



Beesting
Can't touch this

Who am I?

Software engineer turned Cloud Enthusiast ☁️

Kubernetes wizard 🪄

Linux Nerd 🐧



Let me take you back to
August 2024

Context-aware security incident response with Dynatrace Automations and Tetragon

Published May 3, 2024 | Updated December 12, 2024 | 11 min read



Mario Kahlhofer



Simon Ammer



Markus Gierlinger

Application security

Engineering

In this blog post

1. Better, faster application protection and security investigation
2. Step 1: Automating the placement of honeytokens to create strong indicators of compromise
3. Step 2: Alerting with automated context enrichment
4. Step 3: Auto-remediate with network policies and GitOps
5. Bonus step: Deploying the security policy into the live cluster
6. Workflows for security incident response on the Dynatrace platform

For the most severe threat scenarios, you want multiple layers of automated defenses, and not have to rely on humans to analyze the traces of an attack weeks after your system got compromised. Many security teams use runbooks to glue together tools, processes, events, and actions for security incident response. A runbook lays out the step-by-step instructions to follow when a security incident happens, when an emerging threat surfaces, or when your security tool reports suspicious behavior.

But runbooks that stitch together glamorous security tooling are merely decorations without automated workflows for incident detection and response.

The security community agrees on many high-level best practices in such situations, but we need a single platform solution to orchestrate application security, observability, and DevOps practices. Because every situation is a little unique, Dynatrace makes it easy to create custom runbooks using Dynatrace Automations, fine-tuned to your individual business risks.

In this blog post, we'll demonstrate how to use [Dynatrace Automations](#) to build a runbook that combats sophisticated security incidents with honeytokens and eBPF-based detection. We show an end-to-end solution, starting with deploying policies in a Kubernetes cluster and ending in a pull request assigned to the responsible team, all without manual intervention.

To demonstrate the integration of external security tools into the Dynatrace platform, we use [Tetragon](#) for eBPF-based security monitoring. Using [Kyverno](#), we can automatically kick attackers out of our cluster with network policies and harden our configuration with a GitOps workflow to prevent the same incident from happening again.

Share blog post



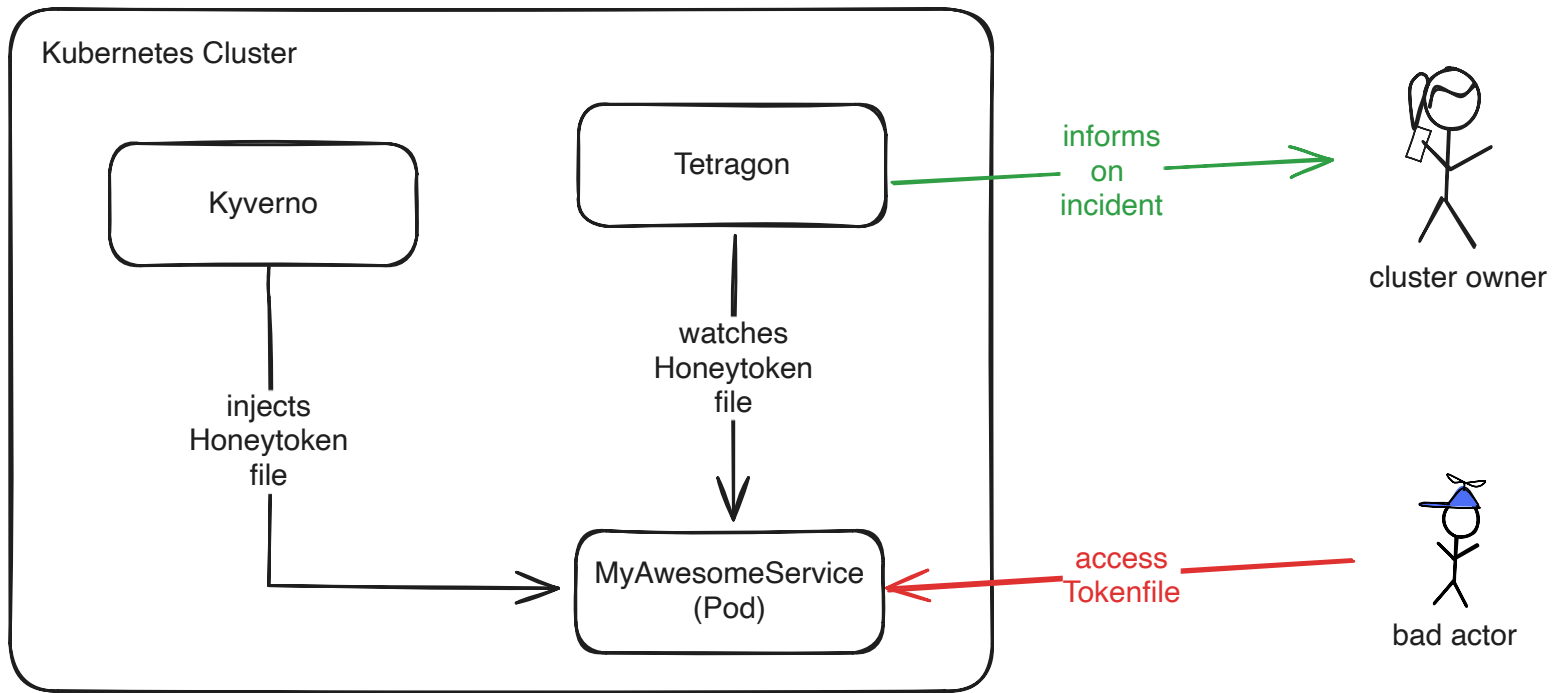
Stay Updated



Enter your email

- ☒ All updates
- ☐ Blog posts
- ☐ Product news

Subscribe now



**What even is a
Honeytoken?**



Digital bait



This is us

commerce:SEO

We are currently updating our CMS database. You will soon find further information about **commerce:SEO CMS** on this page. We will provide you an overview with all details about:

- latest changes and bug fixes
- updated or new plugins
- templates and themes
- developers news

webEdition

We are currently updating our CMS database. You will soon find further information about **webEdition CMS** on this page. We will provide you an overview with all details about:

- latest changes and bug fixes
- updated or new plugins
- templates and themes
- developers news

Cuppa CMS

We are currently updating our CMS database. You will soon find further information about **Cuppa CMS** on this page. We will provide you an overview with all details about:

- latest changes and bug fixes
- updated or new plugins
- templates and themes
- developers news

H.H.G. multistore

DO YOU WANT TO ADVERTISE ON THIS SITE?

[Contact us](#)

CMS CATEGORIES

[Blog](#) (28)

[Clan CMS](#) (4)

[CMF](#) (1)

[CMS / Portals](#) (343)

[CRM](#) (1)

[ECM](#) (1)

[eCommerce](#) (22)

[Forum](#) (15)

[Groupware](#) (7)

[Image Galleries](#) (8)

[Lite / Simple](#) (11)

[LMS / LCMS](#) (5)

[MVC](#) (2)

[Social Dating](#) (1)

[WAF](#) (1)

[Wiki](#) (4)

POPULAR DEMOS

[WordPress Demo](#)

[Drupal Demo](#)

[MODx Demo](#)

[Typo3 Demo](#)

[Joomla Demo](#)

[concrete5 Demo](#)

NEW DEMOS

GetSimple CMS v3.3.16 was discovered to contain a remote...

Critical severity

Unreviewed

Published on Oct 18, 2022 to the GitHub Advisory Database • Updated on May 24, 2023

Package

No package listed— Suggest a package

Affected versions

Unknown

Patched versions

Unknown

Description

GetSimple CMS v3.3.16 was discovered to contain a remote code execution (RCE) vulnerability via the `edited_file` parameter in `admin/theme-edit.php`.

References

- <https://nvd.nist.gov/vuln/detail/CVE-2022-41544>
- [GetSimpleCMS/GetSimpleCMS#1352](#)
- <http://packetstormsecurity.com/files/172553/GetSimple-CMS-3.3.16-Shell-Upload.html>



Published by the **National Vulnerability Database** on Oct 18, 2022



Published to the GitHub Advisory Database on Oct 18, 2022



Last updated on May 24, 2023

Severity

Critical 9.8 / 10

CVSS v3 base metrics

Attack vector	Network
Attack complexity	Low
Privileges required	None
User interaction	None
Scope	Unchanged
Confidentiality	High
Integrity	High
Availability	High

[Learn more about base metrics](#)

CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H

EPSS score

4.568% (92nd percentile)

Weaknesses

No CWEs

CVE ID

CVE-2022-41544

GHSA ID

GHSA-c599-8gm5-c2jh

Source code

No known source code

```
$ root@simpecms-...:/root#
```

```
$ root@simpecms-...:/root# env
KUBERNETES_SERVICE_PORT_HTTPS=443
KUBERNETES_SERVICE_PORT=443
HOSTNAME=simpecms-6c4bfc97cd-wdnhv
PHP_VERSION=7.4.9
APACHE_CONFDIR=/etc/apache2
PHP_LDFLAGS=-Wl,-01 -pie
PWD=/var/www/html
HOME=/root
KUBERNETES_PORT_443_TCP=tcp://10.96.0.1:443
TERM=xterm
PHP_URL=https://www.php.net/distributions/php-7.4.9.tar.xz
...
```

```
$ root@simpecms-...:/root# env
KUBERNETES_SERVICE_PORT_HTTPS=443
KUBERNETES_SERVICE_PORT=443
HOSTNAME=simpecms-6c4bfc97cd-wdnhv
PHP_VERSION=7.4.9
APACHE_CONFDIR=/etc/apache2
PHP_LDFLAGS=-Wl,-O1 -pie
PWD=/var/www/html
HOME=/root
KUBERNETES_PORT_443_TCP=tcp://10.96.0.1:443
TERM=xterm
PHP_URL=https://www.php.net/distributions/php-7.4.9.tar.xz
...
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/  
total 20  
drwxr-xr-x 4 root root 4096 Jan  5 09:57 .  
drwxr-xr-x 1 root root 4096 Jan  5 09:57 ..  
drwxr-xr-x 2 root root 4096 Jan  5 09:57 eks.amazonaws.com  
drwxr-xr-x 3 root root 4096 Jan  5 09:57 kubernetes.io
```



```
$ root@simpecms-...:/root# ls -al /var/run/secrets/  
total 20  
drwxr-xr-x 4 root root 4096 Jan  5 09:57 .  
drwxr-xr-x 1 root root 4096 Jan  5 09:57 ..  
drwxr-xr-x 2 root root 4096 Jan  5 09:57 eks.amazonaws.com  
drwxr-xr-x 3 root root 4096 Jan  5 09:57 kubernetes.io
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/
```

```
total 20
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 1 root root 4096 Jan  5 09:57 ..
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 eks.amazonaws.com
```

```
drwxr-xr-x 3 root root 4096 Jan  5 09:57 kubernetes.io
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/eks.amazonaws.com/
```

```
total 12
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 ..
```

```
-rw-r--r-- 1 root root   16 Jan  5 09:57 access_key_token
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/
```

```
total 20
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 1 root root 4096 Jan  5 09:57 ..
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 eks.amazonaws.com
```

```
drwxr-xr-x 3 root root 4096 Jan  5 09:57 kubernetes.io
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/eks.amazonaws.com/
```

```
total 12
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 ..
```

```
-rw-r--r-- 1 root root   16 Jan  5 09:57 access_key_token
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/
```

```
total 20
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 1 root root 4096 Jan  5 09:57 ..
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 eks.amazonaws.com
```

```
drwxr-xr-x 3 root root 4096 Jan  5 09:57 kubernetes.io
```

```
$ root@simpecms-...:/root# ls -al /var/run/secrets/eks.amazonaws.com/
```

```
total 12
```

```
drwxr-xr-x 2 root root 4096 Jan  5 09:57 .
```

```
drwxr-xr-x 4 root root 4096 Jan  5 09:57 ..
```

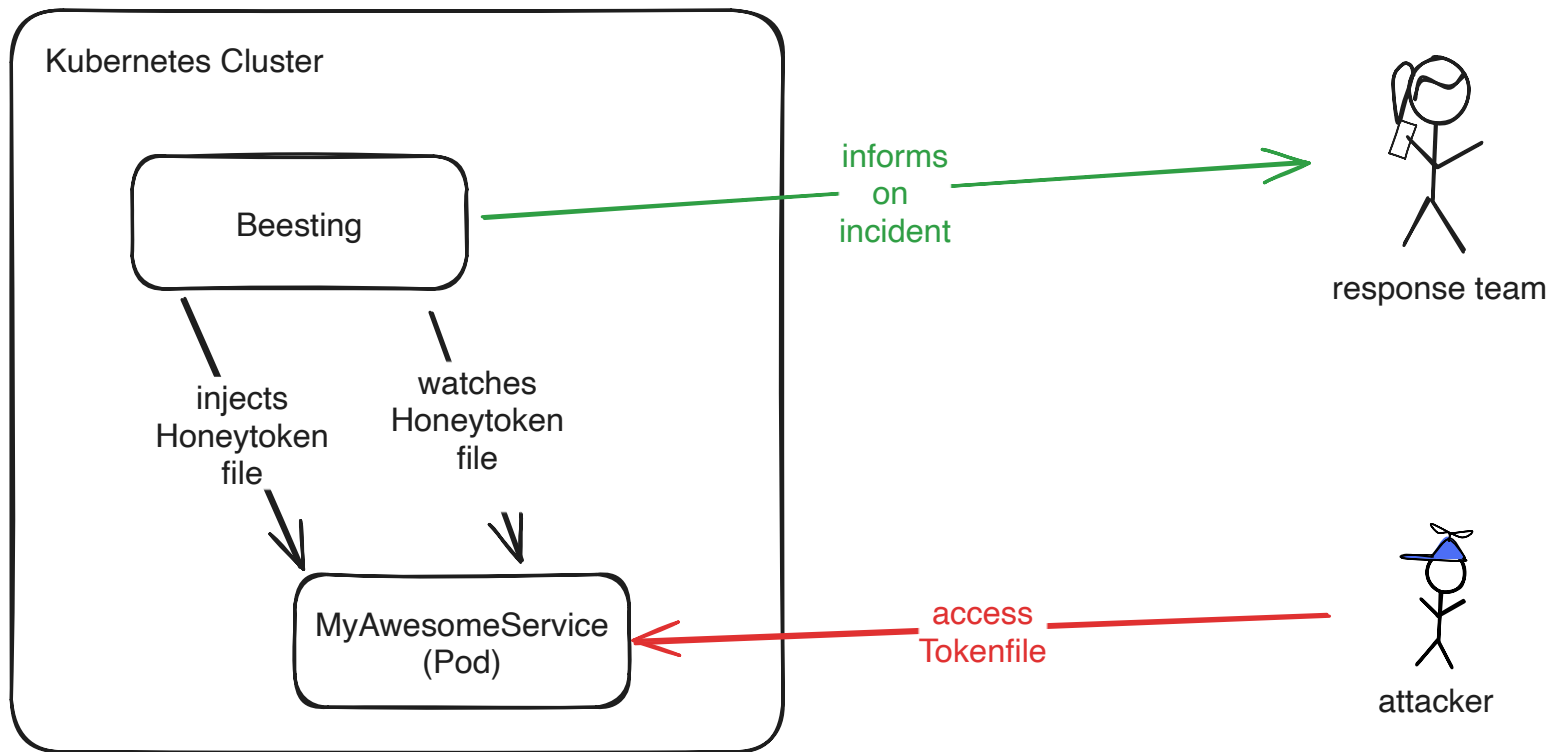
```
-rw-r--r-- 1 root root   16 Jan  5 09:57 access_key_token
```

```
$ root@simpecms-...:/root# cat /v/r/s/eks.amazonaws.com/access_key_token
```

SECURITY TEAM

**YOU JUST ACTIVATED
MY TRAP CARD**





How does
Kubernetes
work?



Control Plane

API Server

EtcD

Scheduler

Data Plane

Node 1

kubelet

my-fancy-app

Node 2

kubelet

awesome-app

Control Plane



API Server

EtcD

Scheduler

Data Plane

Node 1

kubelet

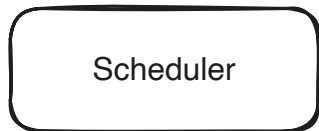
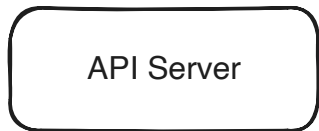
my-fancy-app

Node 2

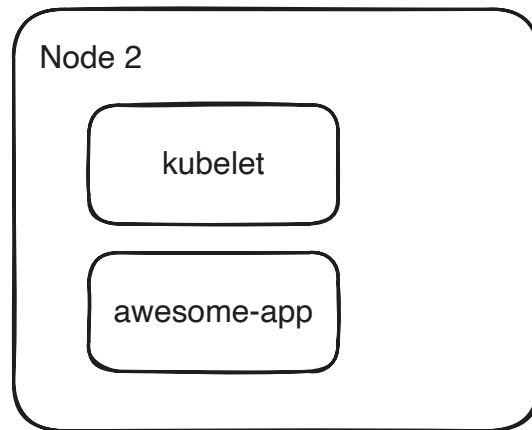
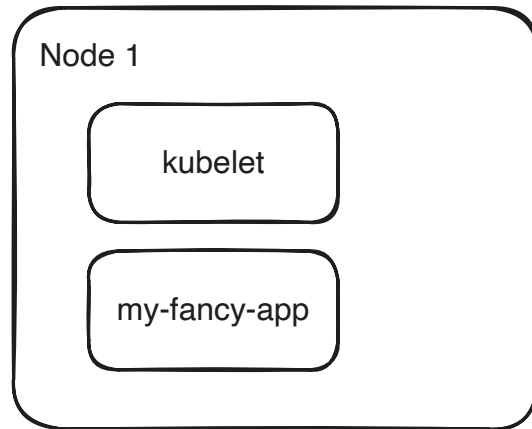
kubelet

awesome-app

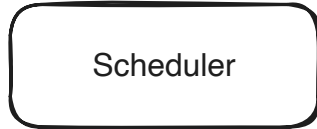
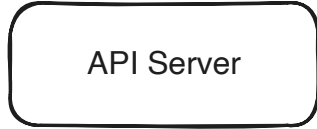
Control Plane



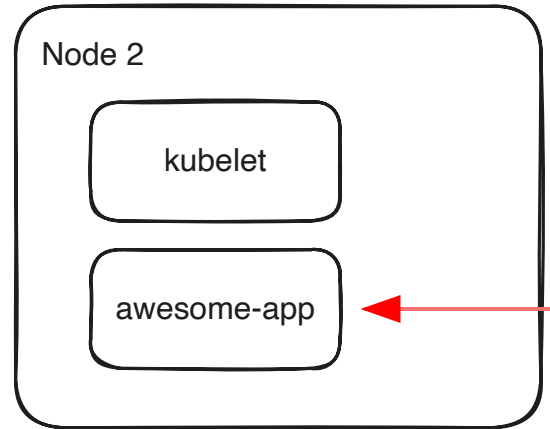
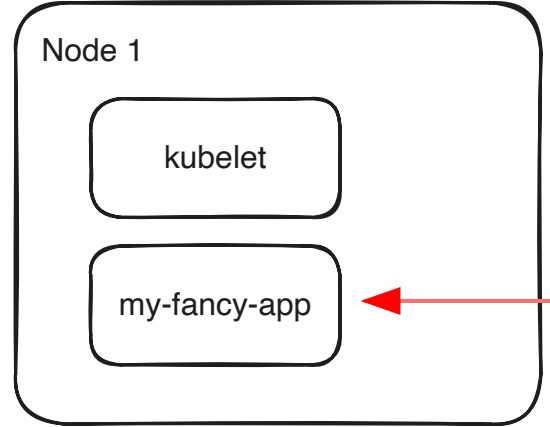
Data Plane



