape is a battlefield of cyber threat actors relentlessly exposing vulnerable government agencies to new security risks. In the second half of 2022, cyber attacks against the government sector increased by 95%¹. These threats aren't going anywhere. Agencies continue to be a target to a growing number of malicious actors, increasing the need for keeping software secure and supply chains clean. By integrating robust security measures throughout the software development lifecycle, agencies can reduce their vulnerability and strengthen operational resiliency. Secure development plays a pivotal role in this endeavor by enabling the creation of safe, reliable, and maintainable software systems.

¹ <https://www.csoonline.com/article/574275/cyberattacks-against-governments-jumped-95-in-last-half-of-2022-cloudsek-says.html>



Why secure containers?

A secure supply chain starts with secure development. Container development tools have revolutionized the way software is packaged, shipped, and executed. These tools provide developers with an efficient and scalable solution to build applications and their dependencies in an isolated environment.

Proactive container development security means stopping anything malicious from getting into your developer environments in the first place. This is where using container technology becomes advantageous. Alongside the amplified scalability, flexibility, and reproducibility that containers offer, in general, running software inside a container is more secure than running it natively on a host. This is because the container is a smaller attack vector and the host is isolated. Most container development tools already come with baseline security features to protect against common security risks. These baseline container security features:

- Allow users to run images unrestricted (no admin rights needed).
- Run all containers in a Linux Virtual Machine (VM) to isolate them from the underlying Mac/Windows/Linux host.
- Prevent containers from accessing host files not owned by the user.
- Offer configurations for network proxies, registries, tool updates, host file sharing, and more.

Container security risks

While container development tools offer baseline security functionality, if a container is breached, it can lead to unauthorized network access, data breaches, and compromise of the entire system. Understanding how you can be breached and how to combat those risks is essential to ensuring the integrity, confidentiality, and availability of containerized applications and their underlying infrastructure.

The top five security risks we see to container development tools today are malware threats, supply chain attacks, inadequate access control, misconfigurations, and insider threat.

RISK #1

Malware threats

If malware enters your network infrastructure, it can exfiltrate sensitive data and credentials and compromise the container and its underlying host system. The most common way for malware to get into your systems is when developers inadvertently download a malicious container image on their local machine and run it, or when a developer exposes the container engine API on the network without protection.

Developers often rely on pre-built images from public repositories or create their own images. Unfortunately, these images can be vulnerable or contain outdated software components. Free and open source software has its advantages, but also means those container images can contain vulnerabilities. While one advantage of using open source software is that developers can achieve

² <https://sysdig.com/2023-cloud-native-security-and-usage-report/>



faster release cycles, one significant disadvantage is that organizations may not collaborate at the speed needed to maintain a secure development life cycle (SDLC) practice. In fact, an alarming 87% of container images running in production have critical or high-severity vulnerabilities, up from 75% a year ago. ²

The same malware threats apply for public registries. If these registries lack proper access controls or authentication mechanisms, they are susceptible to image tampering. An attacker gaining control over a compromised container registry can distribute malicious images leading to widespread security breaches.

RISK #2

Supply chain attack

Supply chain attacks aim to compromise the integrity of containerized applications or their dependencies. Attackers may inject malicious code or tamper with the software supply chain, leading to the distribution of compromised containers.

Once an attacker successfully breaks out of a container, they can move laterally within the containerized environment, escalate privileges, access sensitive resources, and launch further attacks. Container breakout and lateral movement attacks can have severe consequences, jeopardizing the security and integrity of the entire container infrastructure.

RISK #3

Inadequate access controls

Without robust access controls, malicious actors can gain unauthorized access to sensitive containers, manipulate their configurations, or exfiltrate sensitive data. Insufficient authentication practices can lead to unauthorized access to container orchestration systems, management consoles, or containerized applications, allowing attackers to exploit vulnerabilities and compromise the entire environment.

If a developer has too many local admin rights, they can relax firewall rules, turn off anti-malware software, disable updates, and install untrusted software.

If a developer's workstation is not configured securely per organizational policies, attackers can much more easily compromise these machines and from there, the organization's network.

RISK #4

Misconfigurations

Improperly configured containers, orchestrators, or network settings may allow unauthorized access, data leaks, or unintended exposure of sensitive information. Attackers actively scan for misconfigurations and exploit them to gain control over containers or compromise the underlying infrastructure.

Misconfigurations often happen because there is no centralized or mandatory organizational control over the tool's security settings. Even if an IT admin initially configures the tool's settings per organizational security policy, a developer can always modify these settings at will.

