



# How to Promote the use of Best Practices and Automate Security Policies Using Tools Like OPA and Kubernetes Native Declaratives

**Gary Duan**  
Co-Founder & CTO, NeuVector

# Agenda

## Build Security Into Your Pipeline

### Security as Code, CRD and OPA

- Demo 1: OPA and Vulnerability Scan

### Kubernetes Admission Control

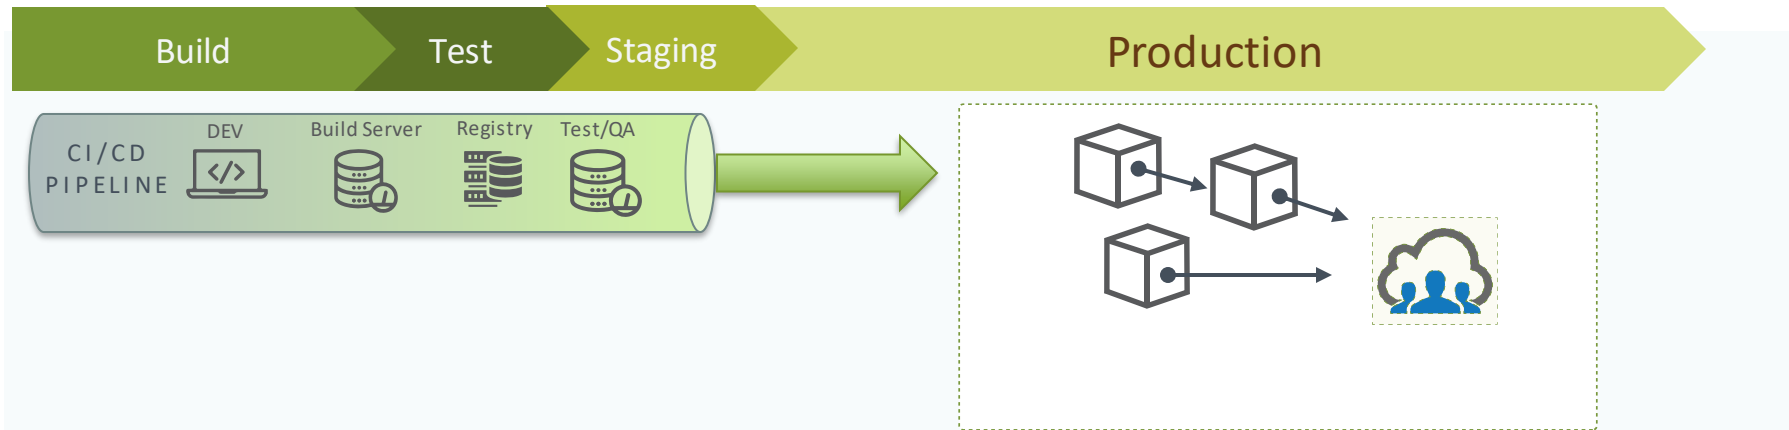
- Demo 2: OPA to Audit Admission Control

