



Technical Paper

# Container Runtime Security *Comparative Insights*

2025 Edition

---

Evaluating Detection, Response, and Prevention  
Capabilities Across Falco, KubeArmor, Tetragon,  
NeuVector, and More

---

Authored By:



**Rahul Jadhav**

SIG Security Chair, Nephio  
CTO, Cofounder, AccuKnox

# Table of Contents

<b>Introduction.....</b>	<b>3</b>
Characteristics of a Runtime Security Solution.....	4
Detection Capabilities.....	4
Response Capabilities.....	5
TOCTOU issues.....	6
Overwhelming the events causing the events to drop.....	7
Advantages of Detect and Response Model.....	7
Disadvantages of Detect and Response Model.....	7
Prevention Capabilities.....	8
eBPF and its role in preventive capabilities.....	8
Prevention Capabilities and Zero Trust.....	9
Sandboxing Capabilities.....	9
Sandboxing techniques.....	10
Performance Impact.....	10
Ease of Deployment on Hardened Distributions.....	11
Ease of Runtime Policy Enforcement.....	11
Policies Adhering to Zero Trust Principles.....	11
<b>Runtime Security with Detect and Response.....</b>	<b>12</b>
Issue 1: Killing a process is not an effective remediation strategy.....	12
Issue 2: The Response depends on the successful execution of a chain of actions	13
<b>Falco Analysis.....</b>	<b>13</b>
Policy Enforcement.....	14
<b>Tetragon Analysis.....</b>	<b>14</b>
Policy Enforcement Mechanics.....	15
<b>NeuVector Analysis.....</b>	<b>15</b>

Policy Enforcement Mechanics.....	16
<b>Palo Alto Prisma/TwistLock Analysis.....</b>	<b>16</b>
Policy Enforcement Mechanics.....	18
<b>gVisor Analysis.....</b>	<b>19</b>
Policy Enforcement Mechanics.....	19
<b>KubeArmor Analysis.....</b>	<b>20</b>
Policy Enforcement Mechanics.....	20
<b>Case Studies.....</b>	<b>21</b>
File Integrity Monitoring/Protection using different tools.....	21
FIM using Falco.....	21
FIM using Tetragon.....	21
FIM using KubeArmor.....	21
<b>Summary.....</b>	<b>22</b>


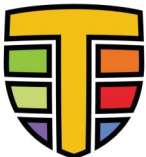
*This guide offers a technical analysis of container runtime security tools, comparing their detection, response, and prevention capabilities. It covers key architectures, TOCTOU issues, and zero-trust principles while examining tools like Falco, NeuVector, and KubeArmor. Ideal for security practitioners looking to secure containerized environments effectively.*





## Introduction

This technical guide aims to compare/contrast the fundamental architectures of the different container runtime security tooling such as Falco, KubeArmor, Tetragon, Tracee, and NeuVector, and understand the primitives used under the hood. Please note that the technical guide focuses only on container runtime security.

At a broader level, the tools can be categorized into:

1. Detection only
2. Detect and Respond
3. Preventive engines with Inline enforcement. Most inline enforcement tooling also has an observability mode that allows one to trace/log the system events.

Tool	Type	Remarks
<a href="#">Falco</a> 	Detect and Respond	Falco is a detection engine. Sysdig recently open-sourced <a href="#">Falco-Talon</a> that provides a response engine on top of Falco. There is no inline mitigation capability.
<a href="#">Tetragon</a> 	Detection + Enforcement	Tetragon provides a detection engine based on eBPF. Tetragon also provides enforcement capabilities using: <ul style="list-style-type: none"><li>• <a href="#">bpf_send_signal()</a> to kill the process.</li><li>• It also uses <a href="#">bpf_override_return()</a> to handle inline mitigation. However, override return is mostly turned off in production env and is not reliable or advised to be used.</li></ul>

<p><a href="#">Tracee</a></p> 	<p>Detection only</p>	<p>Open-source Tracee provides detection only.</p>
<p><a href="#">KubeArmor</a></p> 	<p>Detection + Inline Mitigation</p>	<p>KubeArmor provides a detection engine based on eBPF. KubeArmor uses LSMs (LSM-BPF and AppArmor) for inline mitigation.</p>
<p><a href="#">NeuVector</a></p> 	<p>Detect and Respond</p>	<p>The only engine that does not leverage eBPF. It uses <a href="#">inotify/fanotify</a>. It provides the ability to kill the process from userspace.</p>
<p><a href="#">Palo Alto TwistLock</a> (not open source)</p> 	<p>Detection + Inline Enforcement</p>	<p>TwistLock Defender <a href="#">replaces the original runc with a runc shim binary</a> to achieve runtime blocking rules enforcement. This <a href="#">results in several issues</a> operating on hardened distributions such as Bottlerocket, GKE COS, etc.</p>

**Disclaimer:** The author is one of the maintainers of KubeArmor and the cofounder of [AccuKnox](#). The aim is to dig deeper into the specific primitives used by different runtime security engines for enforcement. The guide is deeply technical and most of the research needed for this technical guide was done by me. If you disagree with any of the statements, please contact me at [rahul@accuknox.com](mailto:rahul@accuknox.com).

## Characteristics of a Runtime Security Solution

### Detection Capabilities

Most of the Runtime solutions depend on eBPF to get runtime visibility across process executions, file system accesses, and network accesses. The most obvious eBPF hooks to target are [kprobes](#), [kretprobes](#), and [tracepoints](#) which are fairly easy to use. There are a few systems-related challenges, for example, getting an absolute path of the file object. [Some solutions](#) depend on the use of LSM hooks to get the full

path and these mechanisms are pretty well-understood now. The other obvious challenge is that all the expected eBPF hook points or capabilities might not be equally applicable across all Linux distributions and across all platforms (x86, ARM, etc).