Control Plane

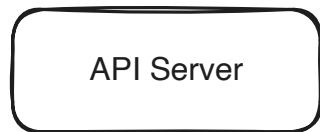


Data Plane



Control Plane

Data Plane



Node 1

containerD
(Container Runtime)

my-fancy-app
("container")

request
start
container
(through CRI)

starts
containers
through

runs

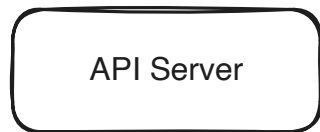
checks
for
work

kubelet

runC
(lowlevel runtime)

Control Plane

Data Plane



Node 1

containerD
(Container Runtime)

my-fancy-app
("container")

request
start
container
(through CRI)

starts
containers
through

runs

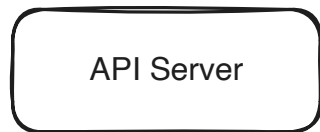
checks
for
work

kubelet

runC
(lowlevel runtime)

Control Plane

Data Plane



Node 1

containerD
(Container Runtime)

my-fancy-app
("container")

request
start
container
(through CRI)

starts
containers
through

runs

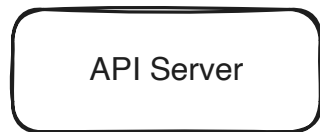
checks
for
work

kubelet

runC
(lowlevel runtime)

Control Plane

Data Plane



Node 1

containerD
(Container Runtime)

my-fancy-app
("container")

request
start
container
(through CRI)

starts
containers
through

runs

checks
for
work

kubelet

runC
(lowlevel runtime)



Cgroups

Limit resource usage (CPU, Memory)

Namespaces

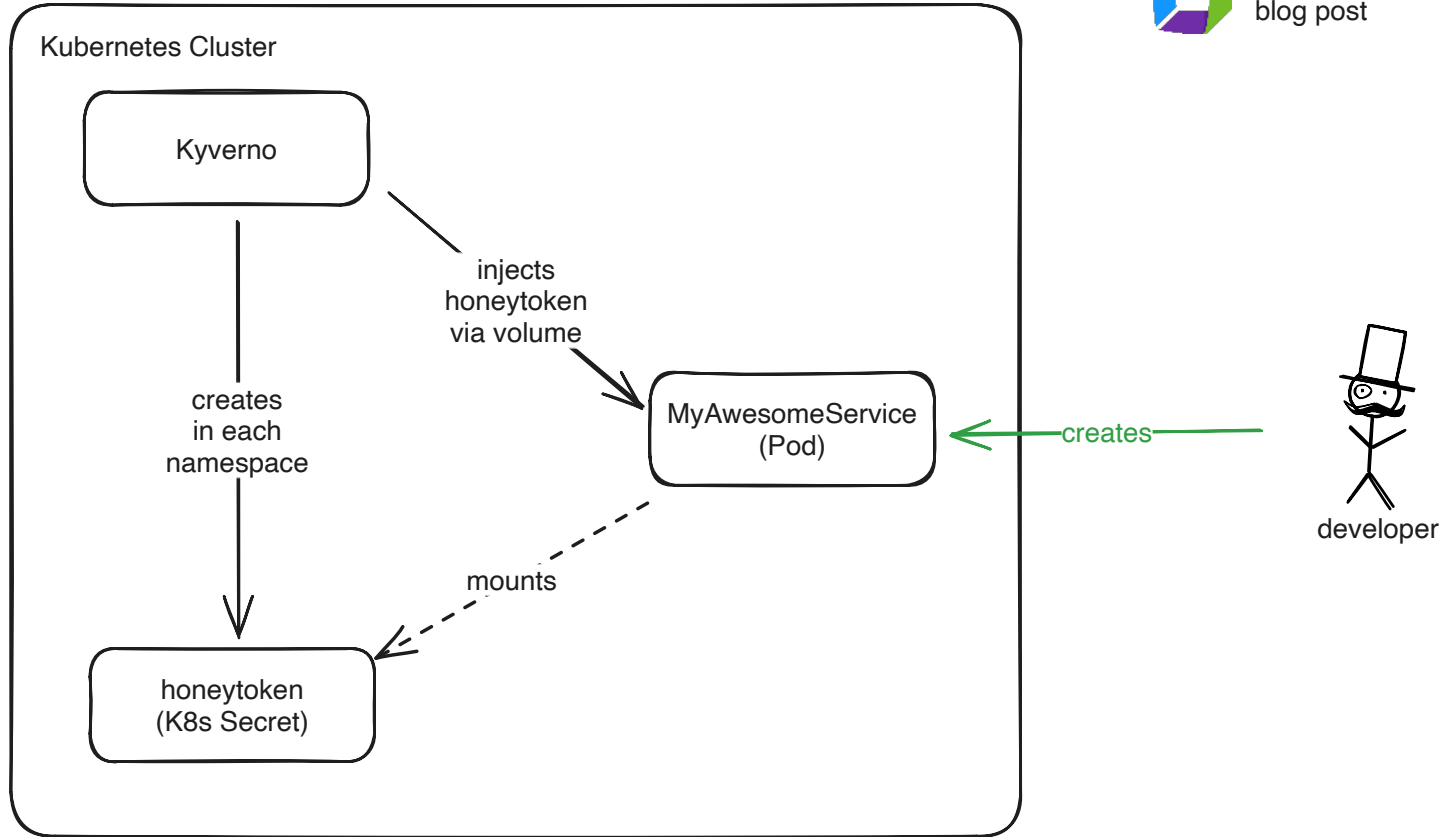
Separate areas (Networking, Mounts)



**How do we get the
Honeytoken
into the container?**

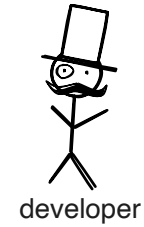
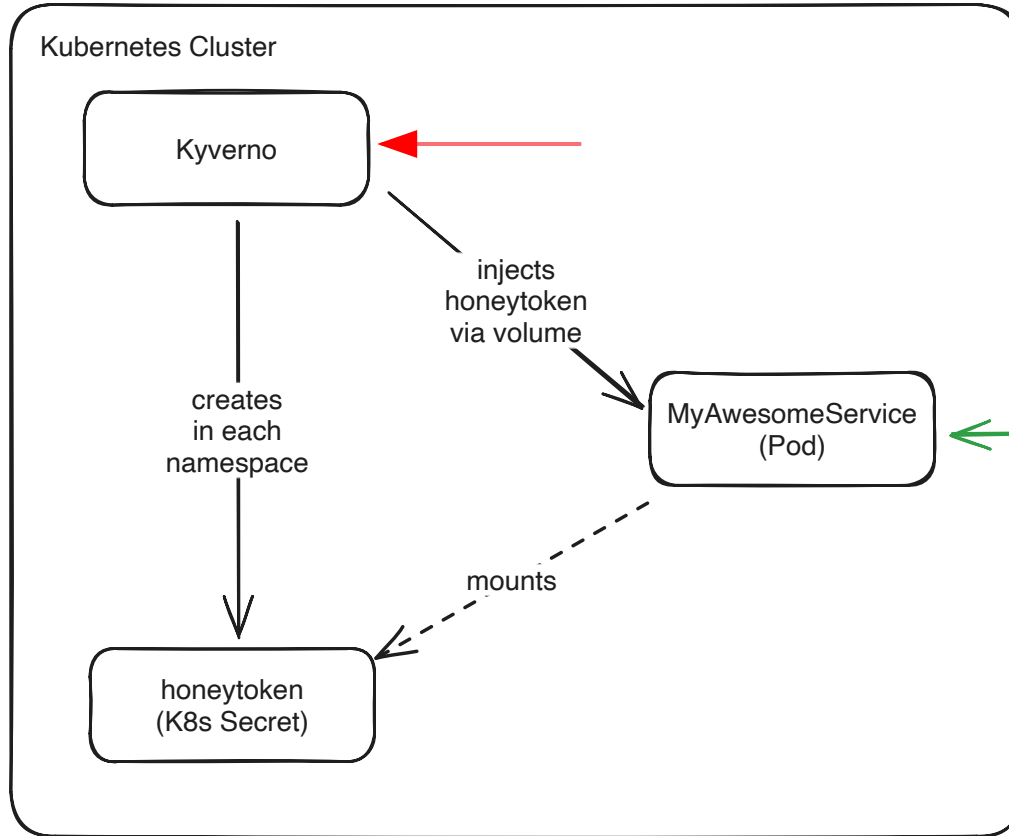


Solution from Dynatrace
blog post



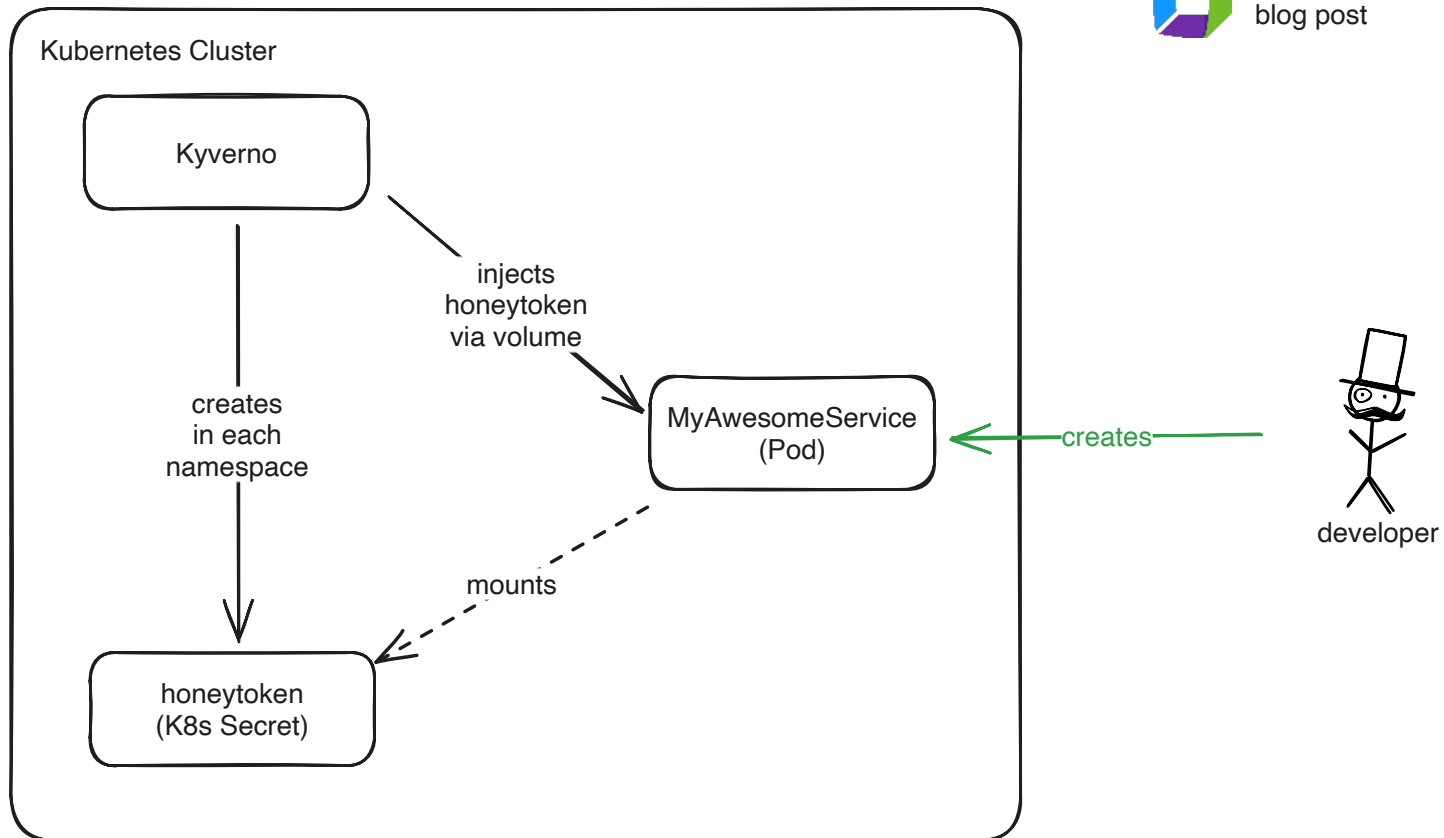


Solution from Dynatrace
blog post





Solution from Dynatrace
blog post




```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```

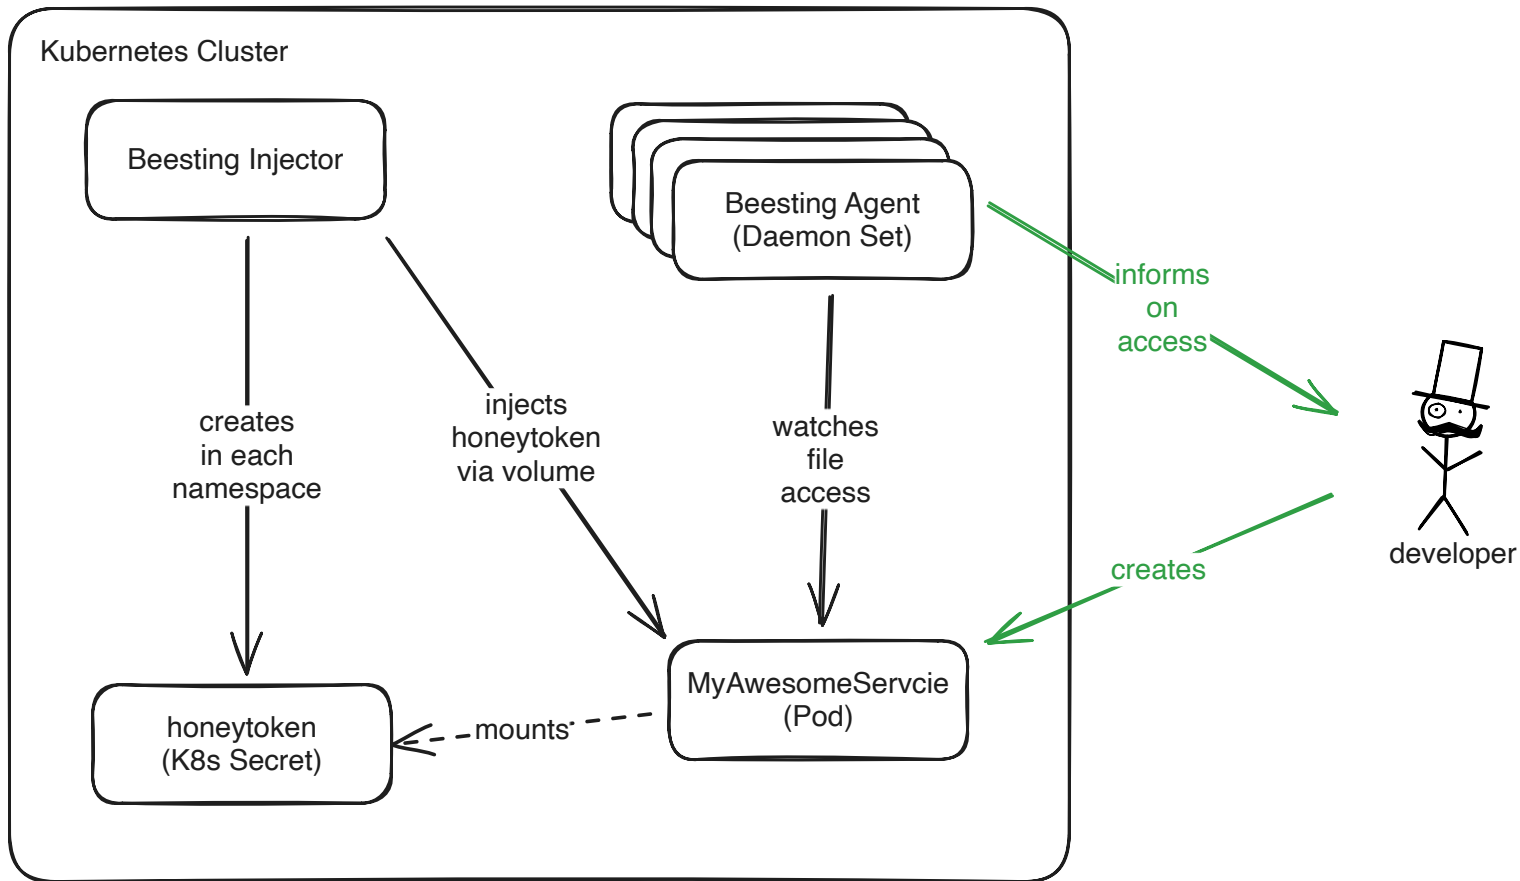
```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```

```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```

```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```

```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```