### Kubernetes Run-Time Security

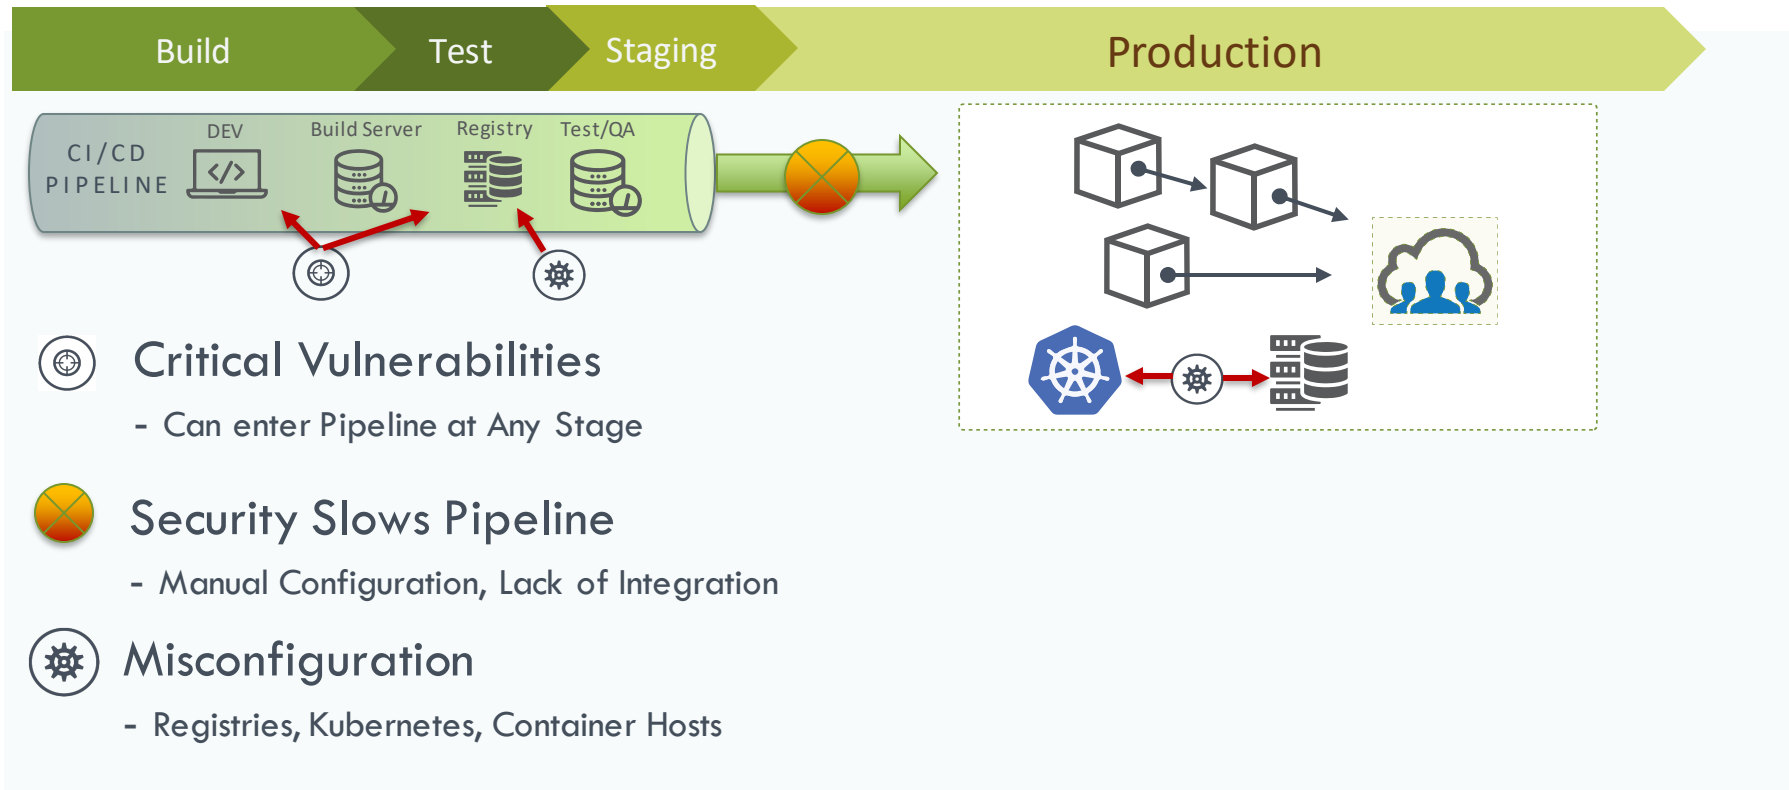
- Demo 3: OPA to Audit Run-Time Policies
- Security Posture Management and Virtual Patching