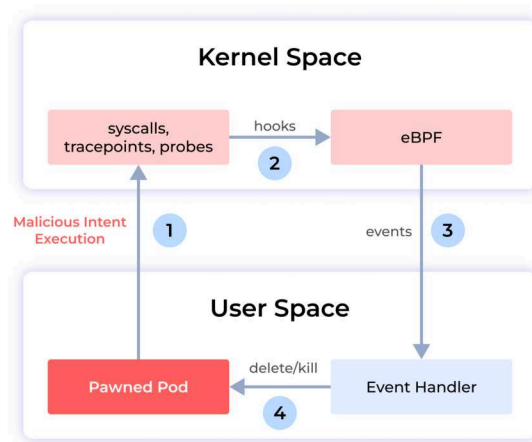
One more challenge when it comes to detection is the ability to ship all the events. Some operations might cause a burst of system activities and in a lot of cases, it might not be possible to ship all of these raw events to the analysis engine in the cloud. It is relatively simple to overwhelm the event loop that operates on kernel perf/ring circular buffer such that events are lost. It is very [common to find solutions \(ref2\)](#) implementing ways to tell users that eBPF events are lost in the circular perf/ring buffer. However, there is no way to prevent the events themselves from getting lost ([ref: Pitfall #5: Event Overload](#)). An attach point can be easily overwhelmed with events because eBPF lacks concurrency primitives and an eBPF probe cannot block the event producer. The kernel will simply stop calling the kprobes if it finds that it is overwhelmed.

Another issue with kprobes is that [it can be easily disarmed](#) if the attacker is allowed to execute its code even for a brief period.

## Response Capabilities

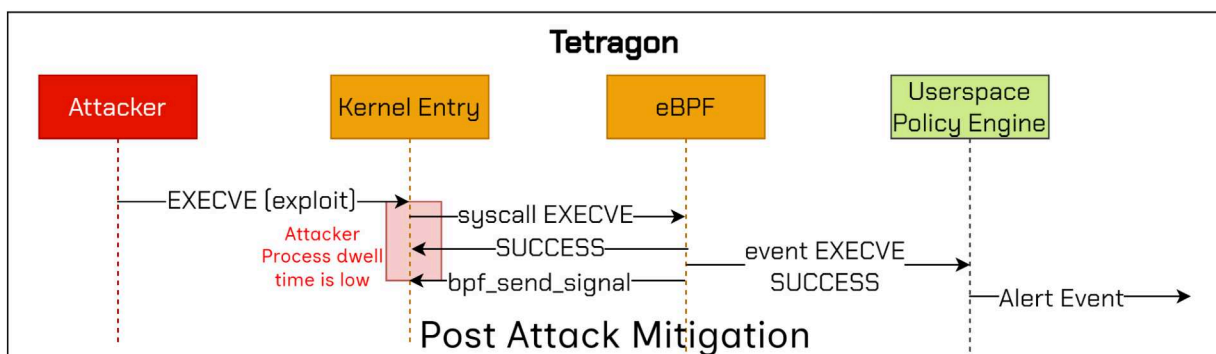
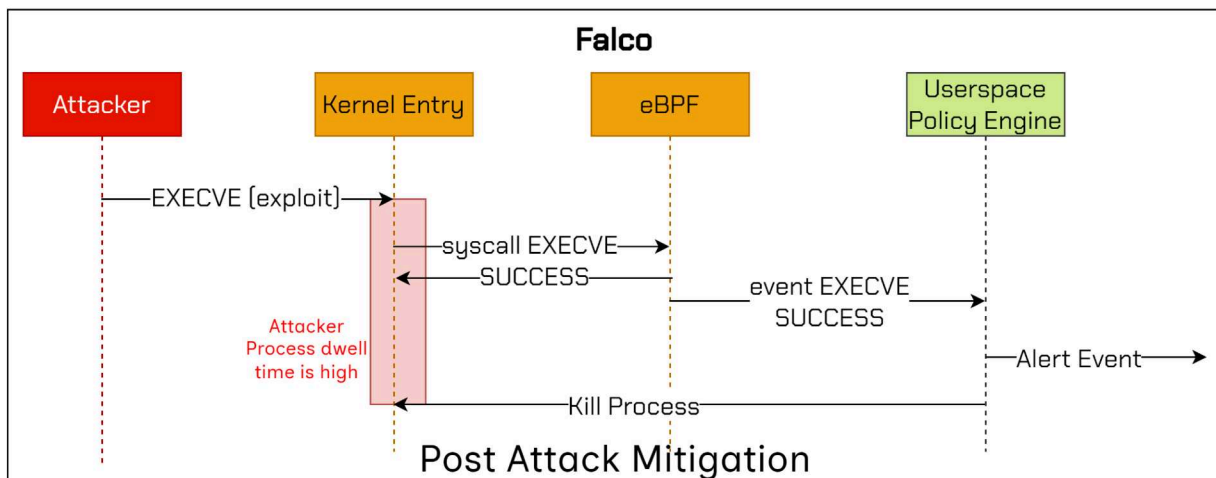
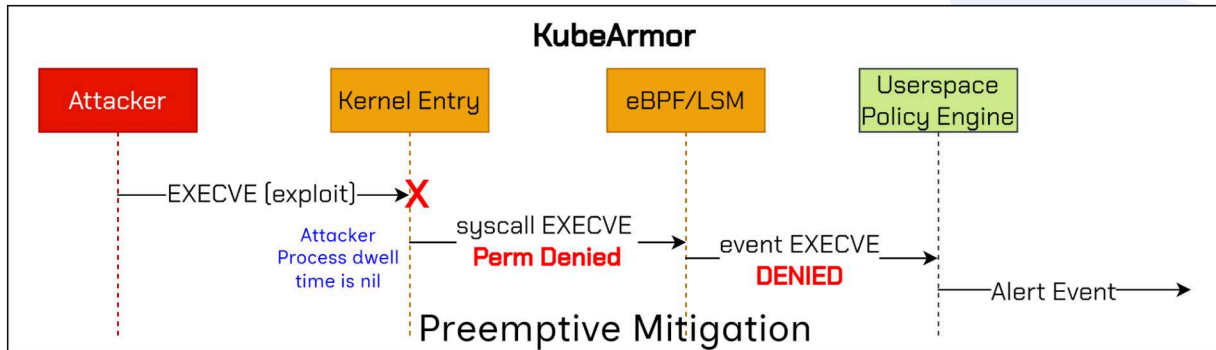
The events from the detection engine are sent to a policy decision process most likely implemented in the userspace that determines the action/response to be taken. Responses involve,

1. killing a target process
2. quarantining a node/pod
3. deleting a pod/node



In security parlance, this model is called post-attack mitigation response since the attacker is allowed to execute their code in the target environment and then a response is taken. In most cases, the response will take several milliseconds to several

minutes to get executed and thus might prove too late to take any effective action. For example, consider a ransomware attacker who is moving/deleting sensitive assets. Typically such an action could be completed in a few milliseconds and a response sent after that might prove ineffective in protecting those assets.