```
apiVersion: v1
kind: Pod
metadata:
  name: "myapp"
  namespace: default
spec:
  containers:
    - name: myapp
      image: "myapp:latest"
      volumeMounts:
        - name: honey-volume
          readOnly: true
          subPath: token
          mountPath: /run/secrets/eks.amazonaws.com/s3_token
  volumes:
    - name: honeytoken
      secret:
        secretName: honeytoken
```



Problems

Problems

- Lot of moving parts

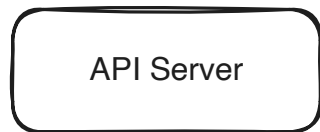
Problems

- Lot of moving parts
- Kubernetes specific

Boring

Control Plane

Data Plane



Node 1

containerD
(Container Runtime)

my-fancy-app
("container")

request
start
container
(through CRI)

starts
containers
through

runs

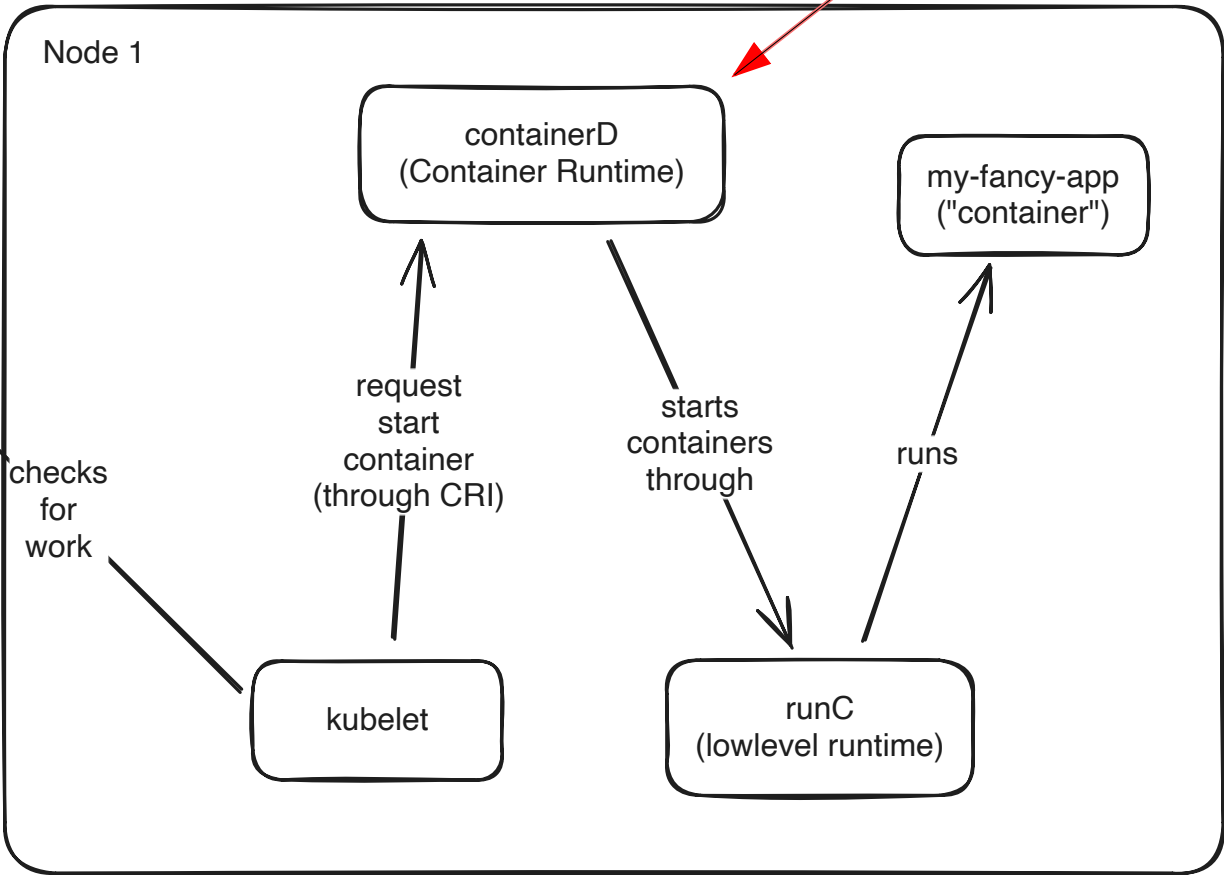
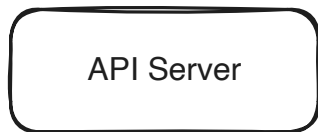
checks
for
work

kubelet

runC
(lowlevel runtime)

Control Plane

Data Plane



containerd Plugins

containerd supports extending its functionality using most of its defined interfaces. This includes using a customized runtime, snapshotter, content store, and even adding gRPC interfaces.

Smart Client Model

containerd has a smart client architecture, meaning any functionality which is not required by the daemon is done by the client. This includes most high level interactions such as creating a container's specification, interacting with an image registry, or loading an image from tar. containerd's Go client gives a user access to many points of extensions from creating their own options on container creation to resolving image registry names.

See [containerd's Go documentation](#)

External Plugins

External plugins allow extending containerd's functionality using an officially released version of containerd without needing to recompile the daemon to add a plugin.

containerd allows extensions through two methods:

- via a binary available in containerd's PATH
- by configuring containerd to proxy to another gRPC service

V2 Runtimes

containerd supports multiple container runtimes. Each container can be invoked with a different runtime.

When using the Container Runtime Interface (CRI) plugin, named runtimes can be defined in the containerd configuration file. When a container is run without specifying a runtime, the configured default runtime is used. Alternatively, a different named runtime can be specified explicitly when creating a container via CRI gRPC by selecting the runtime handler to be used.

When a client such as `ctr` or `nerdctl` creates a container, it can optionally specify a runtime and options to use. If a runtime is not specified, containerd will use its default runtime.

containerd invokes v2 runtimes as binaries on the system, which are used to start the shim process for containerd. This, in turn, allows containerd to start and manage those containers using the runtime shim api returned by the binary.

For more details on runtimes and shims, including how to invoke and configure them, see the [runtime v2 documentation](#)

- Runtime

- Runtime
- Differ

- Runtime
- Differ
- and a few more

Differ Plugins

```
// Diff service creates and applies diffs
service Diff {
    // Apply applies the content associated with the provided digests onto
    // the provided mounts. Archive content will be extracted and
    // decompressed if necessary.
    rpc Apply(ApplyRequest) returns (ApplyResponse);

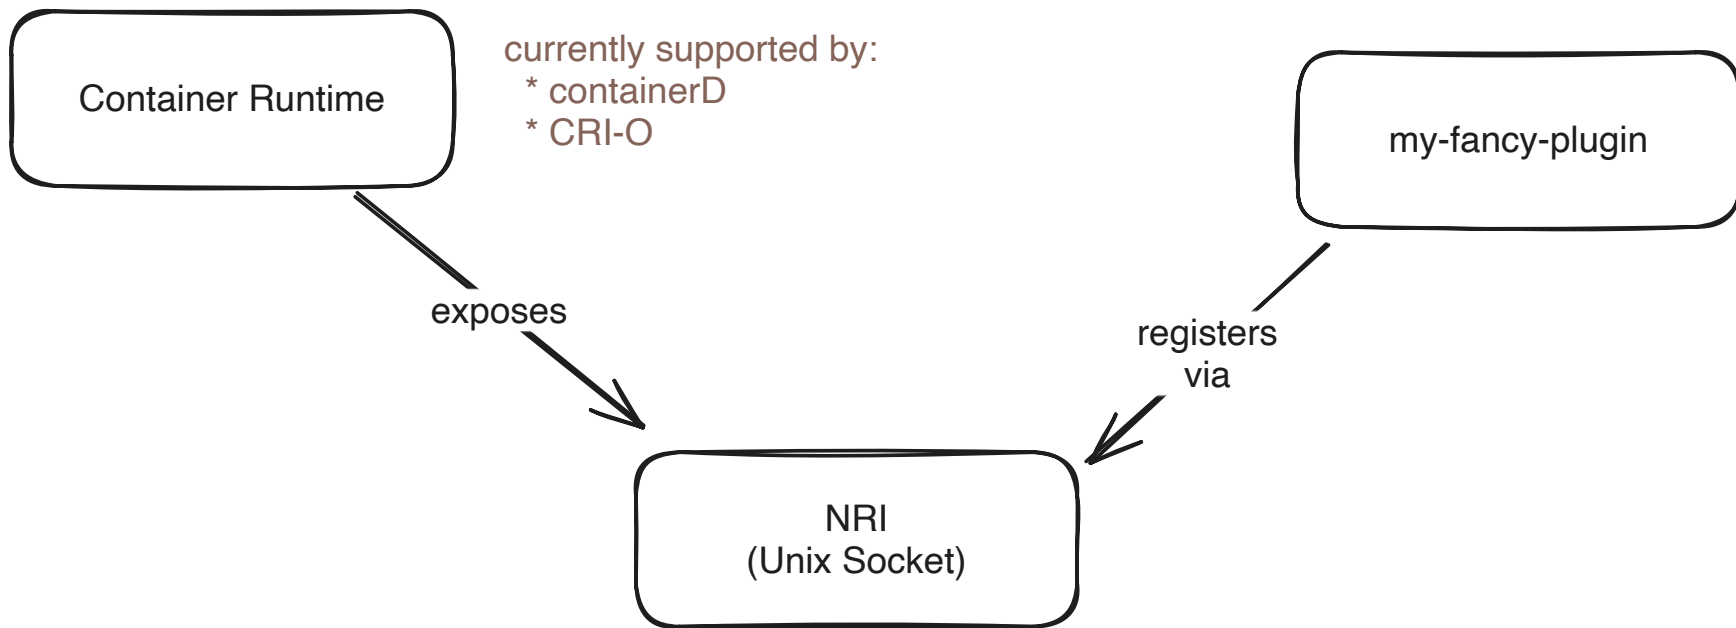
    // Diff creates a diff between the given mounts and uploads the result
    // to the content store.
    rpc Diff(DiffRequest) returns (DiffResponse);
}
```

```
// Diff service creates and applies diffs
service Diff {
    // Apply applies the content associated with the provided digests onto
    // the provided mounts. Archive content will be extracted and
    // decompressed if necessary.
    rpc Apply(ApplyRequest) returns (ApplyResponse);

    // Diff creates a diff between the given mounts and uploads the result
    // to the content store.
    rpc Diff(DiffRequest) returns (DiffResponse);
}
```

```
message DiffRequest {  
    repeated containerd.types.Mount left = 1;  
    repeated containerd.types.Mount right = 2;  
    string media_type = 3;  
    string ref = 4;  
    map<string, string> labels = 5;  
    google.protobuf.Timestamp source_date_epoch = 6;  
}
```

Node Resource Interface



```
type handlers struct {
    Configure      func(...) (api.EventMask, error)
    Synchronize    func(...) ([]*api.ContainerUpdate, error)
    Shutdown       func(...)
    RunPodSandbox  func(...) error
    StopPodSandbox func(...) error
    RemovePodSandbox func(...) error
    CreateContainer func(...) (*api.ContainerAdjustment, []*api.ContainerUpdate, error)
    StartContainer func(...) error
    UpdateContainer func(...) ([]*api.ContainerUpdate, error)
    StopContainer  func(...) ([]*api.ContainerUpdate, error)
    RemoveContainer func(...) error
    PostCreateContainer func(...) error
    PostStartContainer func(...) error
    PostUpdateContainer func(...) error
}
```



```
type handlers struct {  
    Configure      func(...) (api.EventMask, error)  
    Synchronize    func(...) ([]*api.ContainerUpdate, error)  
    Shutdown       func(...)  
    RunPodSandbox  func(...) error  
    StopPodSandbox func(...) error  
    RemovePodSandbox func(...) error  
    CreateContainer func(...) (*api.ContainerAdjustment, []*api.ContainerUpdate, error)  
    StartContainer  func(...) error  
    UpdateContainer func(...) ([]*api.ContainerUpdate, error)  
    StopContainer   func(...) ([]*api.ContainerUpdate, error)  
    RemoveContainer func(...) error  
    PostCreateContainer func(...) error  
    PostStartContainer func(...) error  
    PostUpdateContainer func(...) error  
}
```

```
type ContainerAdjustment struct {  
    Annotations map[string]string  
    Mounts      []*Mount  
    Env         []*KeyValue  
    Hooks       *Hooks  
    Linux       *LinuxContainerAdjustment  
    Rlimits     []*POSIXRlimit  
    CDIDevices  []*CDIDevice  
}
```

```
type ContainerAdjustment struct {  
    Annotations map[string]string  
    Mounts      []*Mount  
    Env          []*KeyValue  
    Hooks        *Hooks  
    Linux        *LinuxContainerAdjustment  
    Rlimits      []*POSIXRlimit  
    CDIDevices   []*CDIDevice  
}
```

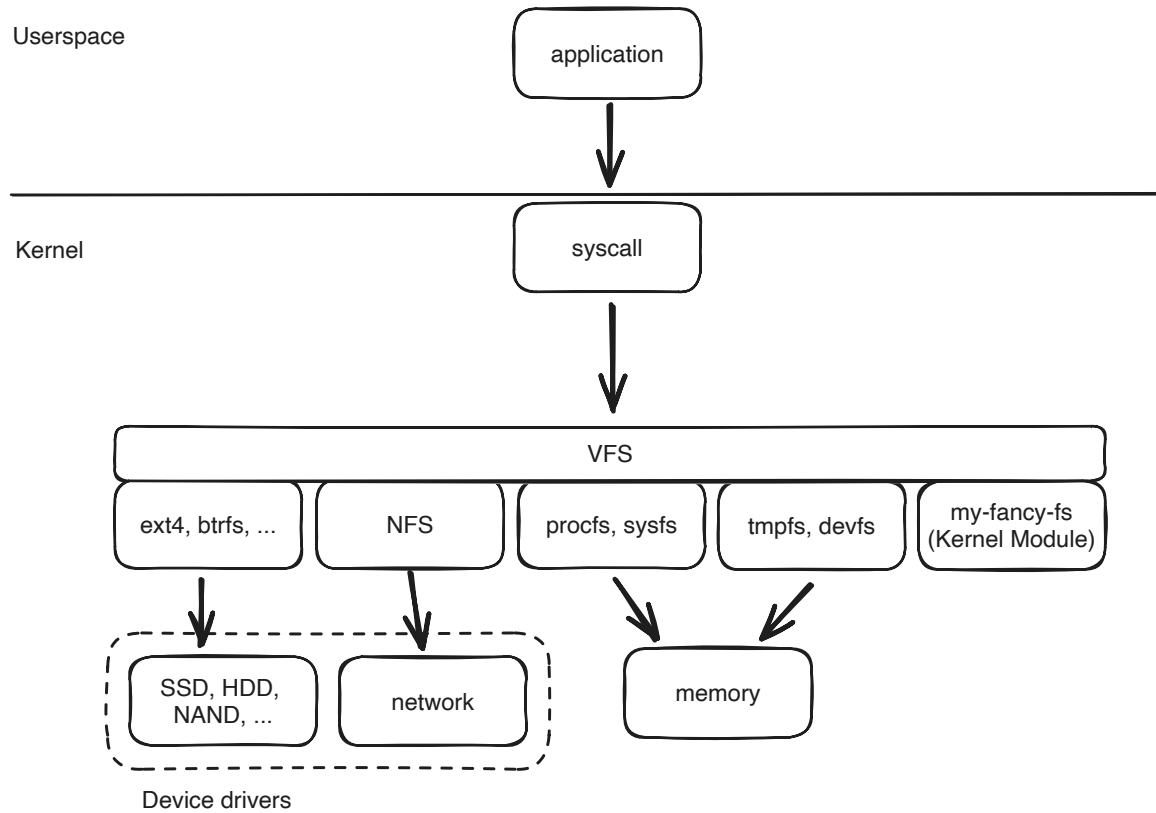
```
type ContainerAdjustment struct {  
    Annotations map[string]string  
    Mounts      []*Mount  
    Env         []*KeyValue  
    Hooks       *Hooks  
    Linux       *LinuxContainerAdjustment  
    Rlimits     []*POSIXRlimit  
    CDIDevices  []*CDIDevice  
}
```

**What to inject
into the
container?**



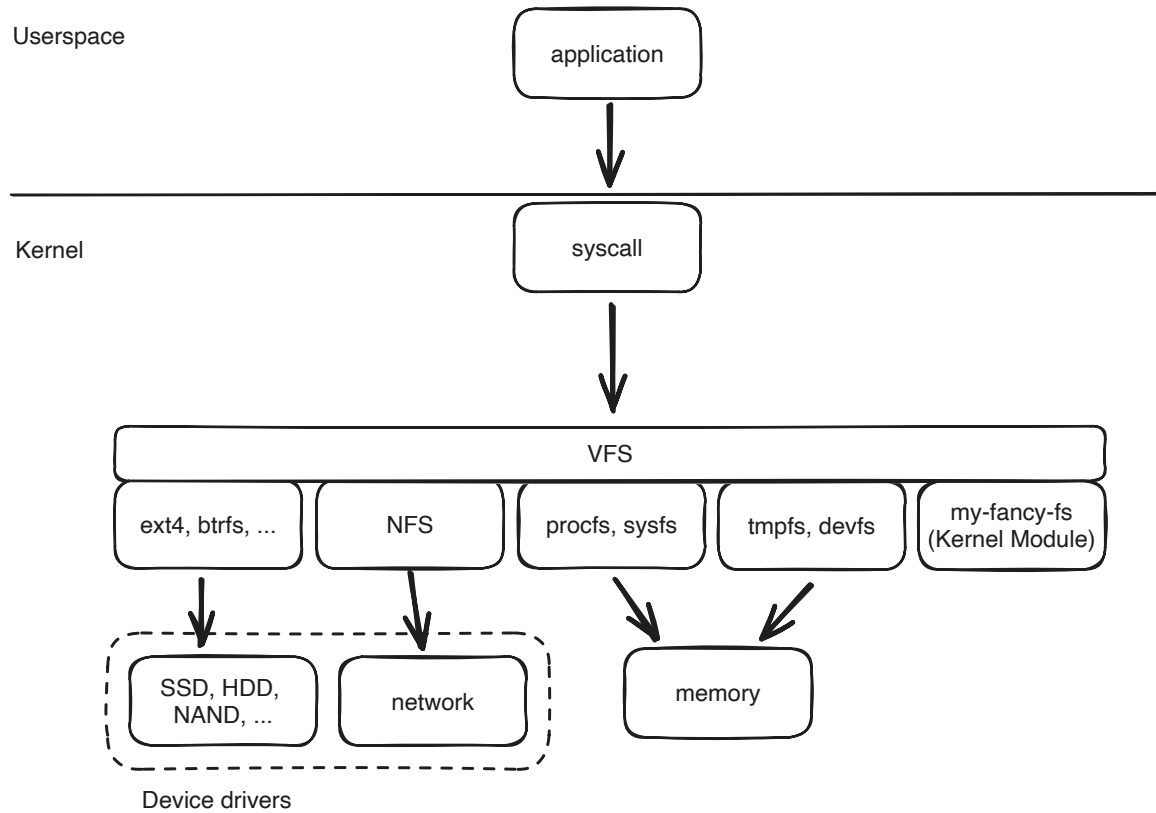
Honeytokens = Files

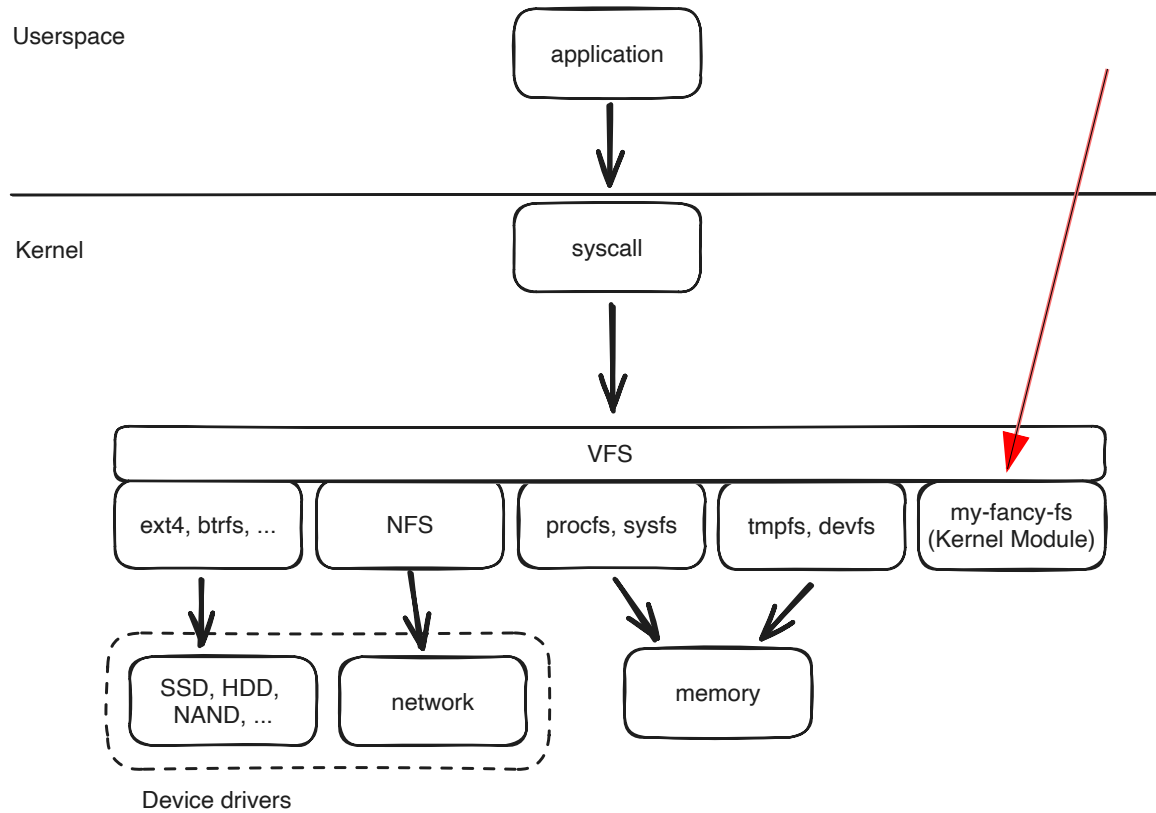
- ext2, ext3, ext3
- BTRFS
- XFS
- ZFS
- and many more



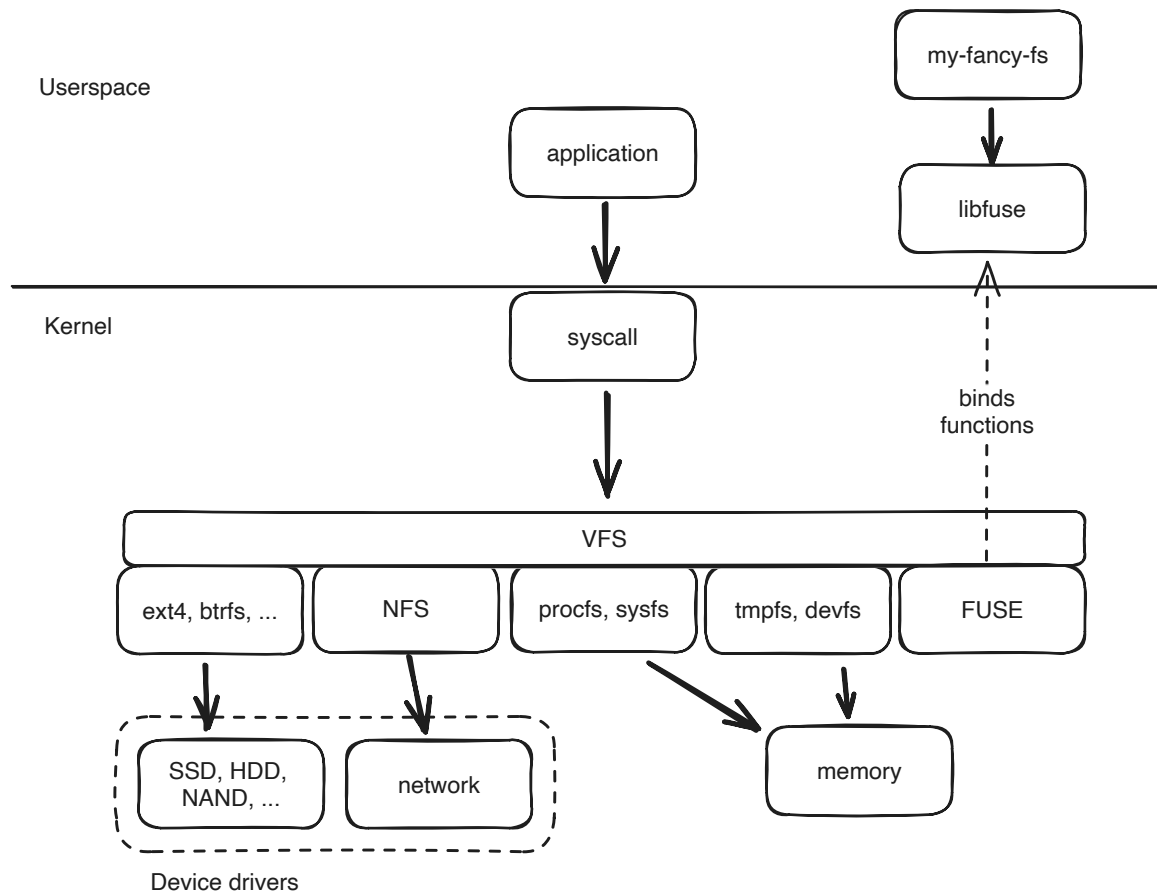


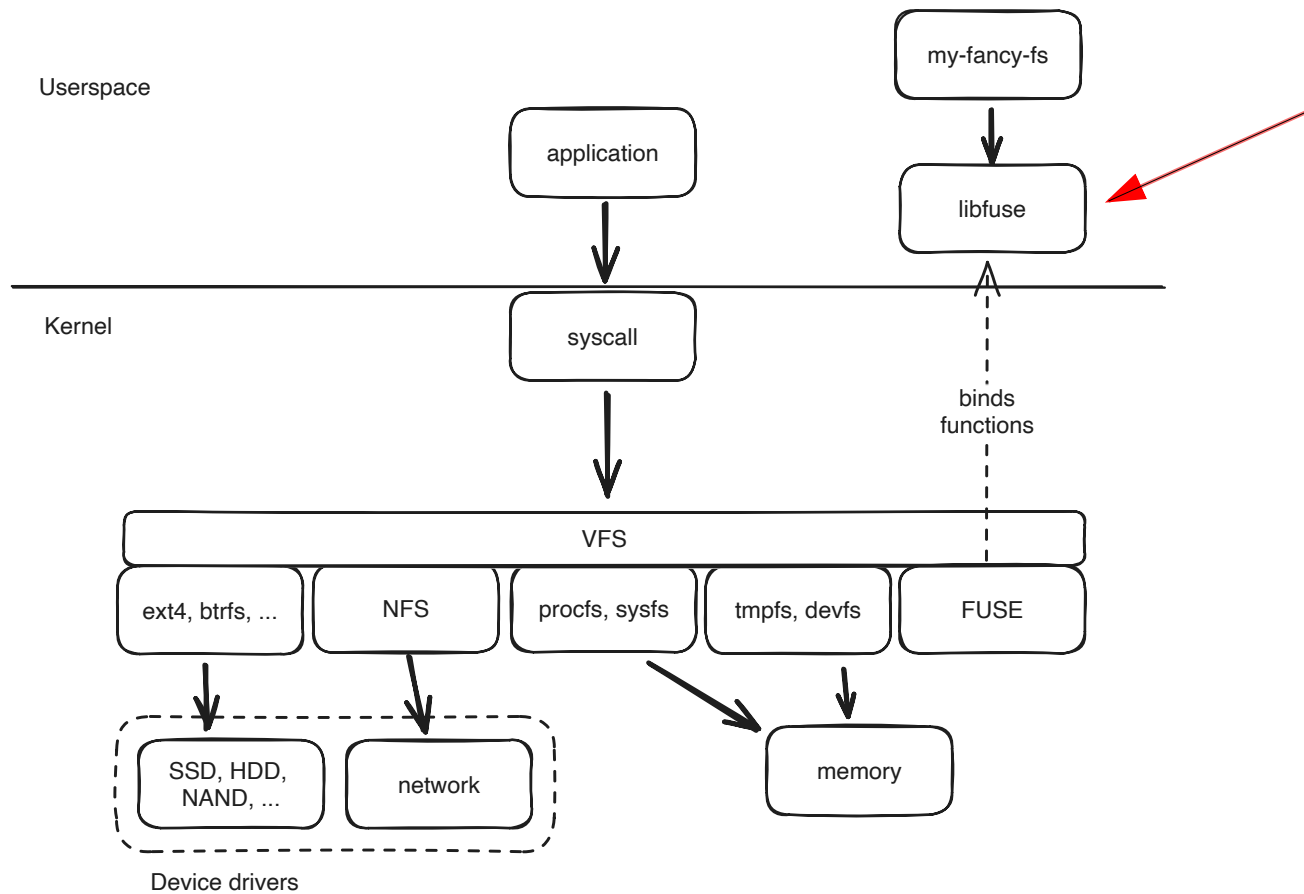
What if ...
we create our own
File System?











What about OCI Hooks?