### Q&A

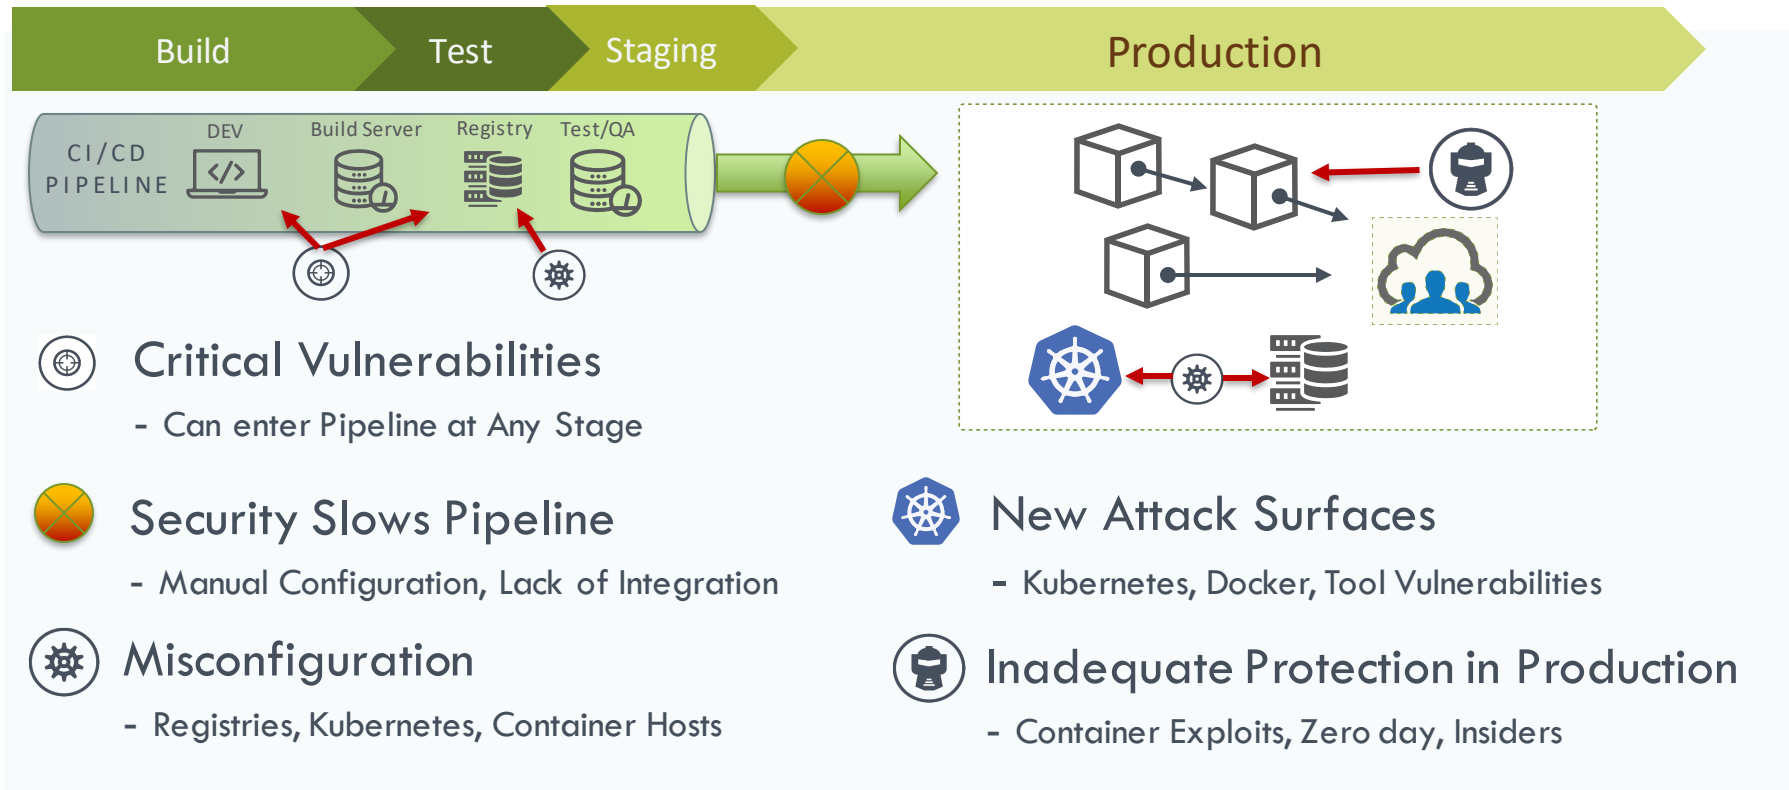
# Cloud-Native Pipeline



# Top Security Issues for Cloud-Native Deployments



# Top Security Issues for Cloud-Native Deployments



# \>whatis Security as Code

*Everything in Kubernetes is a **declarative configuration object** that represents the **desired state** of the system.*