## TOCTOU issues

Time-of-check to time-of-use (TOCTOU) is a condition that occurs when a system's behavior is dependent on the timing between checking a condition and using the result. TOCTOU vulnerabilities can be exploited by attackers to gain unauthorized access to resources, modify data, or elevate privileges.

When handling certain system events such as `connect()`, `open()`, `openat()`, `creat()`, etc, the security engine retrieves some of the arguments by reading userspace buffers upon syscall exit. An attacker running a malicious program on a monitored system could use a variety of techniques to deterministically increase the duration of the syscall execution and modify the arguments in its own address space after the syscall has been invoked and before its execution is complete. The security engine will assume that the modified data is the input argument of the syscall which may lead to rule bypass.

Every detect and response security engine suffers from [this problem](#) including Falco-Talon, Tracee, and NeuVector. Tetragon uses `bpf_send_signal()` to send a kill signal in kernel space itself. However, this also is a case of post-attack mitigation since the attacker is allowed to execute its code even though just for a brief period. A technical guide explaining the pitfalls of this approach is mentioned [here](#). Quoting verbatim:

*“In attempting to mitigate container escapes, Tetragon tries using advanced Linux kernel features like eBPF and kprobes not to protect the very same kernel from getting exploited, but instead to stop an already successful exploit from using its gains.”*

## Overwhelming the events causing the events to drop

Detect and Response systems expect that the system events be successfully transmitted to the decision engine in userspace to handle responses. However, under high load conditions, it might not be feasible to transmit all the events to the userspace because of eBPF perf/ring buffer limitations. The kernel provides a circular ring buffer to propagate events from kernel space to userspace and under heavy load the events might get overwritten/lost. There is no way to know the criticality of events and thus the system might lose critical events. This may result in no response. Attackers might overwhelm the system events queue so that their malicious events never make it to userspace and thus no response is sent.

## Advantages of Detect and Response Model

1. **Multi-Vector Detection:** A detection system can consider multiple aspects of detection from different engines and then respond. For example, L7 API telemetry shows the use of the previously unused `log4j` endpoint and within the same time interval, a process invocation from the `/tmp/` folder can be considered a critical security event. The response could be to quarantine the corresponding pod.
2. **Multi-Dimensional Response:** A malicious system process invocation detected by a runtime security engine can trigger a response that deactivates external network access by changing the AWS VPC Security Group and killing the malicious process.



## Disadvantages of Detect and Response Model

1. **Post Attack Mitigation:** The response is sent once the attacker's code is executed in the target environment. An Attacker can stop the events from getting triggered by either changing the security knobs or by overwhelming the event loop and thus a response will never be initiated.
2. **Unable to protect certain actions:** Consider a k8s pod that has mounted a volume mount containing sensitive assets. An attacker would delete the said sensitive asset and the response would not be able to prevent it since by the time the response is handled, the operation would have been completed.
3. **Bigger impact on services:** The response might result in a service outage depending on the response action. For example, quarantining the node/pod might result in a service outage depending on how the application executing on the corresponding node/pod handles resource unavailability.

## Prevention Capabilities

Prevention or Inline enforcement requires that the system action such as process exec, file access, or network access is vetted/rejected *before the action is executed*. Linux primitives such as seccomp, and LSMs (Linux Security Modules such as AppArmor, SELinux, etc) provide a systematic way of such a vetting process and deny the execution or access of the resource.

There are user-space techniques (such as LD-PRELOAD) that can also achieve prevention capabilities, however, there are security issues surrounding userspace techniques and they can be easily circumvented even by a script kiddie attacker.

### eBPF and its role in preventive capabilities

It is usually assumed that eBPF provides enforcement capabilities as well. However, this is not true in all cases. The enforcement capabilities of eBPF depend on the leveraged hook points. For example, in the case of network packet handling, one can leverage traffic control (TC) hook points or XDP (express data path) hook points where one can redirect/drop a packet and thus handle network packet enforcement rules.

However, the same is not true for *kprobes*, *kretprobes*, and *tracepoints*. These hook points are primarily for observability and one cannot use them to change the system call behavior without major implications.

Note that *bpf\_override\_return()* is often quoted as a bpf-helper primitive available for enforcing controls at the kprobes hook point. This helper works by setting the PC (program counter) to an override function which is run in place of the original probed function. This means the probed function is not run at all. The replacement function just returns with the required value.

Tetragon [quotes](#) this as the primitive it uses. However, `bpf_override_return()` should not be used for security controls as per the [bpf-helper man page](#) itself since it has “security implications”. `bpf_override_return()` is typically used for error injection and is dependent on another kernel configuration called `CONFIG_FUNCTION_ERROR_INJECTION`. Moreover, the system call needs to be [explicitly listed as error-injectable](#) on the target distribution for it to work. Given all this, using `bpf_override_return()` is not worth pursuing in the production environments. [Based on my understanding](#), less than 20% of production systems have this enabled (most of these distributions are used on desktops, the distributions used in k8s and server deployments do not keep these enabled).

## Prevention Capabilities and Zero Trust

Zero Trust expects that policies can be specified in the least permissive mode. For example, the user should be able to specify what is allowed execution, and everything else should be denied (or audited).

A detect-and-response system cannot truly provide a Zero-Trust system because it does not have preventive capabilities.

## Sandboxing Capabilities

[Sandboxing](#) applications during their runtime execution is a practice often followed to ensure that unbounded access is not provided to untrusted applications. Unbounded access in terms of file system access, process invocation, network communications, and the use of advanced capabilities. For example, assume there is a Nginx web server running in a k8s deployment. Typically the nginx is the only process that would be executed within that pod and needs access to sensitive assets such as configuration files (`/etc/nginx`). Only nginx and all its worker/child nginx processes can perform network communication. A sandboxing rule would ensure that.

- Only nginx process is allowed to be spawned within that pod
- Only the nginx process is allowed access to sensitive file system paths such as `/etc/nginx` containing critical configuration files.
- Only nginx process is allowed to do network communication
- No processes are allowed to use advanced capabilities such as `CAP_SYSADMIN`, ...