```
type Hook struct {  
    Path    string  
    Args    []string  
    Env     []string  
    Timeout *OptionalInt  
}
```



```
type Hook struct {  
    Path    string  
    Args    []string  
    Env     []string  
    Timeout *OptionalInt  
}
```

```
type Hook struct {  
    Path    string  
    Args    []string  
    Env     []string  
    Timeout *OptionalInt  
}
```


```
type Hook struct {  
    Path      string  
    Args      []string  
    Env       []string  
    Timeout   *OptionalInt  
}
```


```
type Hook struct {  
    Path    string  
    Args    []string  
    Env     []string  
    Timeout *OptionalInt  
}
```


```
type Hook struct {  
    Path    string  
    Args    []string  
    Env     []string  
    Timeout *OptionalInt  
}
```

```
type Hooks struct {  
    Prestart      []*Hook  
    CreateRuntime []*Hook  
    CreateContainer []*Hook  
    StartContainer []*Hook  
    Poststart      []*Hook  
    Poststop       []*Hook  
}
```

```
type Hooks struct {  
    Prestart      []*Hook  
    CreateRuntime []*Hook  
    CreateContainer []*Hook  
    StartContainer []*Hook  
    Poststart      []*Hook  
    Poststop       []*Hook  
}
```

 v1.2.0 runtime-spec / config.md ↑ Top

Preview Code Blame 1160 lines (1024 loc) · 54.4 KB · 



CreateContainer Hooks

The `createContainer` hooks MUST be called as part of the [create](#) operation after the runtime environment has been created (according to the configuration in `config.json`) but before the `pivot_root` or any equivalent operation has been executed. The `createContainer` hooks MUST be called after the `createRuntime` hooks.

The `createContainer` hooks' path MUST resolve in the [runtime namespace](#). The `createContainer` hooks MUST be executed in the [container namespace](#).

For example, on Linux this would happen before the `pivot_root` operation is executed but after the mount namespace was created and setup.

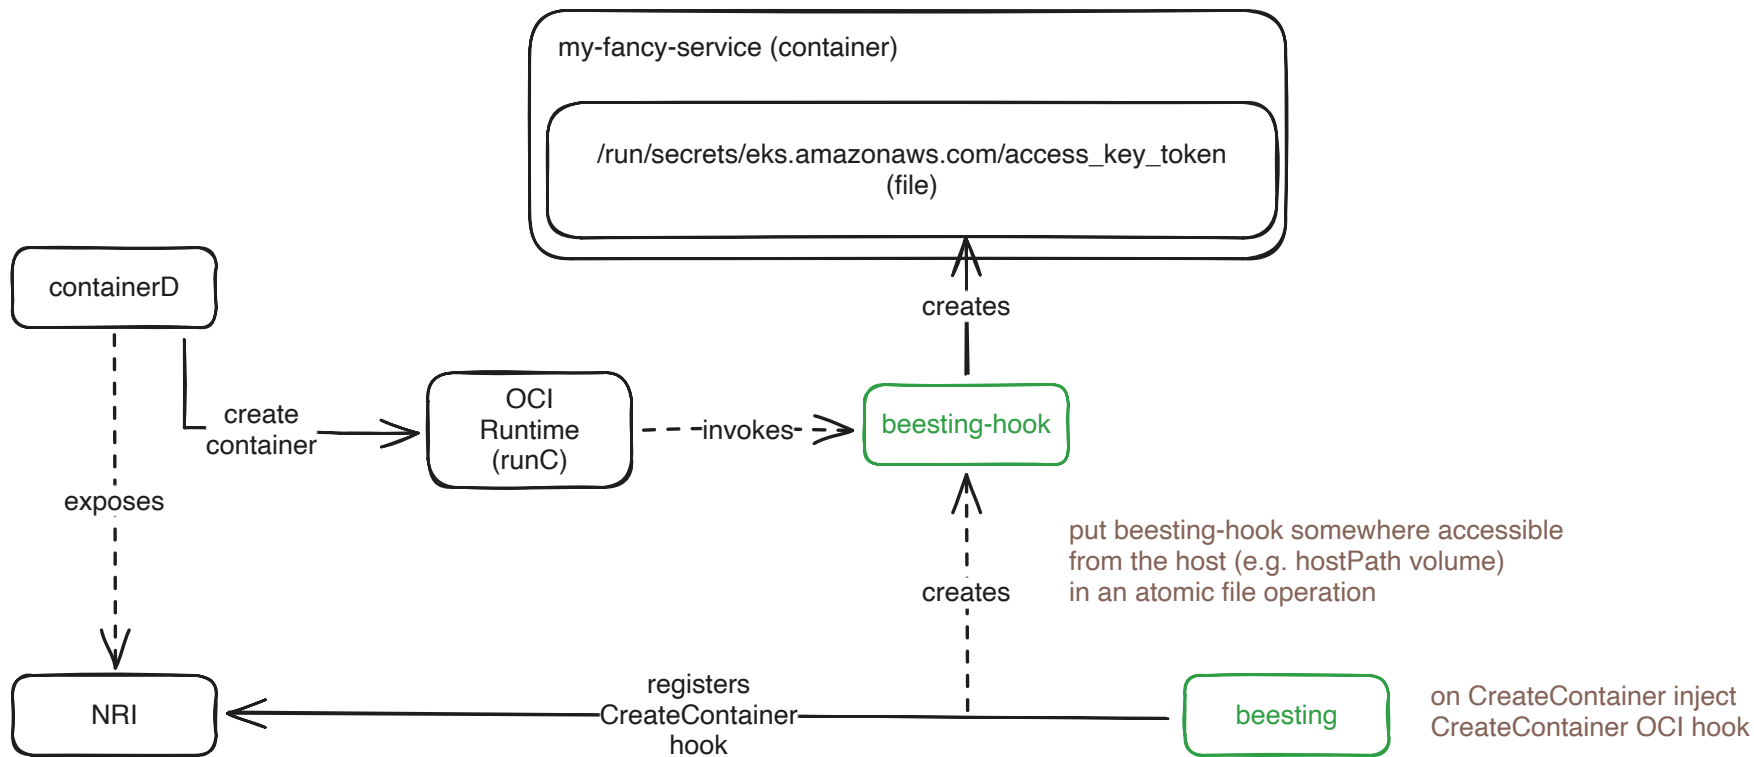
The definition of `createContainer` hooks is currently underspecified and hooks authors, should only expect from the runtime that the mount namespace and different mounts will be setup. Other operations such as `cgroups` and `SELinux/AppArmor` labels might not have been performed by the runtime.

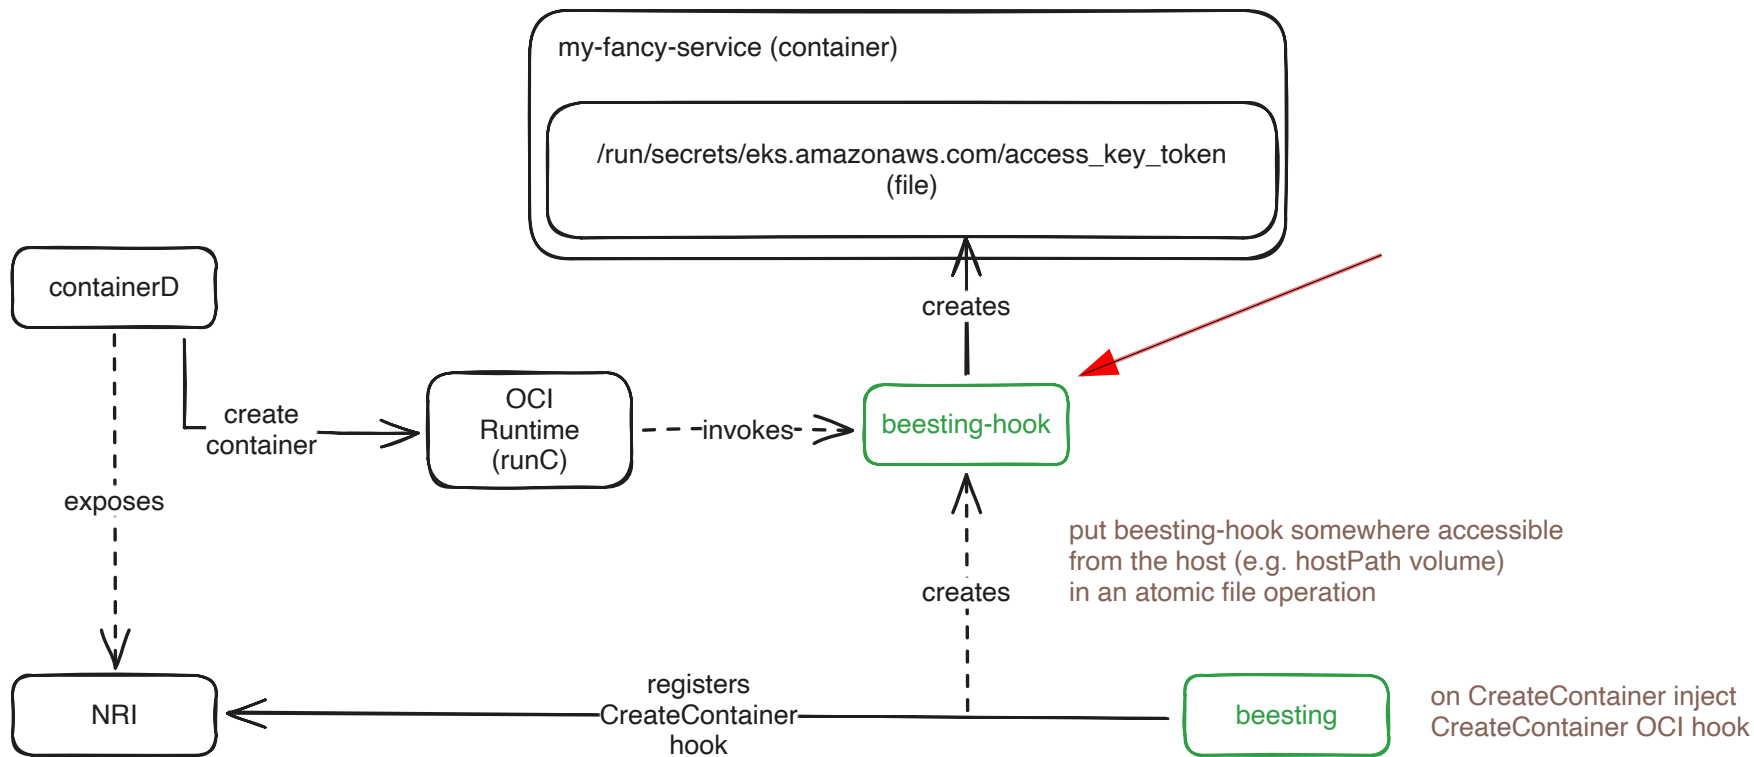
StartContainer Hooks

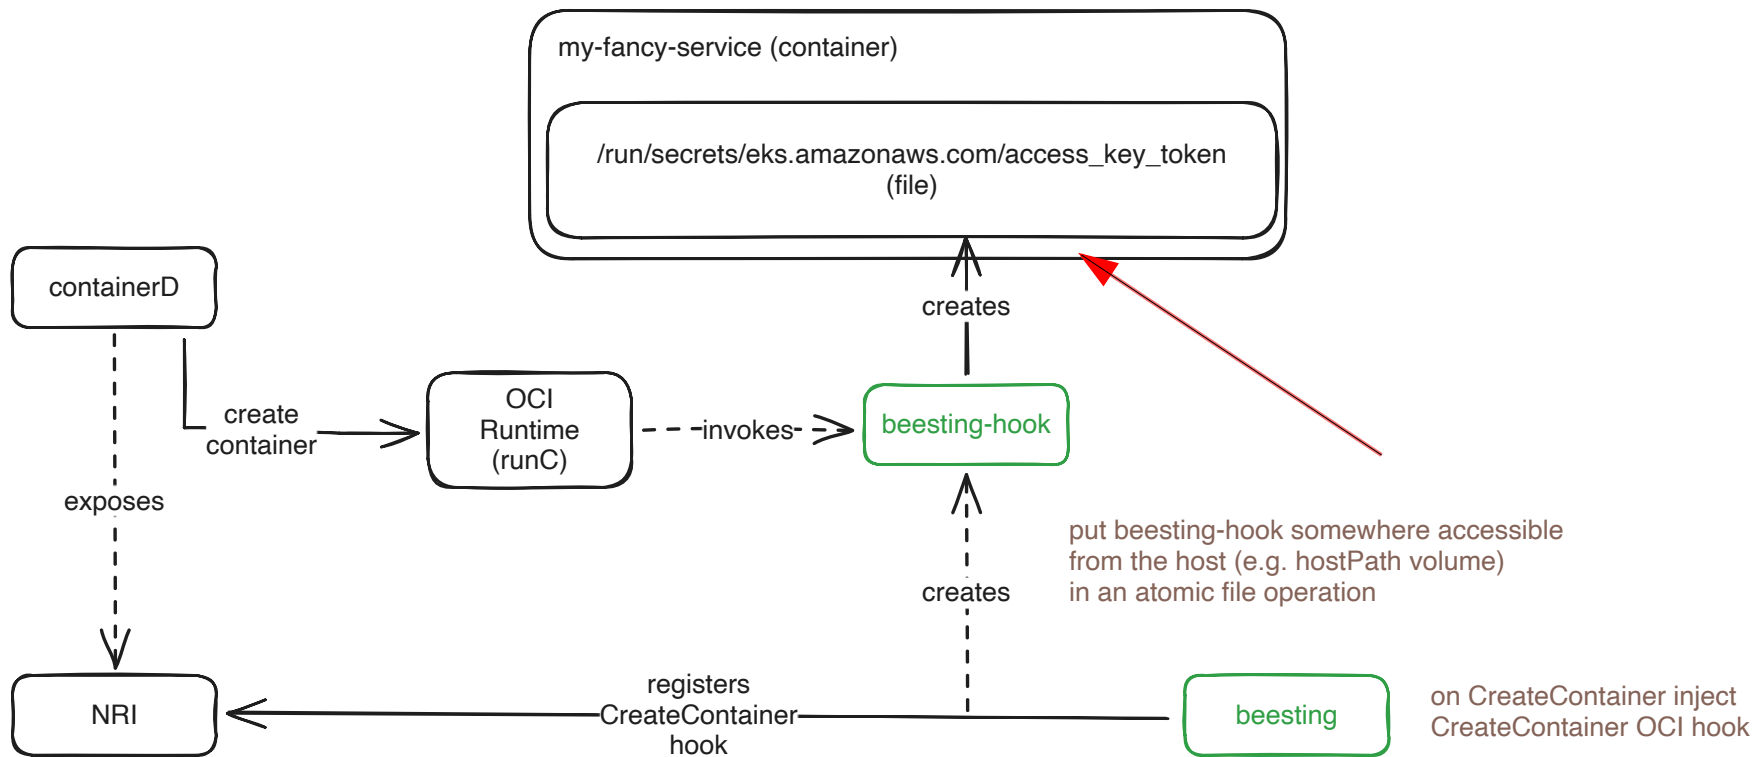
The `startContainer` hooks MUST be called [before the user-specified process is executed](#) as part of the `start` operation. This hook can be used to execute some operations in the container, for example running the `ldconfig` binary on linux before the container process is spawned.

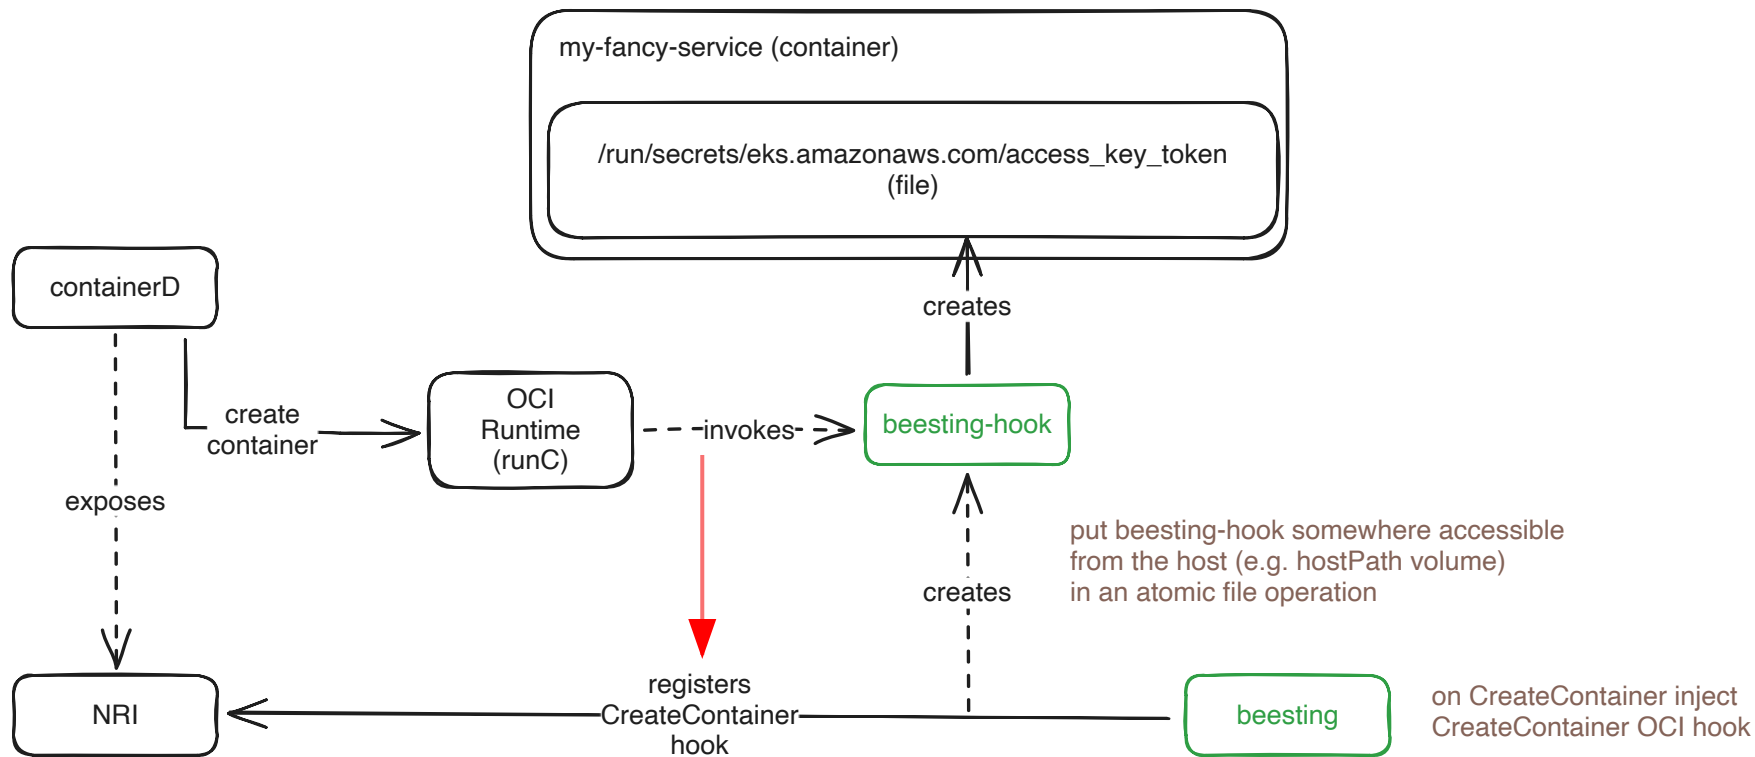
The `startContainer` hooks' path MUST resolve in the [container namespace](#). The `startContainer` hooks MUST be executed in the [container namespace](#).