*- Kubernetes Up and Running - Burns / Beda / Hightower - O'Reilly Media*

Security **stored in declarative configurations** and **delivered** as part of a Continuous Delivery pipeline can be referred to as Security-as-code.



## Challenges for Security as Code in Kubernetes

- What can be / should be a declarative configuration?
- Rapid-deployment in a small DevOps team
- MicroService Communication & Network Visibility

# Kubernetes Custom Resource Definition (CRD)

A *resource* is an endpoint in Kubernetes' API that stores an API object

A *custom resource* is an extension of the Kubernetes API for custom API objects

RBAC-enabled: *namespace* and *cluster* level resources

# Introduction to Open Policy Agent (OPA)

## What is it?

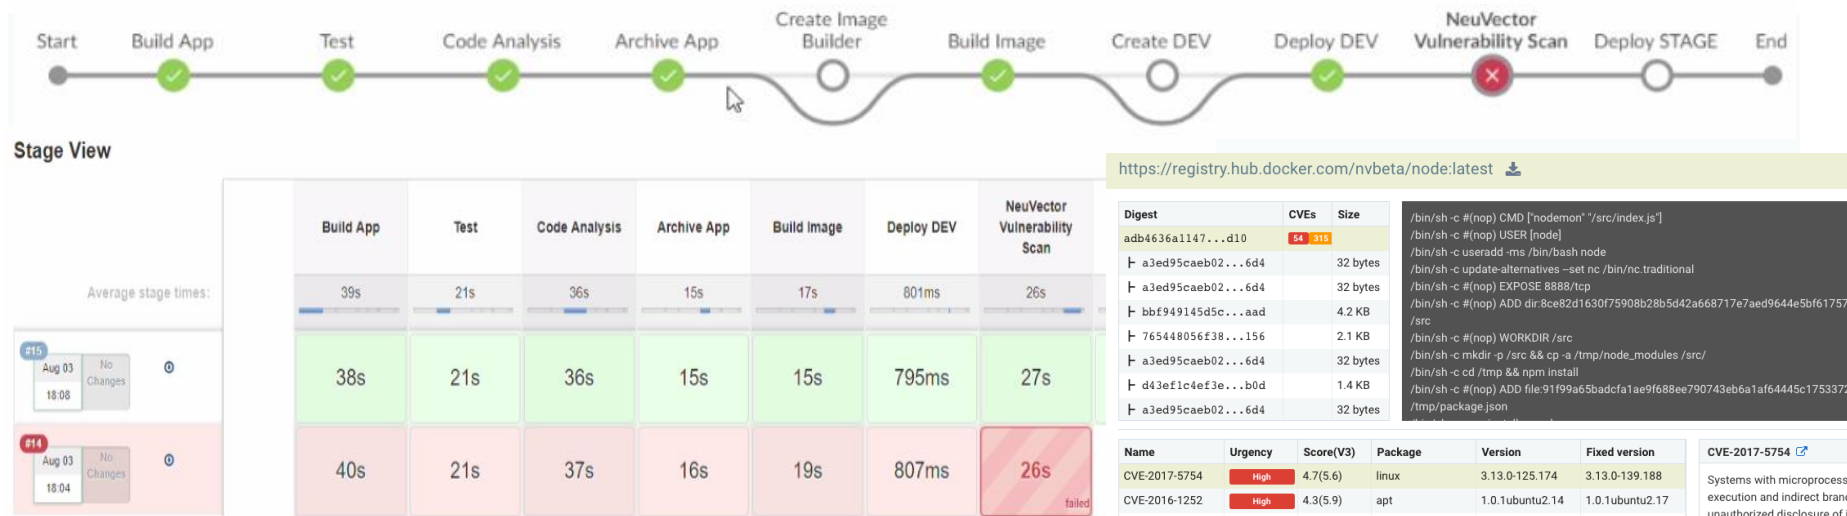
- A general-purpose policy engine
- Offers declarative language to describe the policy
- Decouple the decision-making and the policy enforcement
- Key component of security-as-code workflow