RISK #5

Insider threat

Insider threats don't always have to have ill intent behind them. These risks can come from anyone who has access to container development tools intentionally and unintentionally misusing their priv-



ileges. Proper access controls, monitoring, and security awareness training are useful in mitigating insider threats and preventing data breaches, unauthorized access, or malicious modifications to containerized applications. However, many organizations have a long way to go to mitigate insider threats and embrace cross-team security training: 77% of organizations acknowledged that they lack effective collaboration among developers and security teams.³

How to mitigate these risks

In recent years there has been increasingly heightened attention on putting guardrails in place to further secure the software development lifecycle. At the end of 2022 the Executive Office of the President released a memo addressing the heads of executive departments and agencies on enhancing the security of the software supply chain through secure software development practices⁴. The memo guides agencies to obtain software that conforms to secure software development practices. It also acknowledges that agencies may need to require a software bill of materials (SBOM) in solicitation requirements, based on the unique criticality of the software. SBOMs are an essential component to understanding the provenance of where an artifact came from and employing SBOMs is just one important step to consider in a thorough plan to mitigate risk. Here are some additional steps you can take to prevent security risks in your container development tools:

1. Start with secure base images - begin your container development with secure and trusted base images from reputable sources. These images should be regularly reviewed for vulnerabilities and updated to include the latest security patches and fixes.
2. Receive real time vulnerability updates - Regularly index container images with known vulnerabilities, ensuring timely identification and remediation of security risks.
3. Employ image signing and verification - Ensure the integrity and authenticity of container images using digital signatures.
4. Restrict secure container registries - Implement strong access controls to only allow access to trusted sources for container images and limit access to non-authorized registries. You want to ensure accessible registries are regularly monitored for vulnerabilities and malware.
5. Enforce access controls - Use role-based access controls (RBAC) to grant permissions based on job responsibilities.
6. Fortify network isolation - Minimize the impact of a potential breach/container breakout or lateral movement within the containerized environment by hardening container isolation. Regularly patch and update container runtimes, host systems, and underlying dependencies.

Resolve security issues before they make it into production

All of the measures listed above work great for baseline container security, but in organizations with stringent security and compliance requirements, sometimes you need a bit more. You need a mechanism for IT admins to enforce security settings on the container development platforms (e.g., proxies, registries, Docker Engine configs) with stronger guarantees that won't be easily bypassed by developers or malicious containers.

³ <https:// snyk.io/reports/state-of-cloud-security/>

⁴ <https://www.whitehouse.gov/wp-content/uploads/2022/09/M-22-18.pdf>



Advanced isolation measures

Strong container isolation within the VM makes it that much harder for malicious users, or workloads running in the container, to compromise the host. Container isolation mechanisms should be applied automatically so they do not impact developer workflows and teams can continue using Docker as usual.

Hardened Docker Desktop (HDD) provides an additional layer of security for Docker Desktop users that further fortifies container isolation. With HDD, containers inside the Docker Desktop Linux VM run unprivileged (i.e., root user in the container maps to an unprivileged user inside the VM) to prevent them from taking control of the VM. This keeps containers from bypassing Docker Desktop registry configurations, accessing sensitive data or secrets stored in the VM, accessing the Docker Engine, peeking into other containers, or attacking the host machine.

In addition, HDD's Settings Management feature allows admins to preset and "lock" several Docker Desktop configurations on a developer's machine. Once set, these configurations cannot be altered by developers.

Access management

Access controls and granular permission settings ensure the security and confidentiality of sensitive information. This includes the types of images and repositories developers are allowed to pull. As open source software continues to make up the majority of our applications, developers benefit from these guardrails to significantly reduce the chance of using a container image with a malicious payload.

Docker's Trusted Content provides vetted, reliable, and trusted container images for your teams to build from. When used in combination with Registry Access Management and Image Access Management, you can ensure that your developers are only using images that fit agency compliance requirements.

Trusted Content includes:

- Docker Official Images - Container images curated and maintained by Docker, providing a trusted and reliable foundation for building containers. These images exemplify container image best practices.
- Docker Verified Publisher images - Container images coming from trusted publishers and their repositories. Each repository is vetted by Docker.
- Docker-Sponsored Open Source images - Community-driven projects supported by Docker.

Vulnerability visibility

Proactively identifying and mitigating security risks starts during development. Visibility into vulnerabilities will not only help agencies maintain compliance with necessary security standards and regulations, but will also improve the overall security posture of your software systems, ultimately reducing the likelihood of a successful attack.

Docker Scout helps developers find and fix container vulnerabilities at the earliest stages of software development. It provides a unified, layer-by-layer view of software dependencies, their known vulnerabilities, and recommended remediation paths. Docker Scout includes a software bill of materials (SBOM) that integrates seamlessly with any existing CI/CD pipeline or build process. As a result, admins maintain a verifiable record of their containerized software components.



Preventing gov breaches with Docker

The modern era demands governments to fortify their operations and infrastructure against an ever-growing array of cyber threats. Secure development tools provide a way to create resilient software systems while mitigating the risk of breaches. By adopting preventive measures and implementing best practices, organizations can strengthen the security posture of their containerized environments, ensuring the integrity, confidentiality, and availability of their applications and underlying infrastructure.

[Learn more](#) about these security mechanisms.