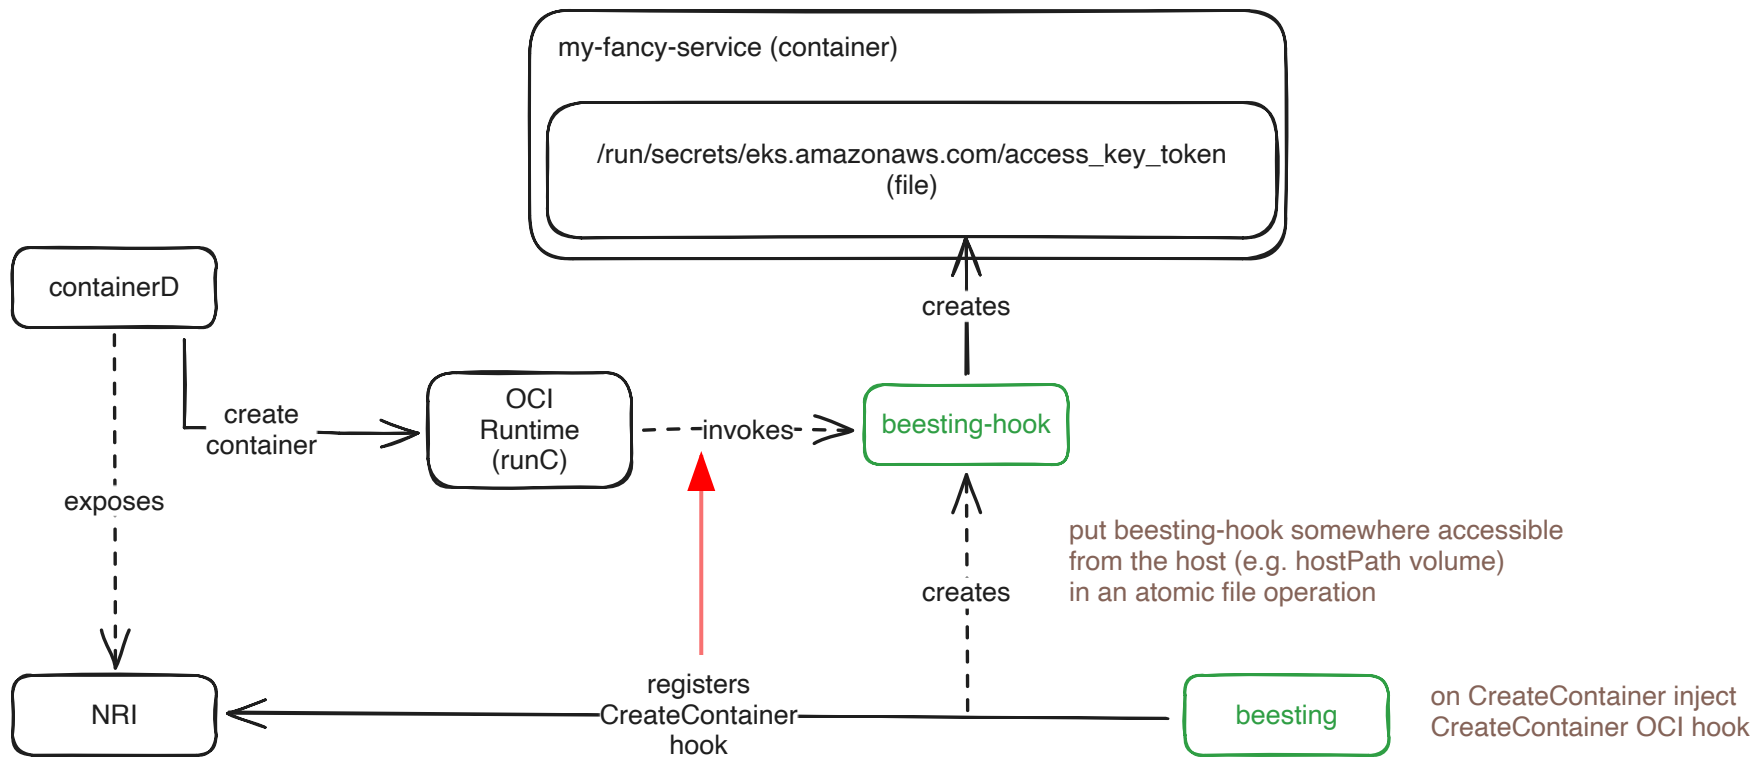
For now just file injection

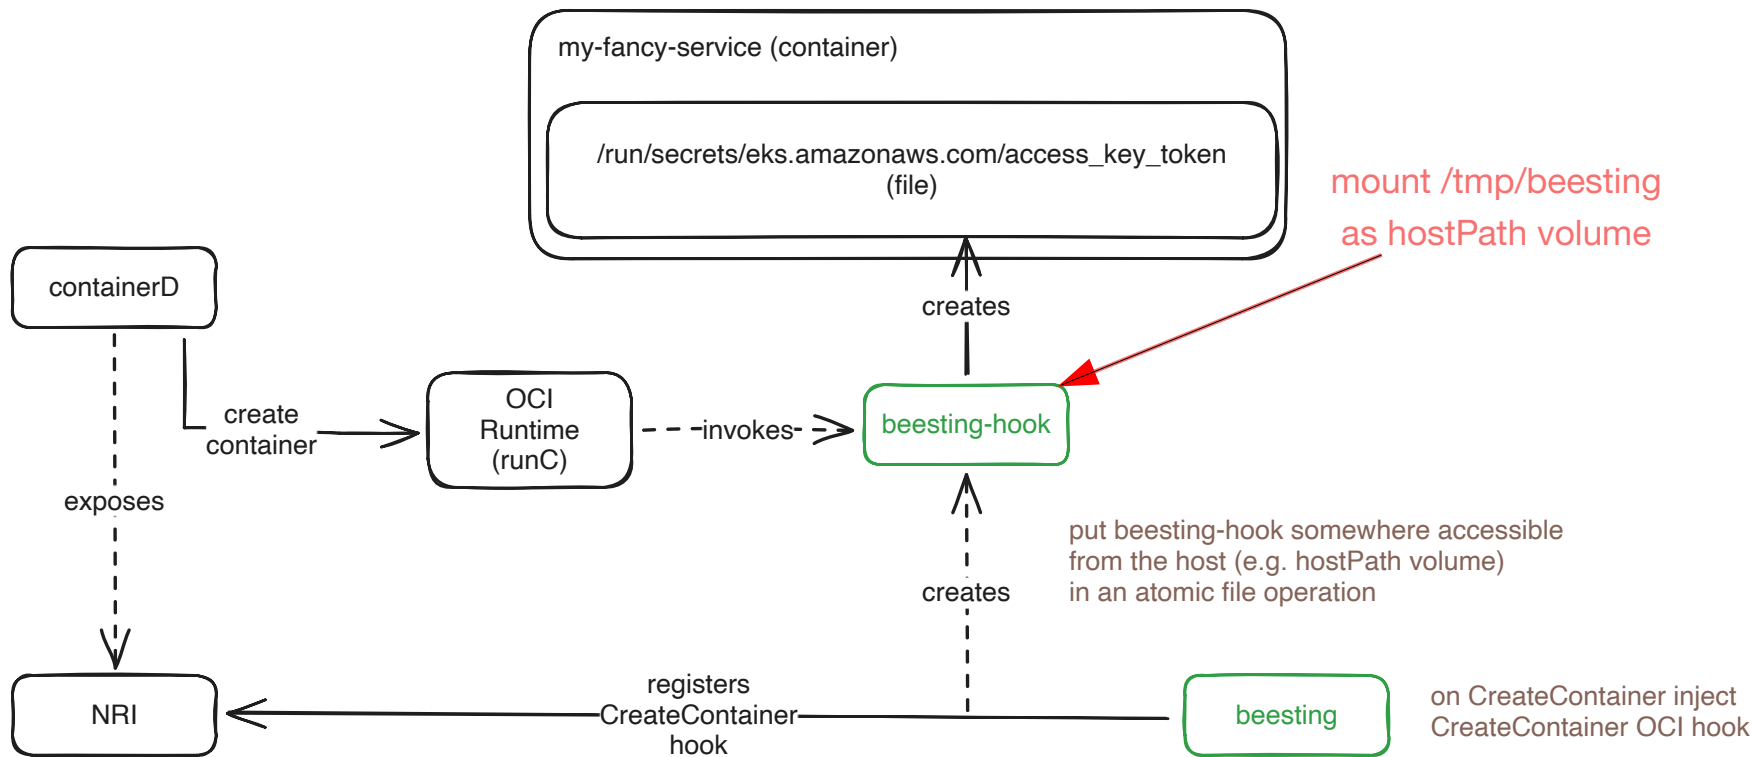












beesting-hook is embedded into
Beesting


```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-q9s2d	1/1	Running	0	54s
dummy-8984df79-zpvn	0/1	RunContainerError	2 (4s ago)	20s

```
$ k describe pod dummy-8984df79-zpvnm
```

```
Name:          dummy-8984df79-zpvnm
```

```
Namespace:     default
```

```
Priority:       0
```

```
Service Account: default
```

```
Node:          beestinger-control-plane/172.18.0.3
```

```
...
```

```
Events:
```

Type	Reason	Age	From	Message
----	-----	----	----	-----

```
...
```

Warning	Failed	8s (x4 over 49s)	kubelet	Error: failed to create containerd task: failed to create shim task: OCI runtime create failed: runc create failed: unable to start container process: error during container init: error running hook #1: fork/exec /tmp/beesting/beesting-hook: permission denied: unknown
---------	--------	------------------	---------	--

```
$ k describe pod dummy-8984df79-zpvnm
```

```
Name:          dummy-8984df79-zpvnm
```

```
Namespace:     default
```

```
Priority:       0
```

```
Service Account: default
```

```
Node:          beestinger-control-plane/172.18.0.3
```

```
...
```

```
Events:
```

Type	Reason	Age	From	Message
----	-----	----	----	-----

```
...
```

Warning	Failed	8s (x4 over 49s)	kubelet	Error: failed to create containerd task: failed to create shim task: OCI runtime create failed: runc create failed: unable to start container process: error during container init: error running hook #1: fork/exec /tmp/beesting/beesting-hook: permission denied: unknown
---------	--------	------------------	---------	--

```
Error: failed to create containerd task: failed to create shim task:  
OCI runtime create failed: runc create failed: unable to start container  
process: error during container init: error running hook #1: fork/exec  
/tmp/beesting/beesting-hook: permission denied: unknown
```

```
$ root@beestinger-control-plane:/# mount
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev type tmpfs (rw,nosuid,size=65536k,mode=755,inode64)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=666)
sysfs on /sys type sysfs (ro,nosuid,nodev,noexec,relatime)
...
mqueue on /dev/mqueue type mqueue (rw,nosuid,nodev,noexec,relatime)
shm on /dev/shm type tmpfs (rw,nosuid,nodev,noexec,relatime,size=65536k,inode64)
tmpfs on /tmp type tmpfs (rw,nosuid,nodev,noexec,relatime,inode64)
/dev/vda1 on /var type ext4 (rw,relatime,discard,errors=remount-ro)
devpts on /dev/console type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=666)
tmpfs on /run/lock type tmpfs (rw,nosuid,nodev,noexec,relatime,size=5120k,inode64)
...
```