Sandboxing allows the deployments to be put in the least permissive mode, making it resilient to zero-day attacks, and remote command injections/executions. Other use cases for sandboxing might require constraining unrestrained access to process executions. For example, LLM frameworks (such as [vllm](#)) utilize external plugins (use of [trust\\_remote\\_code flag with vllm](#)) to achieve certain objectives. The ability to trust

arbitrary code is a huge security concern and thus sandboxing of such framework would be highly desirable.

## Sandboxing techniques

Sandboxing cannot be achieved using a detect and response model since in that model, the untrusted code is allowed execution, and a response action is taken later. Sandboxing requires inline mitigation i.e., if an untrustworthy access/execution is attempted it has to be stopped/blocked/denied inline.

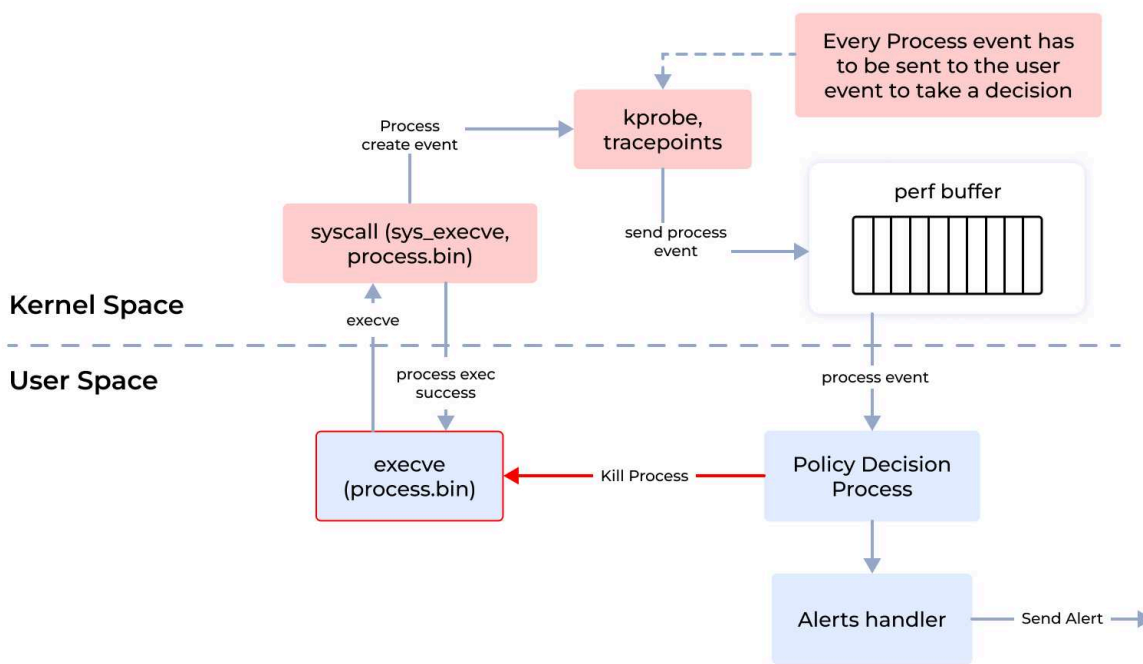
Gating all the action: Sandboxing requires that all the system actions are gated i.e., pass through certain checks, and if the checks fail the action is denied.

[Google gVisor](#) and [KubeArmor](#) provide sandboxing capabilities.

## Performance Impact

Every runtime security engine emits system events telemetry/logs. This can critically impact the user application running on the same cluster or nodes. Thus, it is important to ensure that the runtime performance of the security engine is kept under control.

An architectural issue with detect and response systems is that they have to emit all the systems events to the policy decision process in the cloud/node and thus cause heavy performance impact. Every kernel event shipped to userspace impacts the performance since a context switch is required to ship the event between the kernel to user space boundary and post that there are userspace functions that will further induce performance overhead.



In contrast, a runtime security system that supports profiling and whitelisting of the known system behavior and ensuring that the preventive security is handled in the kernel itself can reduce the overhead significantly. Only unknown or non-whitelisted events will be sent to the cloud, thus reducing the performance overhead significantly.

Another approach to alleviating the performance impact is to do in-kernel aggregation of events before sending it to the userspace reducing kernel context switches. Any aggregation has an impact on the fidelity of the data samples and thus should be carefully done. [AccuKnox has an IPR \(patent\)](#) in the context of handling optimal in-kernel aggregation of such events.

## Ease of Deployment on Hardened Distributions

Enforcing blocking rules requires special primitives to be used. Such primitives might require certain changes at the host or the container runtime interface layer. For example, Palo Alto TwistLock allows block-based policies by overriding the system default container runc binary with their runc-shim binary. Most hardened distributions (Talos Linux, Bottlerocket, GKE-COOS, etc) won't allow this. Runtime Security engines that allow inline/preventive mitigations should be deployable using standard K8s constructs (helm, kubectl apply, k8s operator) without changing anything at the host/node level.

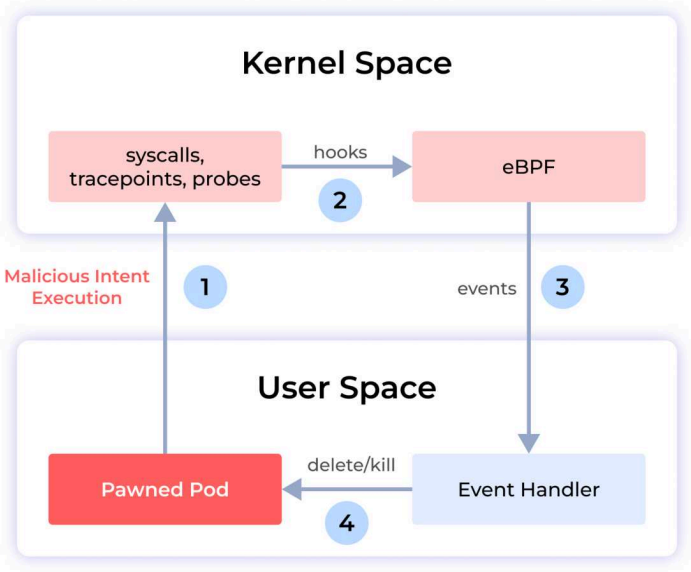
## Ease of Runtime Policy Enforcement

Representing policies as native Kubernetes resources allows one to enable/disable rules at will and manage the lifecycle just as any other k8s resource. It is important that when a rule is enabled/disabled it should not require the containers/pods to be restarted.