## Use Cases

- Kubernetes Admission Control Gatekeeper
- Service Mesh Authorization Policy
- Policy Auditing



# CI/CD Pipeline Integration – Vulnerability Scanning



- ✓ Automated Build Phase Scanning
- ✓ Scanning Public Cloud Registries, Private Registries, 3<sup>rd</sup> Party Registries
- ✓ Layered Image Vulnerability Analysis

# OPA and Vulnerability Scan

## *Policy*

- Image must not have more than 2 vulnerabilities

```
≡ audit.rego
1  package psp
2
3  default allow = false
4
5  allow = true {
6    count(input.report.vulnerabilities) < 3
7  }
8
9
```



makeameme.org

# OPA and Admission Control

## A) Pod Security Policy (PSP)

- Cluster-level resource, authorized to a service account
- Privileged, run-as-root
- Volume, host path/namespace, read-only root fs
- Linux caps, seccomp, SELinux/AppArmor

## B) NeuVector Admission Control

- Monitor and Enforce mode
- Leverage container image vulnerability scan results
- CVE severity, fix status, date and software component
- Leverage container image compliance checks



# Audit Pod Security Profile

## Policy

- No pod can run in privileged mode except for those in the "admin" namespace

```
1  apiVersion: policy/v1beta1
2  kind: PodSecurityPolicy
3  metadata:
4    name: psp.restricted
5  spec:
6    privileged: false
7    allowPrivilegeEscalation: false
8    runAsUser:
9      rule: 'MustRunAsNonRoot'
```

```
≡ psp.rego x
psp > ≡ psp.rego
1  package psp
2
3  import data.roles.bindings as bindings
4  import data.roles.permissions as role
5
6  default allow = false
7
8  allow = true {
9    some spec
10   profile[spec]
11   spec.privileged = false
12 }
13
14 allow = true {
15   some spec
16   profile[spec]
17   spec.privileged != false
18   input.ServiceAccount.namespace = "admin"
19 }
20
21 profile[spec] {
22   some i
23   spec := data.spec
24
25   bindings.subjects[i].namespace = input.ServiceAccount.namespace
26   bindings.subjects[i].name = input.ServiceAccount.name
27
28   role.metadata.name = bindings.roleRef.name
29
30   data.metadata.name = role.rules[_].resourceNames[_]
31 }
32
```

# Audit Admission Control Rules

## *Policy*

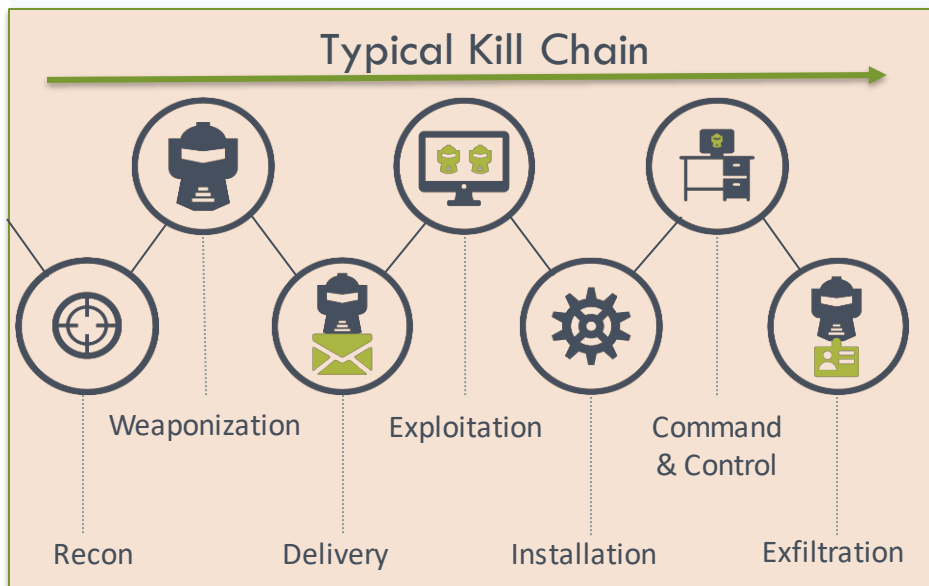
- Number of high severity vulnerabilities that have been reported for more than 30 days must be less than 5
- No restriction in the "staging" namespace