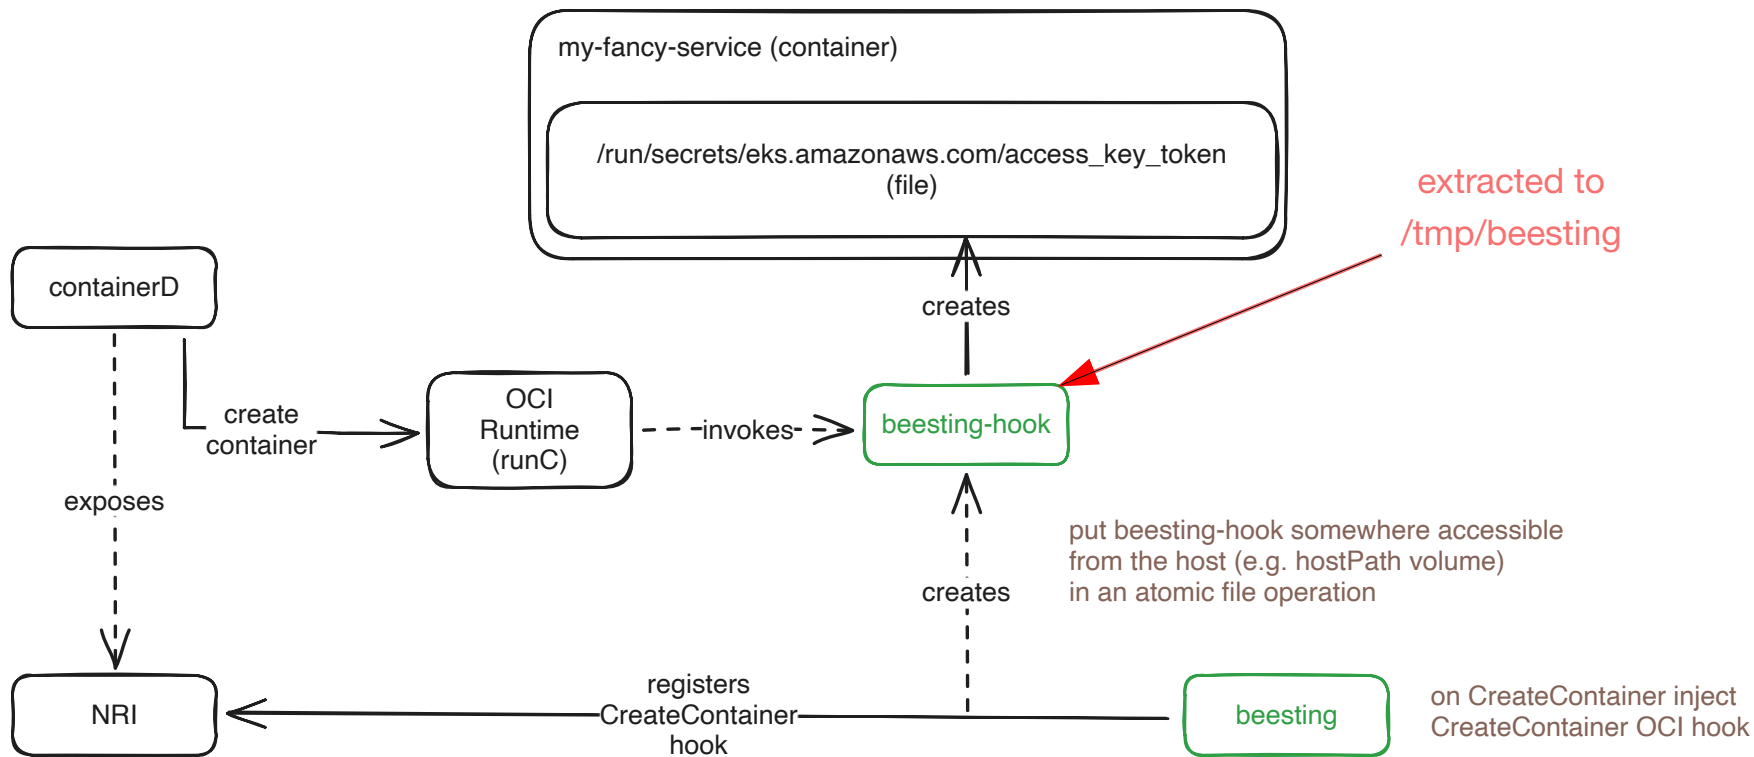
```
$ root@beestinger-control-plane:/# mount
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev type tmpfs (rw,nosuid,size=65536k,mode=755,inode64)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=666)
sysfs on /sys type sysfs (ro,nosuid,nodev,noexec,relatime)
...
mqueue on /dev/mqueue type mqueue (rw,nosuid,nodev,noexec,relatime)
shm on /dev/shm type tmpfs (rw,nosuid,nodev,noexec,relatime,size=65536k,inode64)
tmpfs on /tmp type tmpfs (rw,nosuid,nodev,noexec,relatime,inode64)
/dev/vda1 on /var type ext4 (rw,relatime,discard,errors=remount-ro)
devpts on /dev/console type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=666)
tmpfs on /run/lock type tmpfs (rw,nosuid,nodev,noexec,relatime,size=5120k,inode64)
...
```

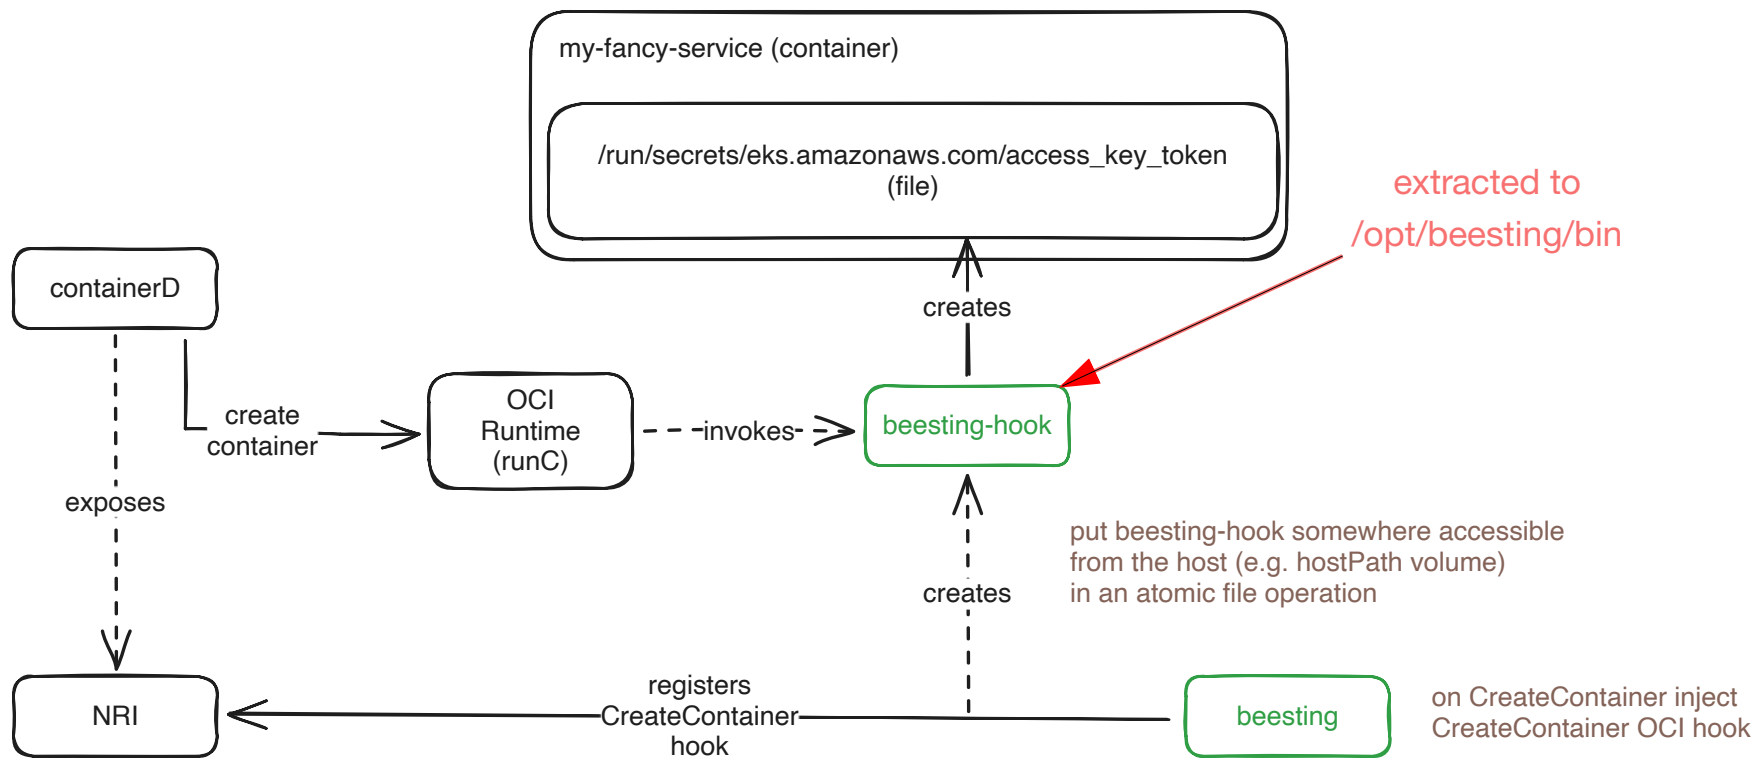

Code Blame 361 lines (331 loc) · 13.5 KB · 🔍

Raw 📄 ⬇️ ✎ ⌵

```
40 type Installer struct {
41 func newInstaller(cfg *Config.InstallConfig, isReady *atomic.Value) *Installer {
42     return &Installer{
43         cfg:          cfg,
44         kubeconfigFilepath: filepath.Join(cfg.CNIAgentRunDir, constants.CNIPluginKubeconfName),
45         isReady:      isReady,
46     }
47 }
48
49 func (in *Installer) installAll(ctx context.Context) (sets.String, error) {
50     // Install binaries
51     // Currently we _always_ do this, since the binaries do not live in a shared location
52     // and we harm no one by doing so.
53     copiedFiles, err := copyBinaries(in.cfg.CNIBinSourceDir, in.cfg.CNIBinTargetDirs)
54     if err != nil {
55         cniInstalls.With(resultLabel.Value(resultCopyBinariesFailure)).Increment()
56         return copiedFiles, fmt.Errorf("copy binaries: %v", err)
57     }
58
59     // Write kubeconfig with our current service account token as the contents, to the Istio agent rundir.
60     // We do not write this to the common/shared CNI config dir, because it's not CNI config, we do not
61     // need to watch it, and writing non-shared stuff to that location creates churn for other node agents.
62     // Only our plugin consumes this kubeconfig, and it resides in our owned rundir on the host node,
63     // so we are good to simply write it out if our watched svcacct token changes.
64     if err := writeKubeConfigFile(in.cfg); err != nil {
65         cniInstalls.With(resultLabel.Value(resultCreateKubeConfigFailure)).Increment()
66         return copiedFiles, fmt.Errorf("write kubeconfig: %v", err)
67     }
68
69     // Install CNI netdir config (if needed) - we write/update this in the shared node CNI netdir,
70     // which may be watched by other CNIs, and so we don't want to trigger writes to this file
71     // unless it's missing or the contents are not what we expect.
72     if err := checkValidCNIConfig(in.cfg, in.cniConfigFilepath); err != nil {
73         installLog.Infof("configuration requires updates, (re)writing CNI config file at %q: %v", in.cniConfigFilepath, err)
74     }
75 }
```

/opt/cni/bin





```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-q9s2d	1/1	Running	0	30s
dummy-8984df79-zpvn	1/1	Running	0	52s

```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-q9s2d	1/1	Running	0	30s
dummy-8984df79-zpvn	1/1	Running	0	52s

```
$ k exec deploy/dummy -it -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root      root          4.0K Jan  6 15:11 .  
drwxr-xr-x    4 root      root          4.0K Jan  6 15:11 ..  
-r--r--r--    1 root      root           16 Jan  6 15:11 access_key_token
```

Too much
complexity

Next try

Replace hook with bind mount

host-1

file is accessible from both locations

root (mount namespace)