## Policies Adhering to Zero Trust Principles

The user should be allowed to specify a policy allowing specific action and deny/audit everything else.

# Runtime Security with Detect and Response



The primary model with such systems is that all the system events are shipped to a userspace PDP (Policy Decision Process) or event handler. The user-specified policies/rules are matched and optionally a response is sent that can:

- a. Kill the target process
- b. Quarantine the target pod/node
- c. Delete the target pod/node
- d. Change the CSP VPC Security group setting

## Issue 1: Killing a process is not an effective remediation strategy

Handling a remedial action takes from a few seconds to a few minutes. The fundamental notion of a detect and respond system is to analyze execution events and then kill the malicious process. However, once the malicious code is allowed to execute in the target environment, then the attacker most likely would turn off the security knobs or overwhelm the event engine such that it reaches the event threshold and starts dropping events. Consider another scenario, where a ransomware attacker is moving or deleting the sensitive files. This typically takes time in a few milliseconds. A detect and response model will never be able to thwart such attacks since by the time the response action is taken the damage is already done.

## Issue 2: The Response depends on the successful execution of a chain of actions

The Chain of action:

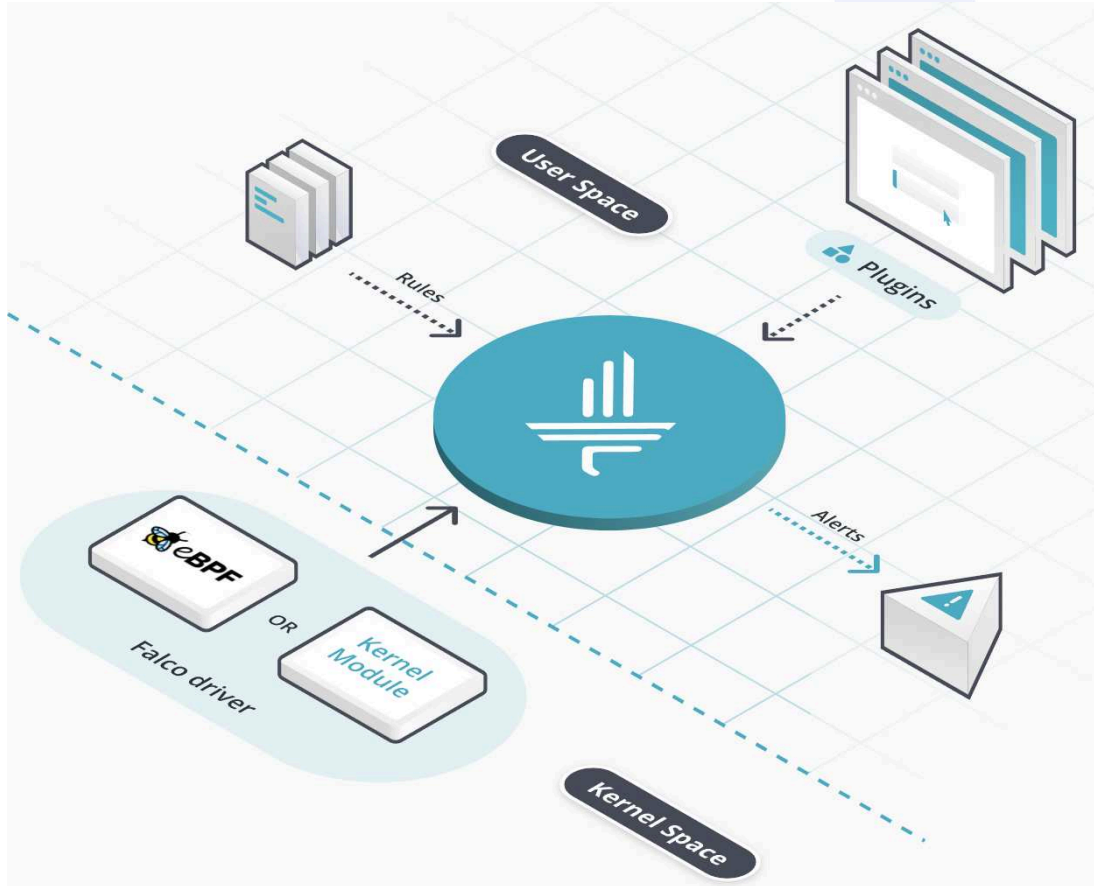
### Event Detection in kernel

- Send the event to userspace
  - Send the event to Policy Decision Process (PDP), in Cloud/Node
    - PDP sends a response action
      - The response needs to be shipped to the target node

The attacker can impede or prevent a response from being sent by compromising multiple aspects of this chain.

## Falco Analysis

[Falco](#) is a CNCF graduated runtime security tool designed to detect and alert on abnormal behavior and potential security threats in real-time. Falco is a detection engine, and [Falco-Talon](#) provides response mechanisms on top of Falco and other events. Falco provides flexible ways to specify rules based on which the system events are filtered.



## Policy Enforcement

Falco does not provide any preventive or inline policy enforcement. For example, one can add rules to detect if particular processes have been executed but cannot add rules to deny the execution of those processes. Falco-Talon provides an asynchronous response engine that operates on top of Falco and qualifies as a detect-and-response system.

## Tetragon Analysis

[Tetragon](#) provides real-time, eBPF-based Security Observability and Runtime Enforcement. One of the primary differences as compared to Falco when it comes to filtering the system events is that most of the filtering happens in kernel and thus saves costly kernel to userspace context switches.

Tetragon can hook into any function in the Linux kernel and filter on its arguments, return value, and associated metadata that Tetragon collects about processes (e.g., executable names), files, and other properties. By writing tracing policies users can solve various security and observability use cases.