≡ audit.rego ×

admission > ≡ audit.rego

```
1  package admission
2
3  default allow = false
4
5  allow = true {
6    some i
7    r := input.rules[i]
8    r.disable = false
9    r.rule_type = "deny"
10
11    some c1
12    r.criteria[c1].name = "namespace"
13    r.criteria[c1].op = "notContainsAny"
14    r.criteria[c1].value = "staging"
15
16    some c2
17    r.criteria[c2].name = "cveHighWithFixCount"
18    r.criteria[c2].op = ">="
19    r.criteria[c2].value = "5"
20
21    some c2s
22    r.criteria[c2].sub_criteria[c2s].name = "publishDays"
23    r.criteria[c2].sub_criteria[c2s].op = ">="
24    r.criteria[c2].sub_criteria[c2s].value = "30"
25  }
```

# Run-Time Protection – Critical Attack Prevention in Production



	Network	Container	Host
Cryptojacking Bitcoin Mining	✓		✓
Port Scan	✓		
EQUIFAX DATA BREACH	✓	✓	
DDOS ATTACK	✓		
SQL	✓		
DIRTY COW		✓	✓
mongoDB	✓	✓	

# Using OPA to Audit Run-Time Policies

## 1. Kubernetes Network Policies

- Namespace scoped, pods selected by labels
- Pod-to-pod policy. Limited flexibility of cluster ingress/egress.
- Defined by protocols and ports

## 2. NeuVector CRD Policies

- Monitor and Protect mode
- Pods selected by service, label, namespace ...
- Pod-to-pod, pod-to-host and cluster ingress/egress
- RBAC-aware, global policy
- CIDR, FQDN and L7 protocols
- Network, process and file system activities



# Audit Kubernetes Network Security

## *Policy*

- Applications in the demo namespace can only listen on at most 1 port
- Applications in the demo namespace cannot make outbound connections

```
≡ audit.rego
1  package policy
2
3  default allow = false
4
5  allow = true {
6    ... data.kind = "NetworkPolicy"
7    ... data.metadata.namespace = "demo"
8    ... ingress_port
9    ... egress_exist
10 }
11
12 ingress_port = true {
13   ... count(data.spec.ingress[0].ports) <= 1
14 }
15
16 ingress_port = true {
17   ... not data.spec.ingress[0].ports
18 }
19
20 egress_exist = true {
21   ... e := data.spec.egress[_]
22   ... e.to[0].podSelector
23 }
24
25 egress_exist = true {
26   ... e := data.spec.egress[_]
27   ... e.to[0].namespaceSelector
28 }
```



# Run-time Policy CRD Explained

## ✓ Define Application Behaviors in Kubernetes-native yaml

- Network Connections and Protocols
- Ingress/egress controls
- Process & File System Activity

## ✓ Enforce Global Security Rules

- Cluster-level policy on Ingress/egress control

## ✓ RBAC Integrated

- Kubernetes enforcement of CRD creation permissions
- Support Multi-team use cases

## ✓ Eases migration from staging to production

## ✓ Support Open Policy Agent (OPA), other integrations

```
- apiVersion: neuvector.com/v1
  kind: NvClusterSecurityRule
  metadata:
    name: containers
  spec:
    process:
      - action: deny
        name: sshd
        path: ""
        target:
          Selector:
            criteria:
              - key: container
                op: =
                value: '*'
              name: containers
            policymode: null
          version: v1
```

```
- apiVersion: neuvector.com/v1
  kind: NvClusterSecurityRule
  metadata:
    name: Oracle-0
  spec:
    ingress:
      - Selector:
          criteria:
            - key: service
              op: =
              value: myapp-pod.default
            - key: domain
              op: =
              value: default
            name: nv.myapp-pod.default
          action: allow
        applications:
          - Oracle
        ports: any
        target:
          Selector:
            criteria:
              - key: address
                op: =
                value: oracledb.acme.com
              name: OracleDB
```