/tmp/beesting/container1-token

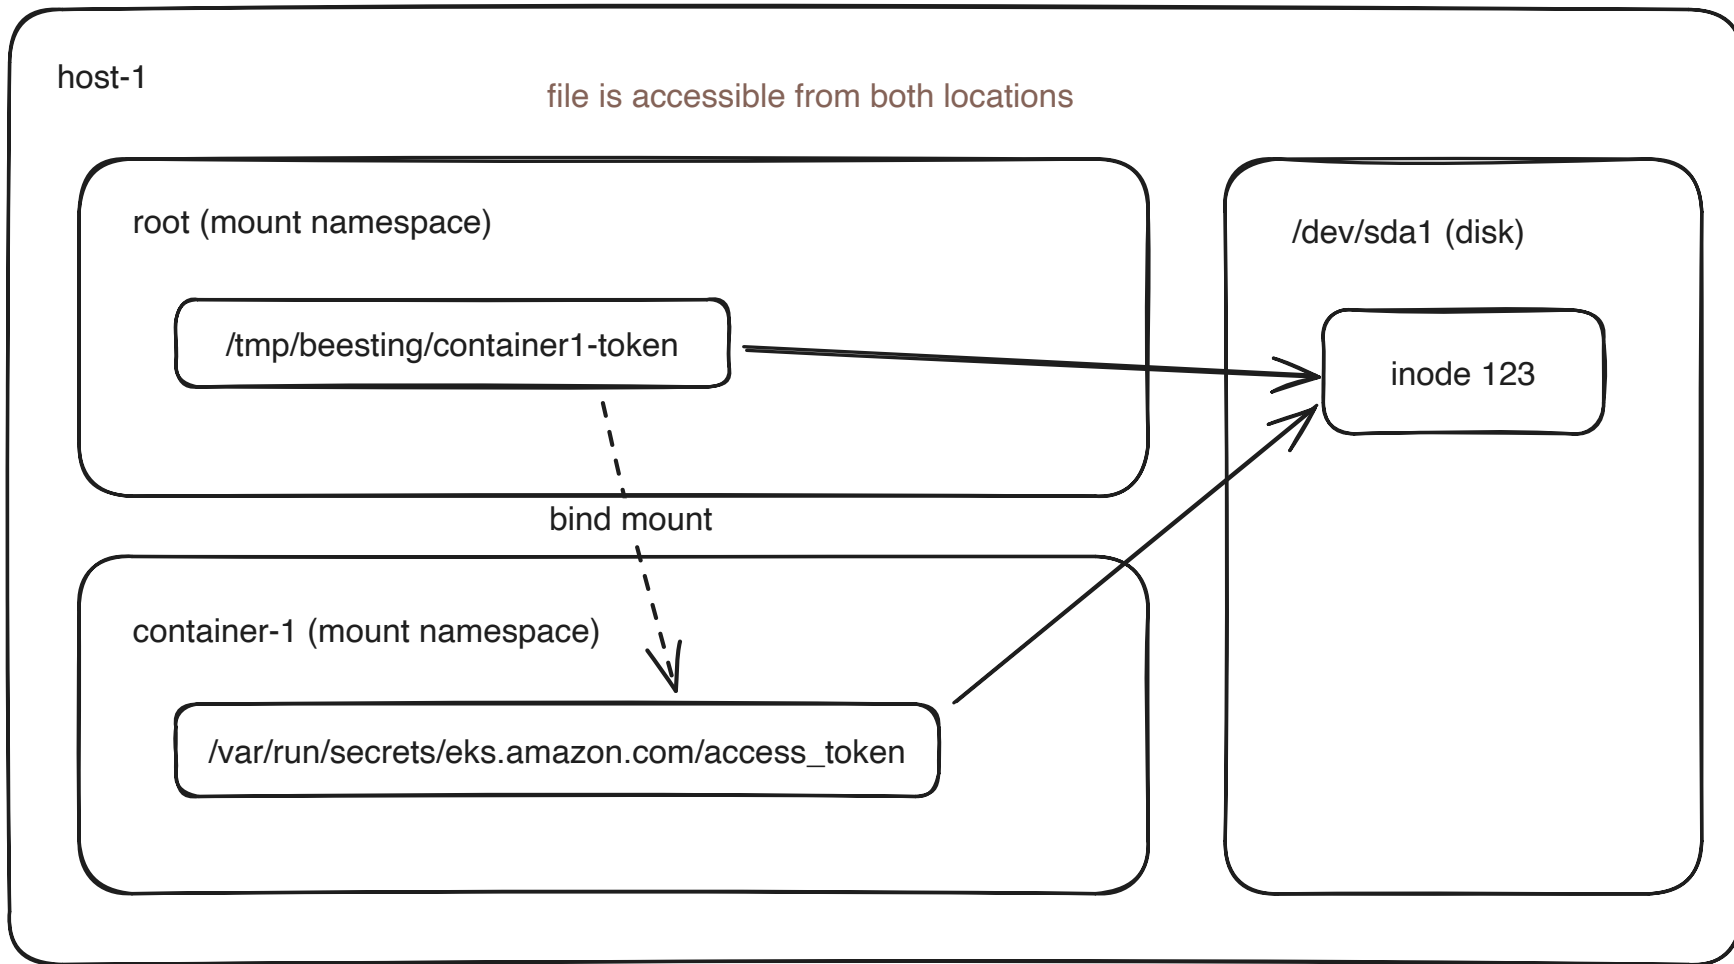
bind mount

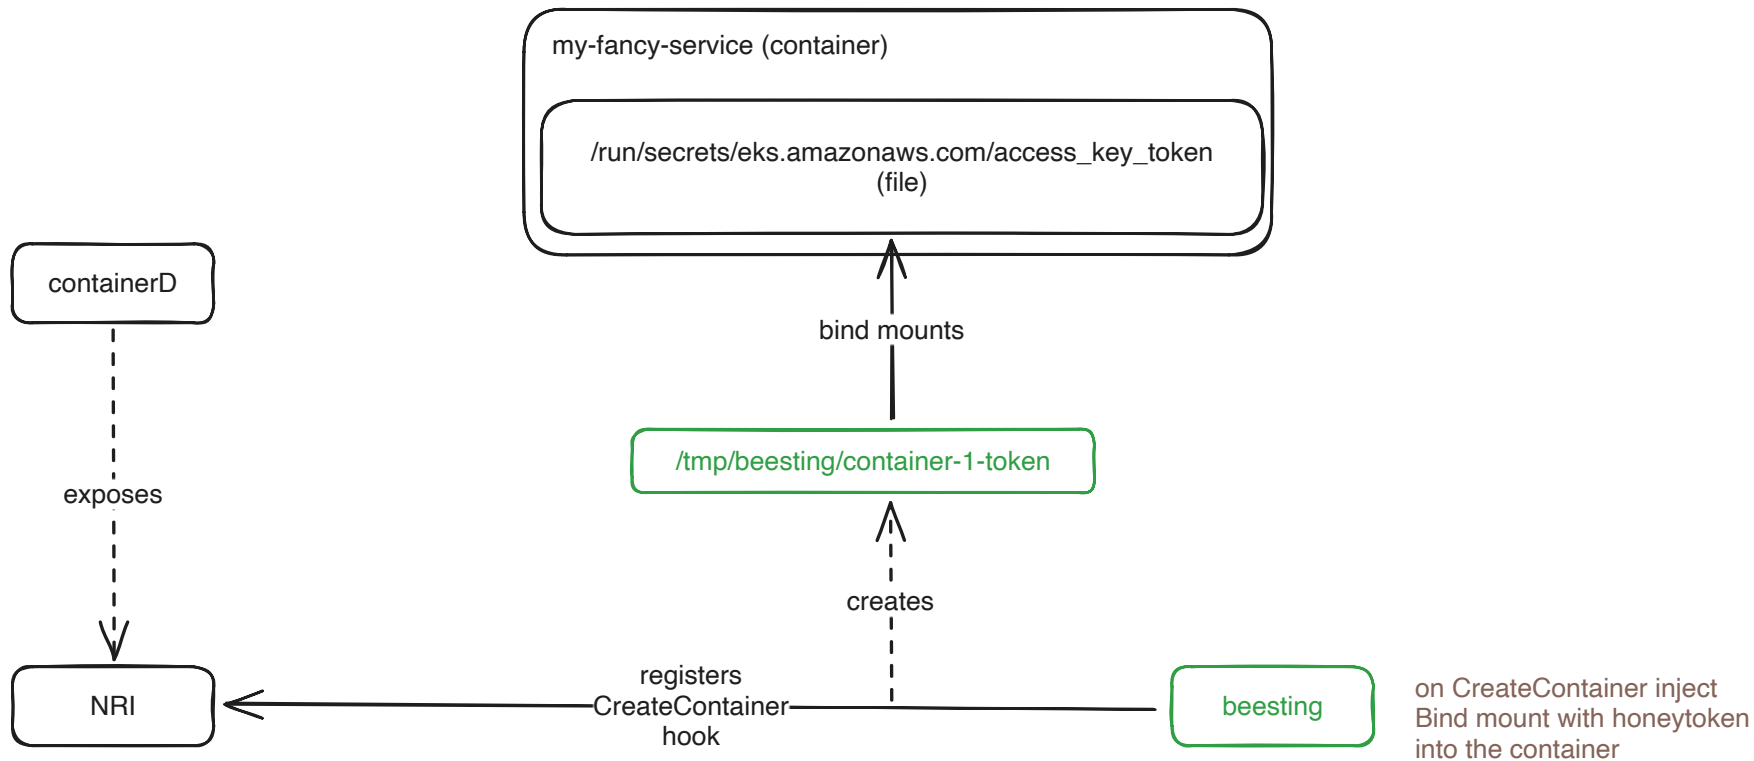
container-1 (mount namespace)

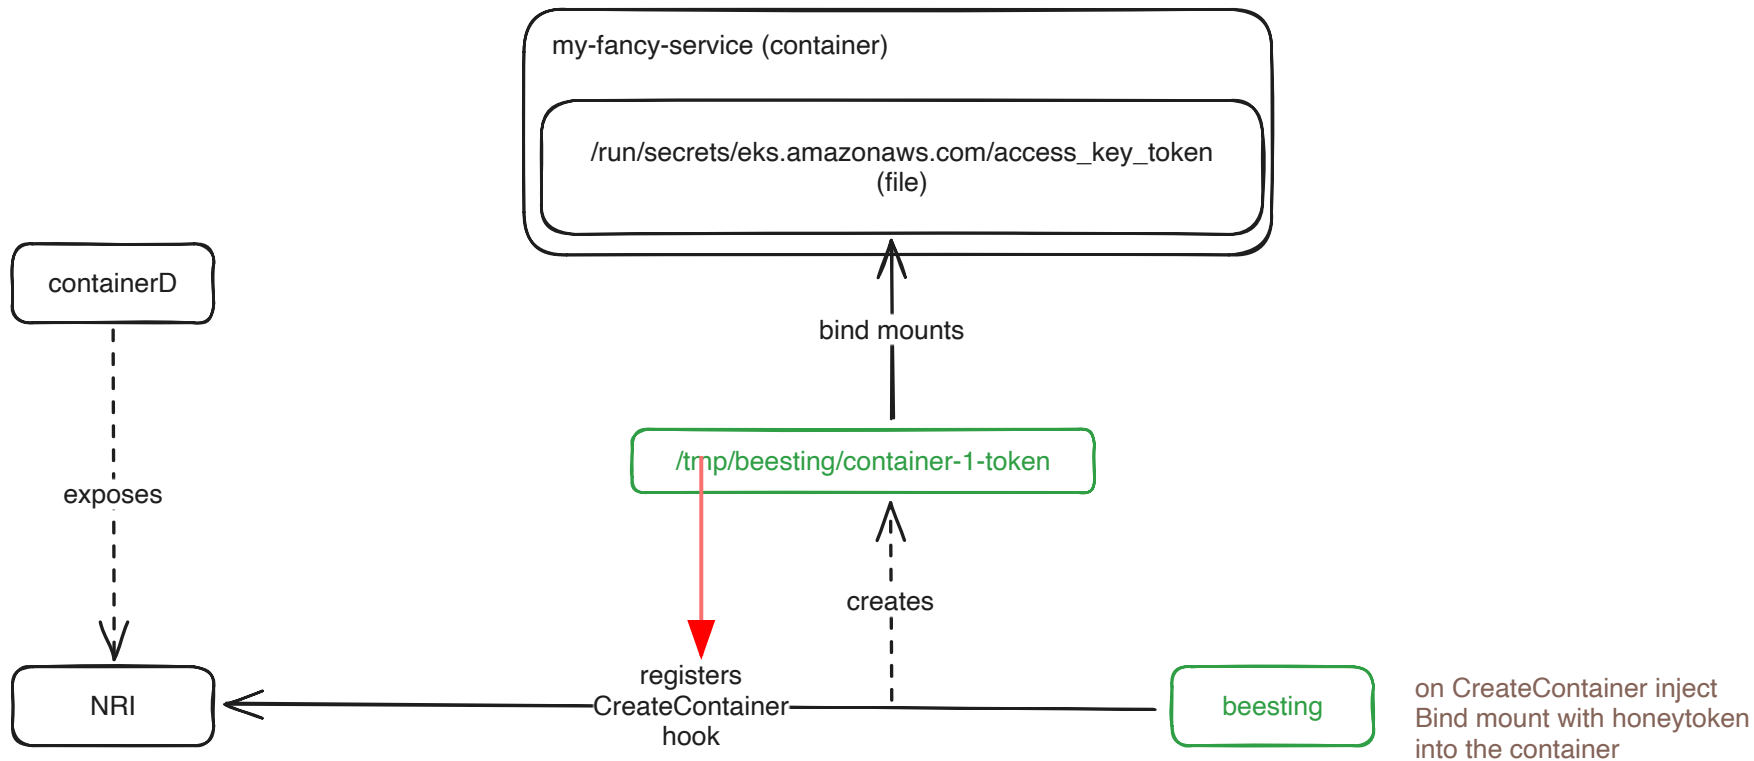
/var/run/secrets/eks.amazon.com/access_token

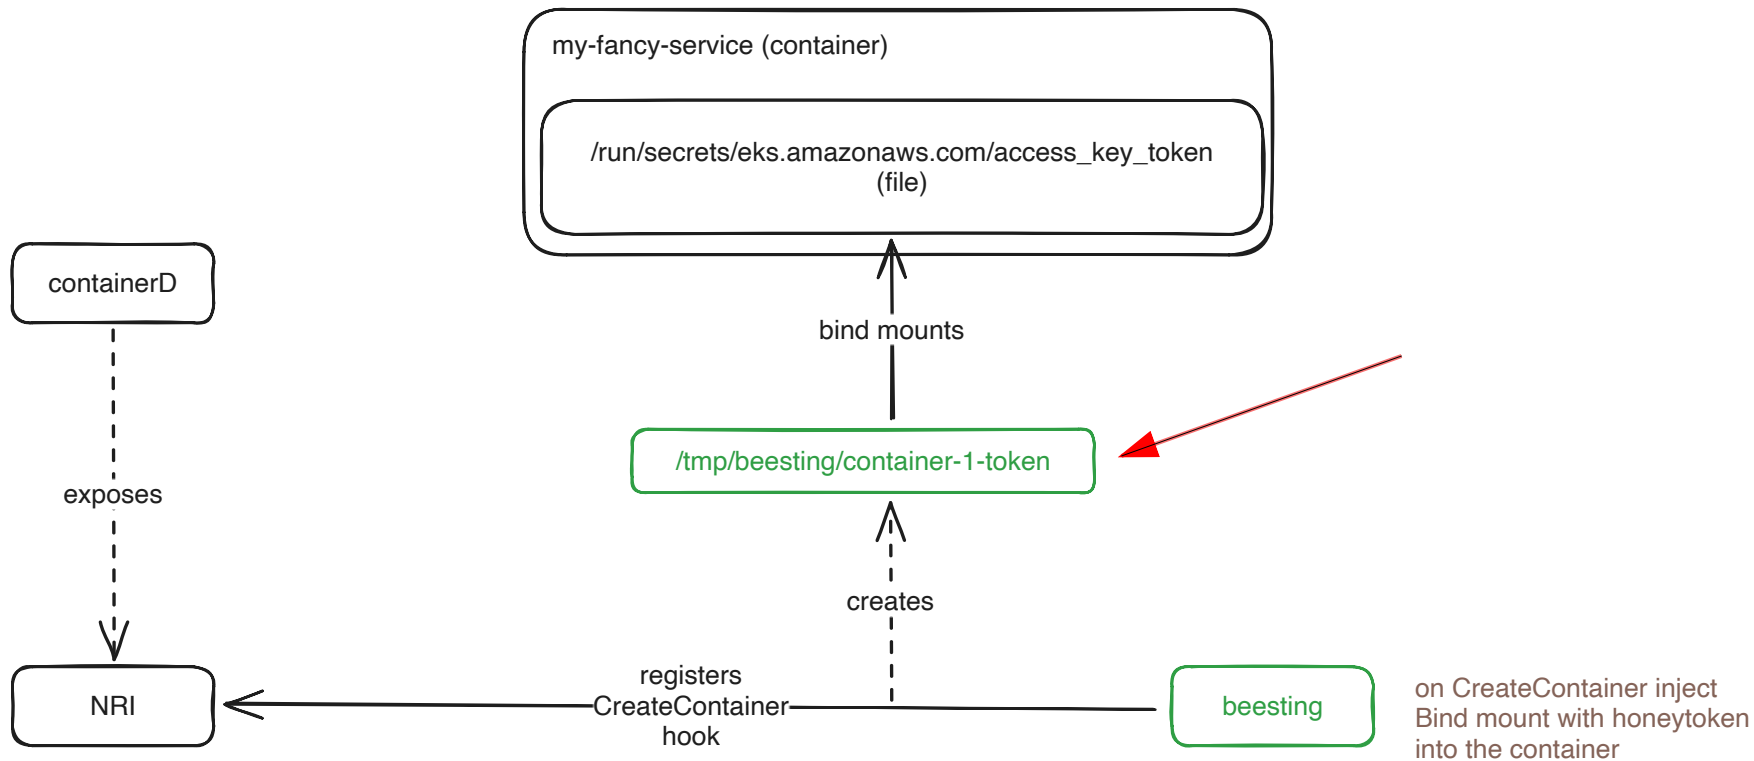
/dev/sda1 (disk)

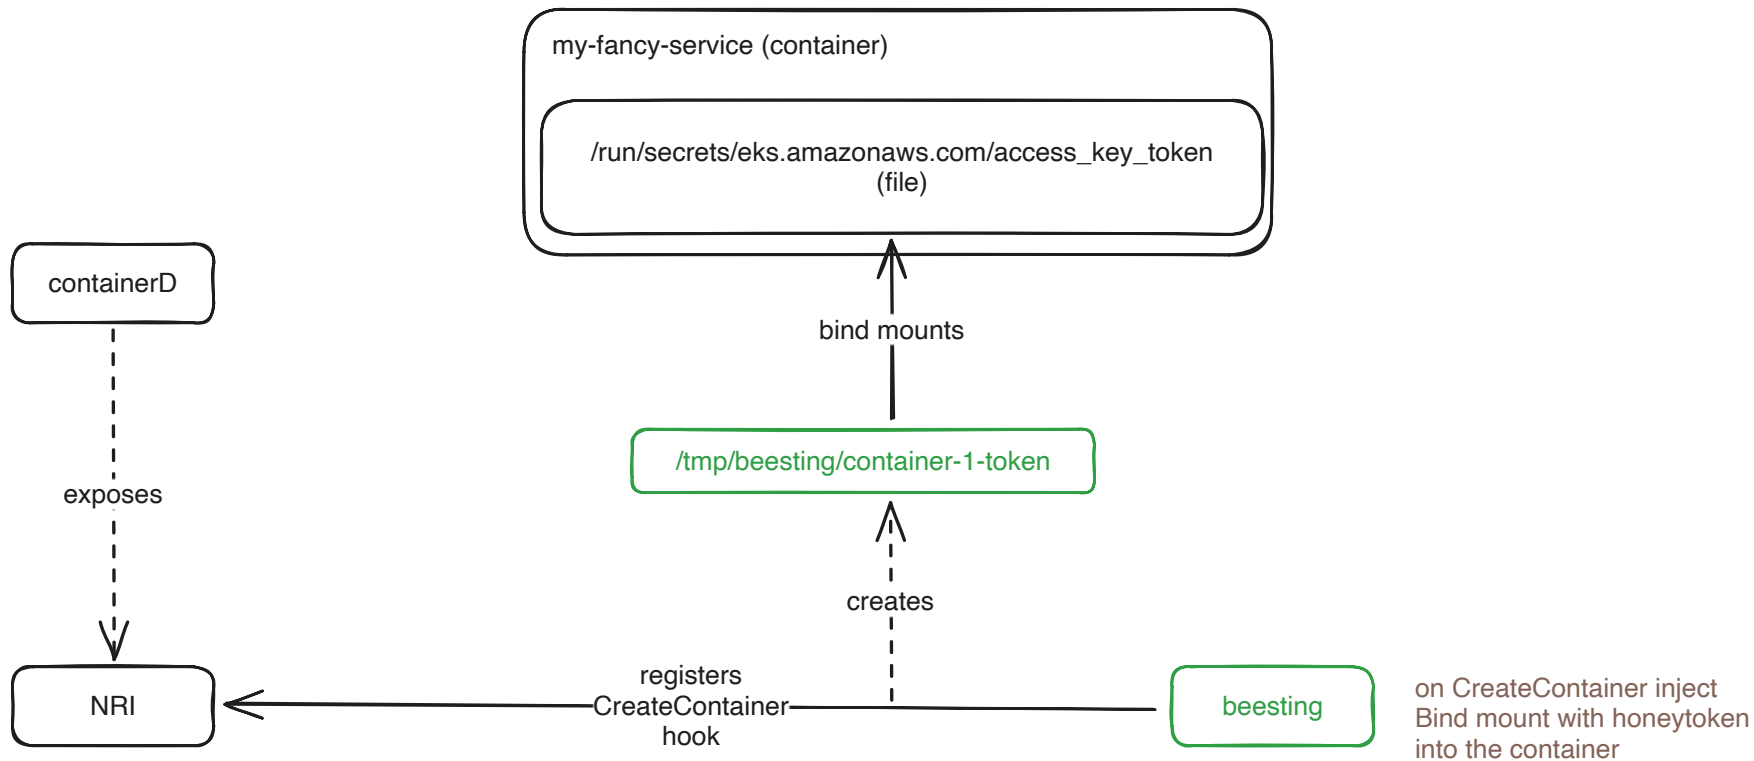
inode 123











```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-x2g8w	1/1	Running	0	30s
dummy-8984df79-isz1s	1/1	Running	0	52s

```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.024708ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-x2g8w	1/1	Running	0	30s
dummy-8984df79-isz1s	1/1	Running	0	52s

```
$ k exec deploy/dummy -it -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root      root          4.0K Jan  6 16:51 .  
drwxr-xr-x    4 root      root          4.0K Jan  6 16:51 ..  
-r--r--r--    1 root      root           16 Jan  6 16:51 access_key_token
```



```
$ k exec deploy/dummy -it -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root    root      4.0K Jan  6 16:51 .  
drwxr-xr-x    4 root    root      4.0K Jan  6 16:51 ..  
-r--r--r--    1 root    root      16 Jan  6 16:51 access_key_token
```





How do we detect File access?





eBPF

Traditional Kernel Development

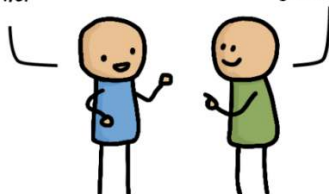
Application Developer:

I want this new feature to observe my app



Hey kernel developer! Please add this new feature to the Linux kernel

OK! Just give me a year to convince the entire community that this is good for everyone.

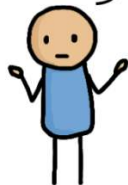


1 year later...

I'm done. The upstream kernel now supports this.



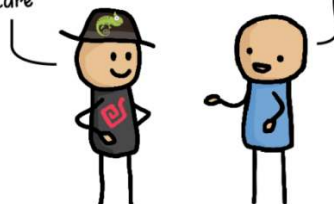
But I need this in my Linux distro



5 years later...