## Policy Enforcement Mechanics

Unlike Falco, Tetragon provides a policy enforcement mechanism as well where users can specify enforcement policy actions such as:

- [Send Signal](#): Signal action (such as Sigkill) sends a specified signal to the current process.
- [Override](#): Override action allows to modify the return value of the call. While Sigkill will terminate the entire process responsible for making the call, Override will run in place of the original kprobed function and return the value specified in the argError field.

While Tetragon allows local in-kernel policy enforcement, its policy enforcement mechanics suffer from certain problems.

- Send Signal: Tetragon leverages [bpf\\_send\\_signal\(\) bpf-helper function](#) to send a kill signal to the current process. While the signal is sent from the kernel space itself, it still qualifies as a post-attack mitigation primitive since the malicious process would still be allowed to execute. [This article from Grsecurity](#) explains the risks associated with post-attack mitigation in general and tetragon in particular.
- Override return: Tetragon leverages [bpf\\_override\\_return\(\) bpf-helper function](#) to return arbitrary user-specified value to the calling system call. Typically this policy action won't work on most production systems because
  - It depends on multiple kernel configurations that are typically not enabled on production systems.
  - There is also a dependency on the target system call to be tagged under ALLOW\_ERROR\_INJECTION list for the helper function to work.
  - [bpf\\_override\\_return\(\) function man page](#) mentions that the bpf-helper has security implications, and thus is subject to restrictions.

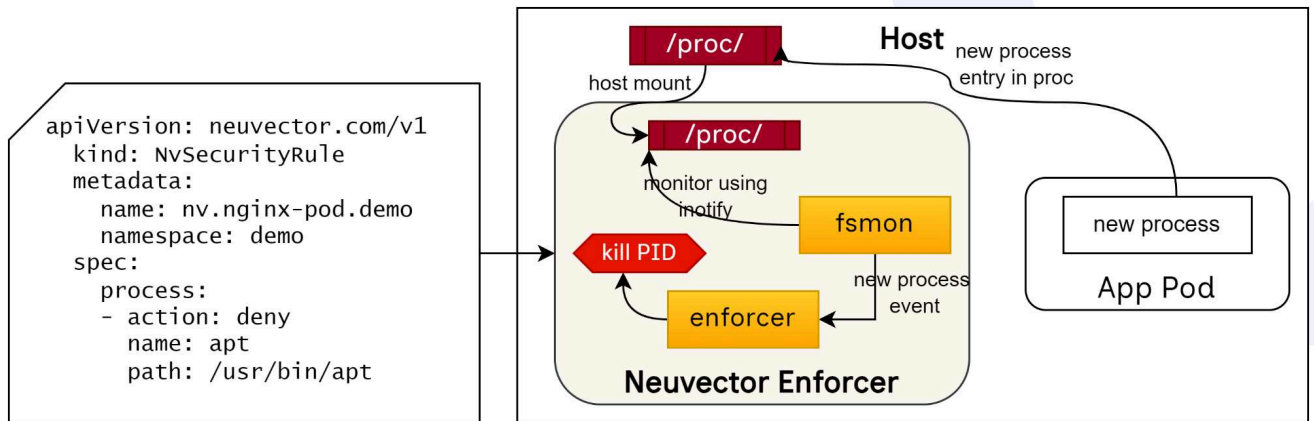
## NeuVector Analysis

[Neuvector](#) provides process, file, and network-based container security rules to be specified. Neuvector also allows local remediation action to be taken. For example, users can specify deny actions for specific processes.

## Policy Enforcement Mechanics

Please note that there is no online documentation that explains the internal architecture of Neuvector container runtime security. However, the code was open source after its acquisition by SUSE, and the following text is based on [the analysis of the code repo](#) itself.





I created the above diagram by understanding the code flow to explain the internal architecture of how policy enforcement works. Based on the above diagram:

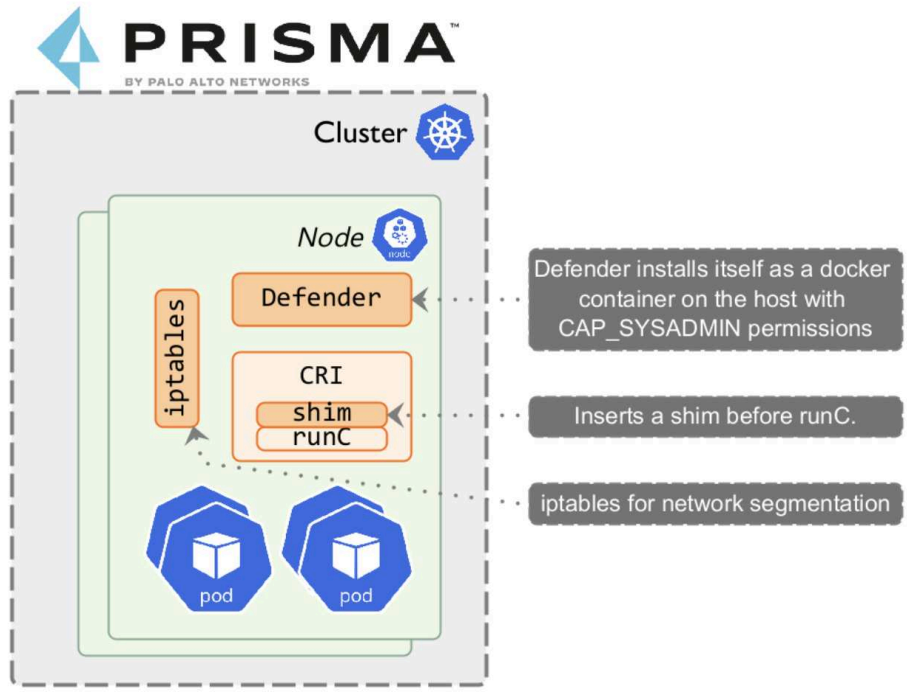
- Neuvector mounts the host procs within the container itself and monitors it using fsmon/inotify.
- Any execution of a new process event can now be tracked since it would result in a new folder created in the procs with the corresponding PID.
- Neuvector enforcer then sends the kill signal to the corresponding PID if the process is blacklisted.

Thus, Neuvector also deploys post-attack mitigation techniques (such as killing the process) from an enforcement perspective.