# Audit NeuVector run-time CRD

## Policy

- Only MyApp application can connect to the Oracle DB outside of the cluster
- No containers can run sshd

```
≡ audit.rego
1  package crd
2
3  default allow = false
4
5  allow = true {
6    # target of the network policy
7    some r1
8    target := data.items[r1].spec.target
9    target.selector.criteria[0].key = "address"
10   target.selector.criteria[0].op = "="
11   target.selector.criteria[0].value = "oracledb.acme.com"
12
13   # ingress L7 application rule to the target
14   some j
15   ingress := data.items[r1].spec.ingress[j]
16   ingress.applications[_] = "Oracle"
17   ingress.selector.criteria[0].key = "service"
18   ingress.selector.criteria[0].op = "="
19   startswith(ingress.selector.criteria[0].value, "myapp-pod")
20
21   # only 'allow' rules are needed, network rules are implicit deny
22   ingress.action = "allow"
23
24   # cluster level process rule that prevents ssh daemon
25   # from running in the containers
26   some r2
27   data.items[r2].kind = "NvClusterSecurityRule"
28
29   # target of the process rule
30   t := data.items[r2].spec.target
31   t.selector.criteria[0].key = "container"
32   t.selector.criteria[0].op = "="
33   t.selector.criteria[0].value = "*"
34
35   some p
36   data.items[r2].spec.process[p].action = "deny"
37   data.items[r2].spec.process[p].name = "sshd"
38 }
```

# Run-Time Policy Migration using CRD



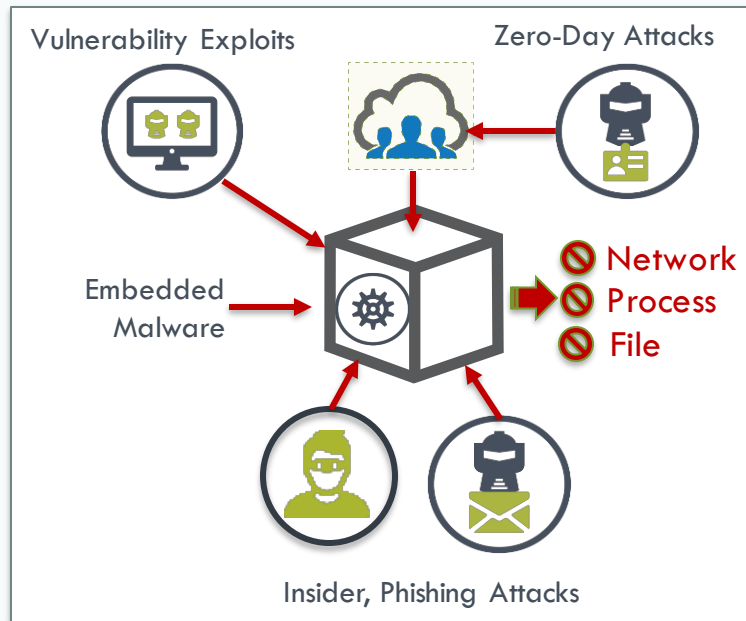
# Virtual Patching Protects Production Workloads & Hosts

## ✓ Detects & Prevents

- Vulnerability Exploits
- Zero-day Attacks
- Embedded Malware
- Insider, Phishing Attacks

## ✓ Learns baselines & Blocks

- Unauthorized Network Connections
- Unauthorized Processes
- Unauthorized File Access



# Summary

- Build Security into your pipeline from the beginning
- Automate Vulnerability Scanning
- Use admission control to gate the deployment
- Protect applications at run-time
- Employ Kubernetes-native constructs, such as CRD
- Using OPA to audit and enforce the security policies



# Questions?

*Thank You!*

*Contact NeuVector*  
*[info@neuvector.com](mailto:info@neuvector.com)*  
*<https://neuvector.com>*