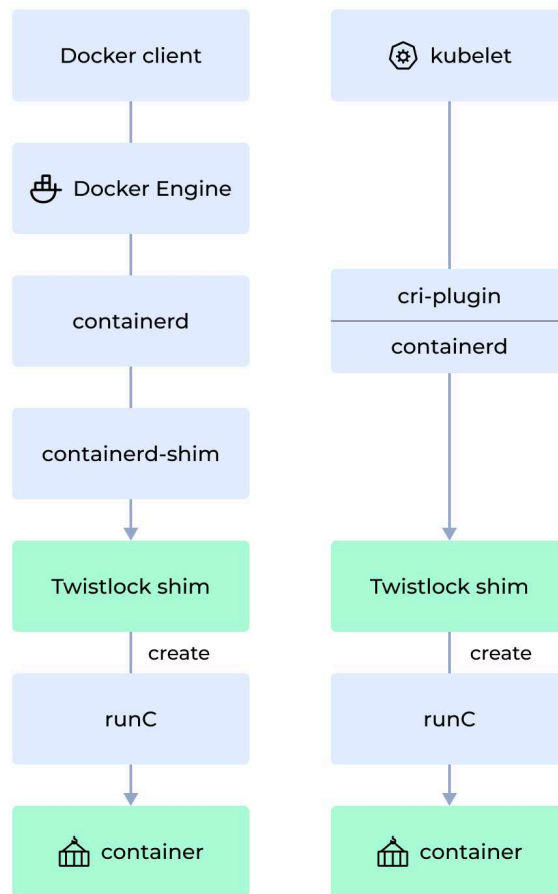
Notably, Neuvector does not employ eBPF-based techniques for detection purposes and thus the overall performance impact would be significantly high since it would be difficult for any userspace-based techniques to scale for such runtime security needs.

## Palo Alto Prisma/TwistLock Analysis

Prisma Cloud provides container runtime security as one of the features and uses its [defender architecture](#) to fulfill this feature. Notably, Prisma Defender also does not use eBPF-based architecture. Instead, it relies on container runtime-based integration.



Prisma Defender is responsible for enforcing vulnerability and compliance blocking rules. When a blocking rule is created, Defender moves the original runC binary to a new path and [inserts a Prisma Cloud runC shim binary](#) in its place.



## Policy Enforcement Mechanics

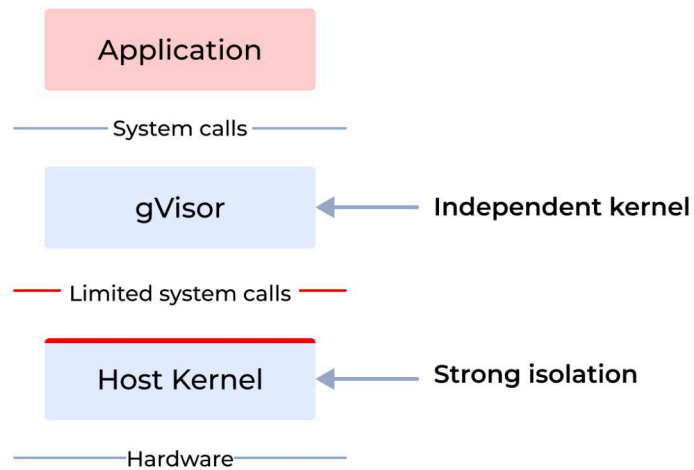
Prisma replaces the system runC with its runC binary to apply block-based rules. The Prisma runC thus becomes the topmost-parent process for all the container-based execution and thus can control the execution. When enforcing block-based enforcement rules, there are multiple issues:

1. Difficult to operate on a hardened distribution such as Bottlerocket, Talos, COS, etc, since it won't allow insertion of runC shim. Thus manual host-specific changes have to be made.
2. Applying block-based rules would require node restarts.
3. The performance impact would be significantly higher since all the decisions are handled in the userspace.

## gVisor Analysis

gVisor provides a strong layer of isolation between running applications and the host operating system. It is an application kernel that implements a Linux-like interface.

gVisor includes an Open Container Initiative (OCI) runtime called runc that enables it to work with existing container tooling. The runc runtime integrates with Docker and Kubernetes, making it simple to run sandboxed containers.



gVisor intercepts application system calls and acts as the guest kernel, without the need for translation through virtualized hardware. gVisor may be thought of as either a merged guest kernel and VMM, or as seccomp on steroids. This architecture allows it to provide a flexible resource footprint (i.e. one based on threads and memory mappings, not fixed guest physical resources) while also lowering the fixed costs of virtualization. However, this comes at the price of reduced application compatibility and higher per-system call overhead.

## Policy Enforcement Mechanics

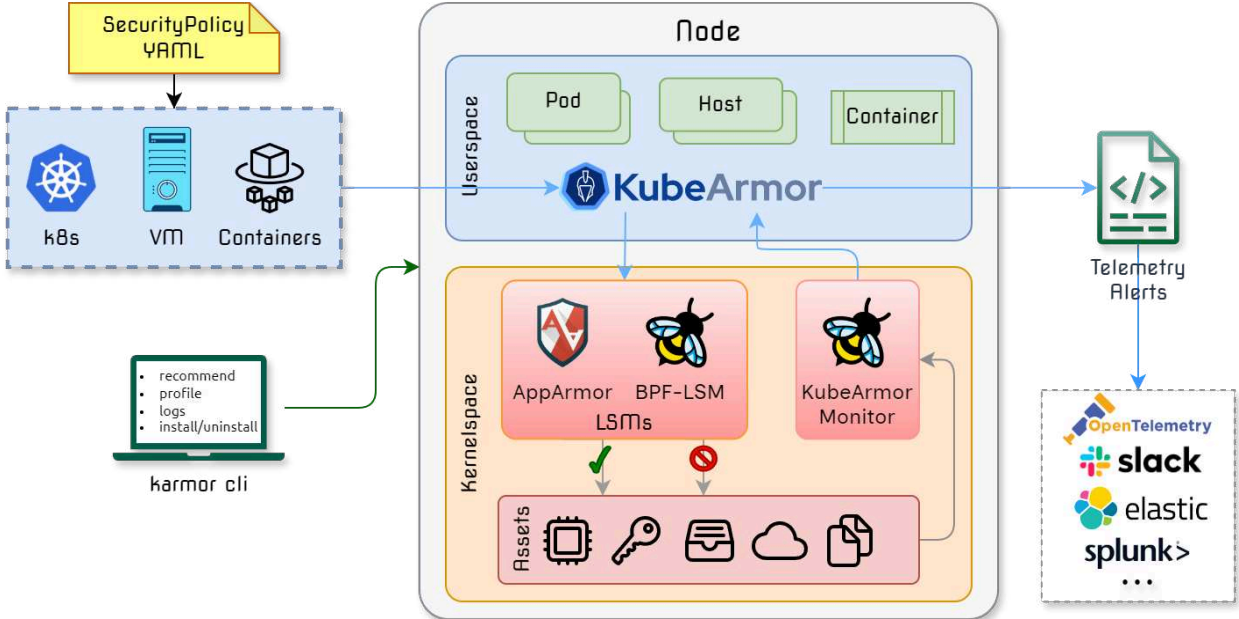
Compared to any other security engine, gVisor takes a very different approach. It introduces an intermediate system call layer through which all the calls are made and the policy logic is implemented at this intermediate layer. gVisor provides a clean sandboxing environment and extremely powerful and flexible policy enforcement options as compared to any other policy engine. However, the other issues are:

- Intrusive deployment process: Replacing runC is difficult, especially on hardened distributions such as Bottlerocket, Talos, COS, etc.
- Performance impact: Since every call has to be directed through the sentry wall introduced by gVisor

# KubeArmor Analysis

KubeArmor is a cloud-native runtime security enforcement system that restricts the behavior (such as process execution, file access, and networking operations) of pods, containers, and nodes (VMs) at the system level.

KubeArmor leverages Linux security modules (LSMs) such as AppArmor, SELinux, or BPF-LSM to enforce the user-specified policies. KubeArmor generates rich alerts/telemetry events with container/pod/namespace identities by leveraging eBPF.



## Policy Enforcement Mechanics

KubeArmor leverages eBPF for detection and audit rules. For block-based rules, it leverages (in the order of priority):

- LSM-BPF
- AppArmor
- SELinux (only for host-based rules).

[KubeArmor is the first engine to leverage LSM-BPF](#) to enforce block-based rules consistently for process, file, and network. LSM-BPF provides an extremely flexible way for converting user-specified rules (managed by k8s native resources) into eBPF bytecode that is then injected at LSM hooks. LSM hooks ensure that the enforcement does not suffer from post-attack mitigation, TOCTOU, or Semantic poisoning issues. Since all the decision-making, including enforcement happens in kernel space, the impact on the performance is limited.

KubeArmor cannot be operated on environments that do not support LSM-BPF or AppArmor. However, there are no distributions where either of them is not enabled.

# Case Studies

## File Integrity Monitoring/Protection using different tools

### FIM using Falco

Falco provides open-source [File Integrity Monitoring rules](#) based on eBPF. However, it does not support enforcement of block rules. For example, there is no way to deny changes in system folders using these rules. Using Talon, however, one can add response actions in post-attack mitigation (detect and response) style if there are changes detected in the system folders.

### FIM using Tetragon

Tetragon provides both File Integrity Monitoring and [enforcement/protection](#). In enforcement/protection, it sends a kill signal to any process trying to make changes to the system folders. However, we didn't find any rule to prevent deletion/unlinking of assets from the same system folders. Hence, we added a security rule using kprobe security\_path\_unlink to ensure that the unlinks are prevented. However, Tetragon could not prevent the deletion/unlink of the assets because of post-attack mitigation issues i.e., the target process was killed after the assets were deleted. The detailed note with the changes as well as behaviour is captured [here](#).

### FIM using KubeArmor

KubeArmor leverages LSMs (LSM-BPF) to prevent changes to sensitive folders. The sample policy can be found [here](#). KubeArmor does not suffer from TOCTOU issues or any post-attack mitigation issues.

# Summary

There are three categories of Runtime Security Engines:

1. Engines that provide detection capabilities and handle threat analysis in the cloud/control plane providing remediation/response capabilities. (Falco-Talon + Sysdig, Tracee, ORCA, Wiz, ...)
2. Engines that provide detection capabilities and provide localized response capabilities (NeuVector, Tetragon)
3. Engines that provide detection capabilities, response capabilities from the cloud/control plane, and provide inline mitigation capabilities (KubeArmor, [security-profile-operator](#)).

Engine \ Feature	Falco + Talon	Tracee	Tetragon	KubeArmor	NeuVector	gVisor	Prisma (Enterprise only)
<b>Detect &amp; Response</b>	Yes	No <sup>(ent)</sup>	No	No <sup>(ent)</sup>	No <sup>(ent)</sup>	No	Yes
<b>Preventive Capabilities</b>	No	No	Limited	Yes	No	Yes	Yes
<b>Hardened distros support</b>	Yes	Yes	Limited	Yes	Yes	No	Limited
<b>Performance Impact</b>	High	High	Low	Low	High	High	High
<b>Zero Trust Policy</b>	No	No	Yes	Yes	Yes	Yes	Yes
<b>Sandboxing</b>	No	No	No	Yes	No	Yes	Yes
<b>Deployment complexity</b>	Low	Low	Low	Low	Low	High(requires runC change)	High (requires runC change)
<b>Implementation Language</b>	C++	go	go	go	go	go	Not-Known

<sup>(ent)</sup>: Available in Enterprise version

Performance Impact is high for every Detect and Response engine that does not do local/in-kernel aggregation.

Containerized environments are becoming the backbone of modern application deployment, bringing agility and security challenges. This guide aims to provide a comparison of the top runtime container security tools, highlighting their capabilities, key architectural tenets, and trade-offs of those architectural tenets. By

aligning the right tools with your security strategy and operational needs, you can strengthen your containerized workloads against evolving threats. Effective runtime security is a journey of continuous improvement, not a one-time solution.

If you have any feedback or suggestions, feel free to share them with Rahul at [rahul@accuknox.com](mailto:rahul@accuknox.com).



Featured by



ORACLE

Gartner



IBM

LFEDGE



Extra 30 Days Free Trial



\*No strings attached, limited period offer!



Scan for Demo

## About AccuKnox

AccuKnox delivers agentless zero trust security for public and private cloud platforms to secure modern and traditional workloads.



in [linkedin.com/accuknox](https://www.linkedin.com/company/accuknox)

X @AccuKnox