Good news. Our Linux distribution now ships a kernel with your required feature

OK but my requirements have changed since...



eBPF Revolution

Application Developer:

I want this new feature to observe my app



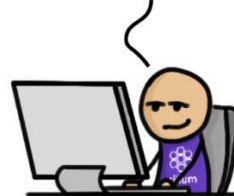
eBPF Developer:

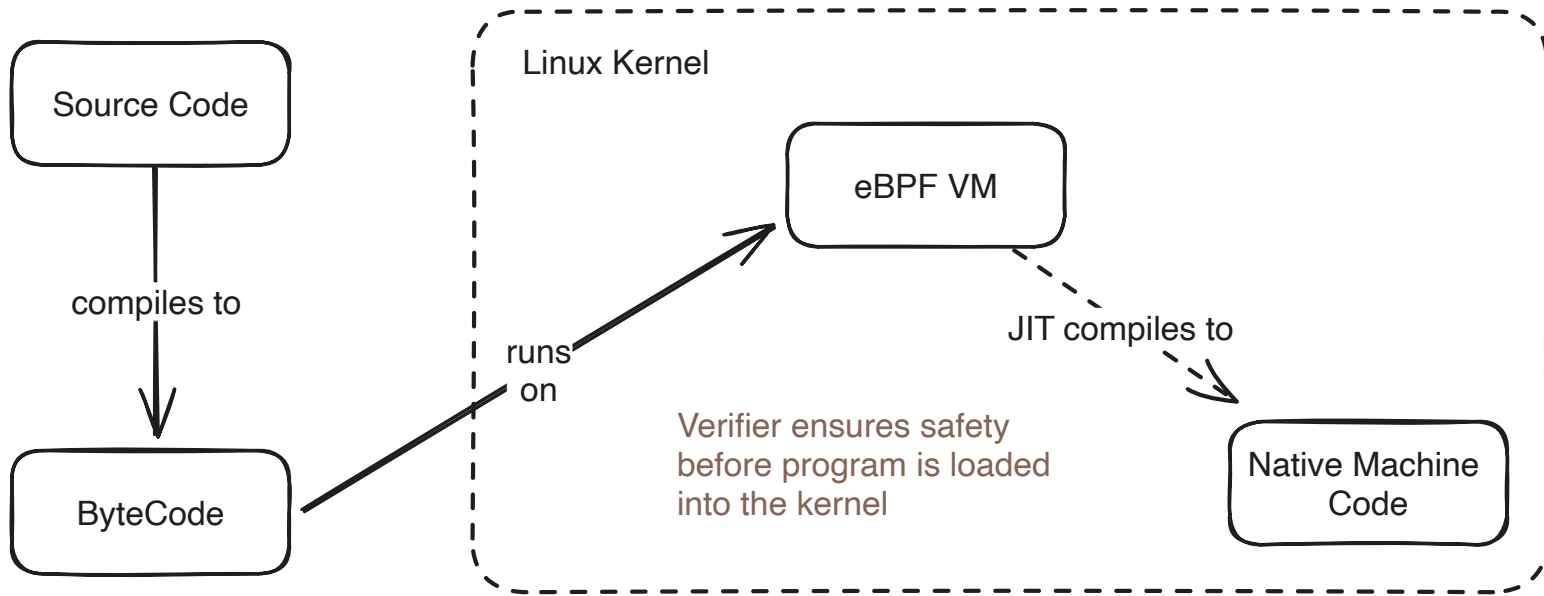
OK! The kernel can't do this so let me quickly solve this with eBPF.



A couple of days later...

Here is a release of our eBPF project that has this feature now. BTW, you don't have to reboot your machine.



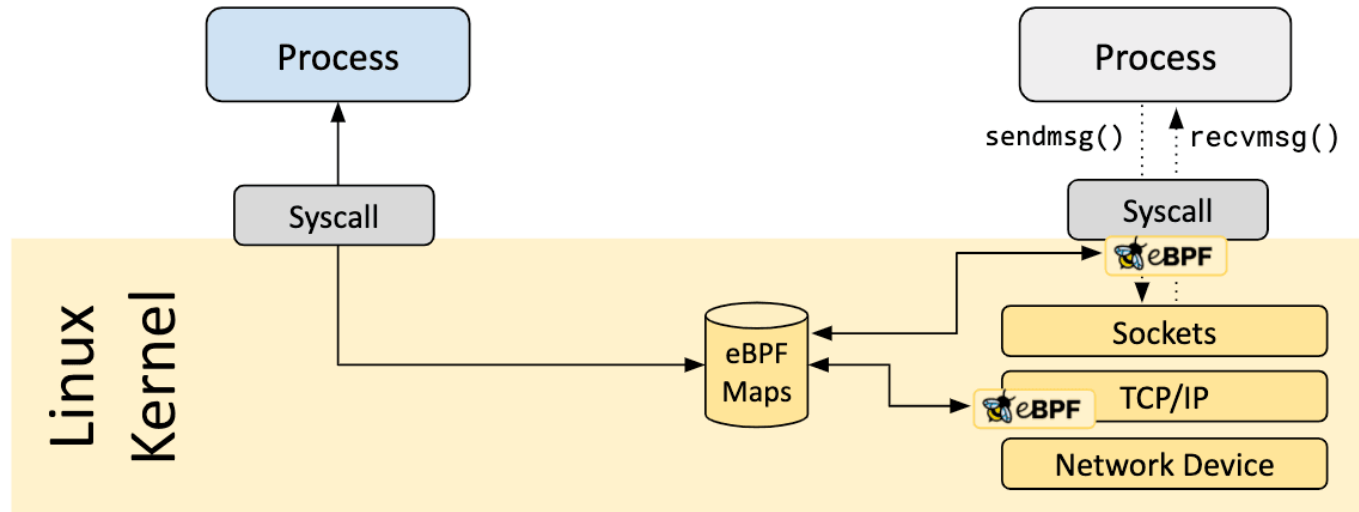


Unbounded Loops

Unbounded Loops





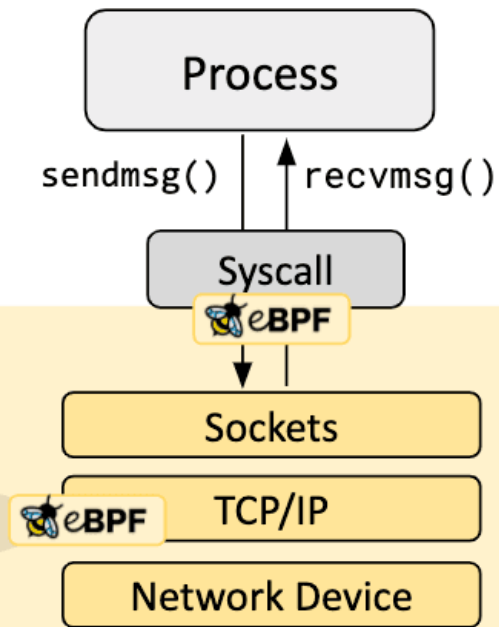


Map Types

- HashTable, Arrays
- LRU (Least Recently Used)
- Perf and Ring Buffer
- ...

Linux Kernel

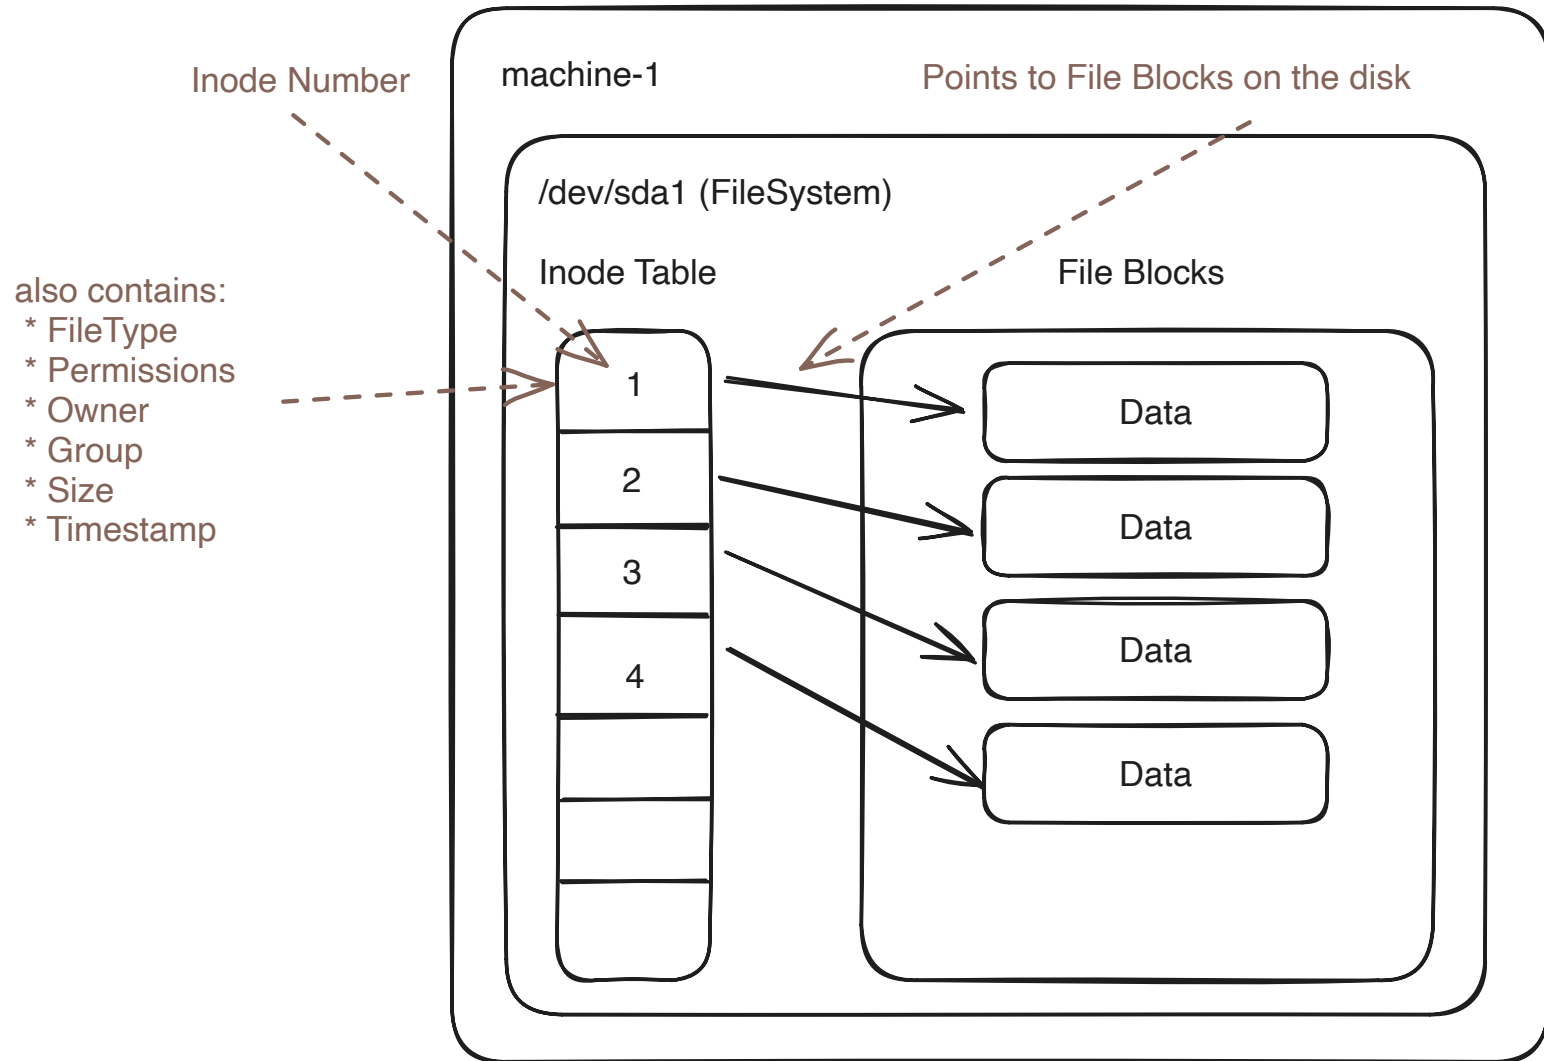
```
[...]  
num = bpf_get_prandom_u32();  
[...]
```



Helpers

- `bpf_get_current_pid_tgid`
- `bpf_map_lookup_elem`
- `bpf_map_delete_elem`
- ...

What is a file?



- open
- openat
- symlinks
- ...



Linux Security

Modules

Beesting doesn't use LSM directly

Userspace

my-fancy-service

open file
/tmp/token

Kernel

open
(syscall)

Beesting
eBPF Program

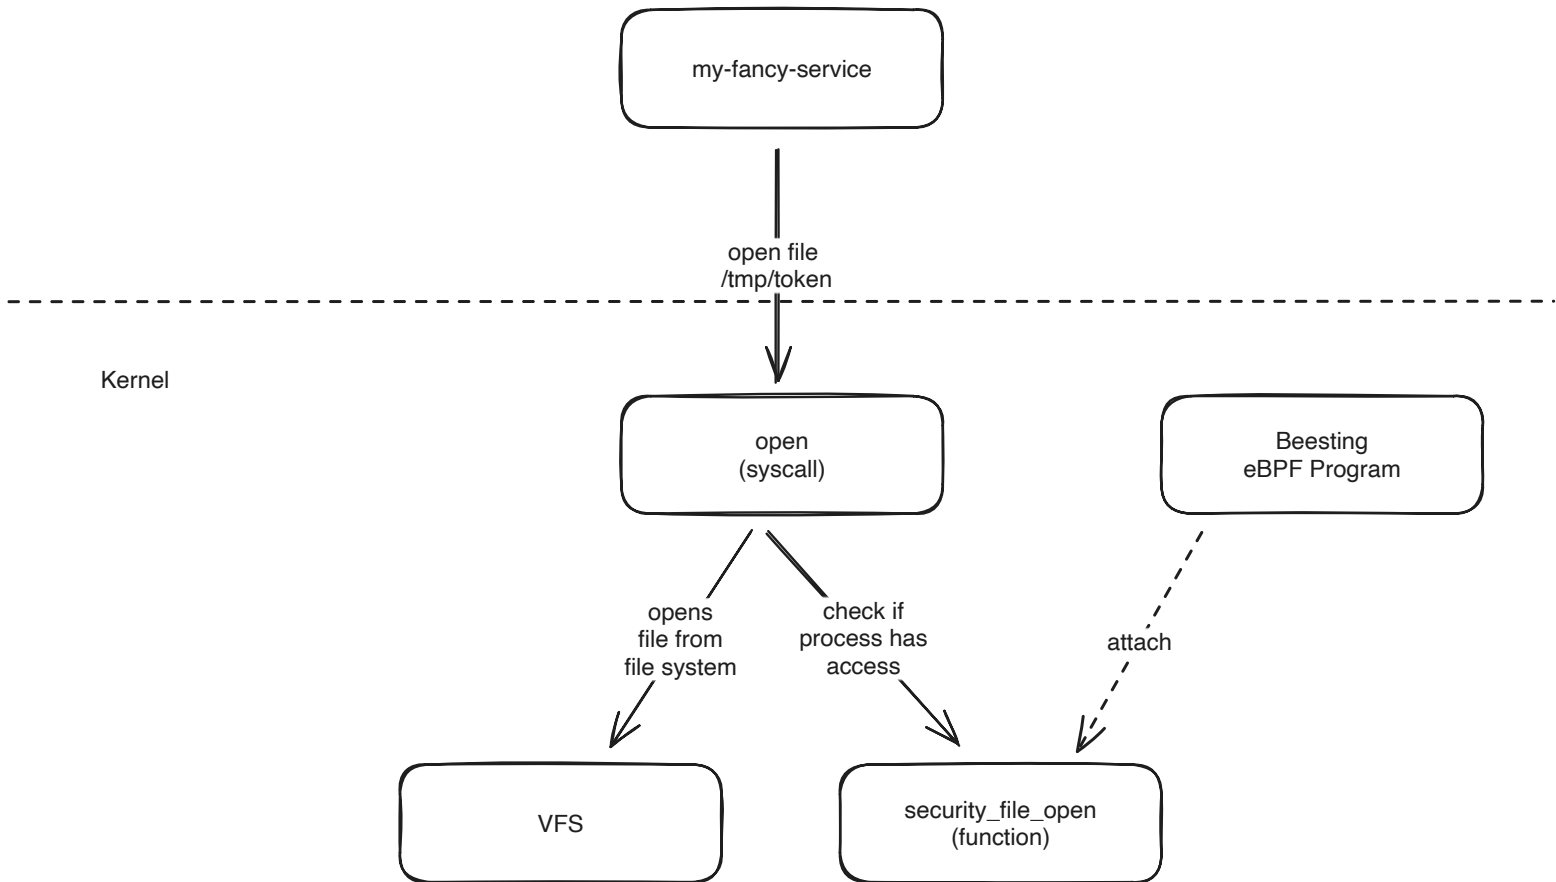
opens
file from
file system

check if
process has
access

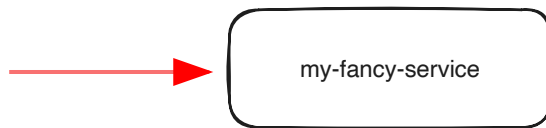
attach

VFS

security_file_open
(function)

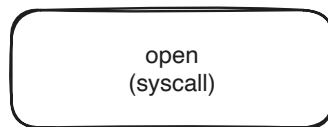


Userspace



open file
/tmp/token

Kernel

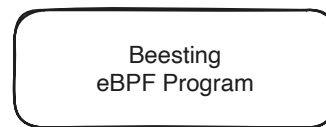


opens
file from
file system

check if
process has
access



security_file_open
(function)



attach

Userspace

my-fancy-service

open file
/tmp/token

Kernel

open
(syscall)

Beesting
eBPF Program

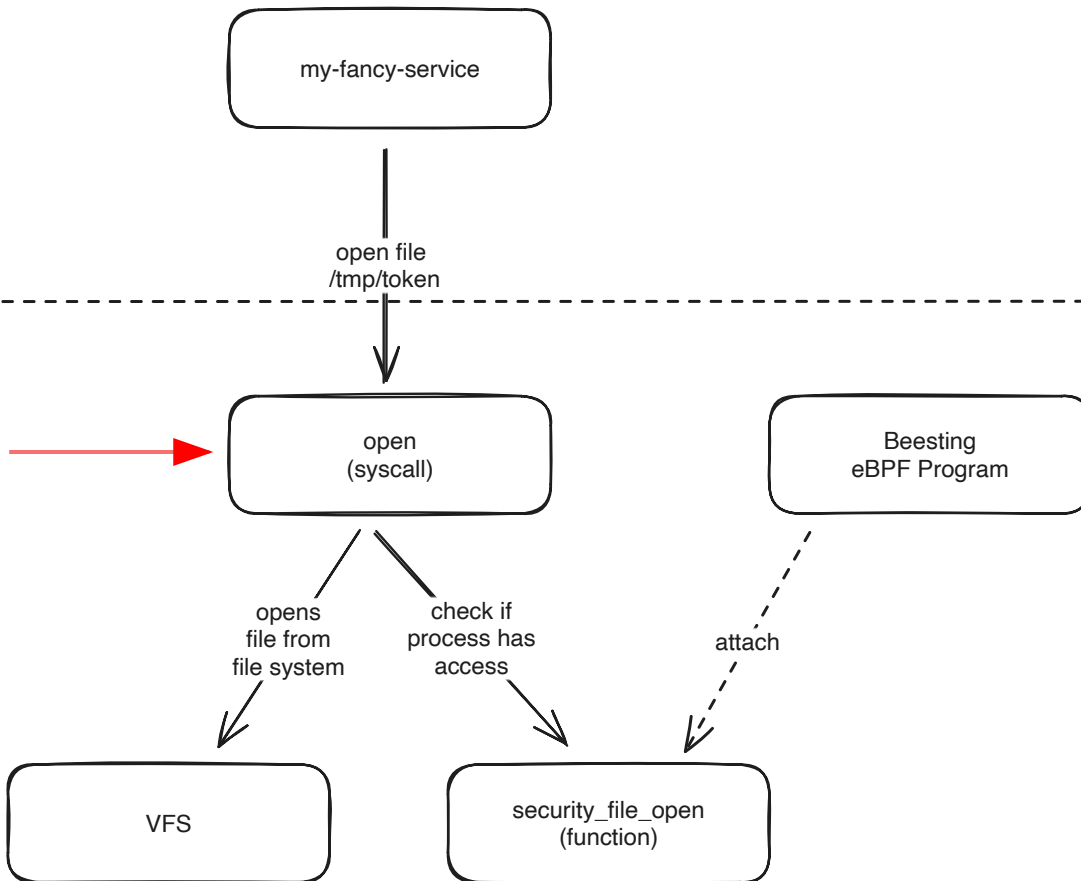
opens
file from
file system

check if
process has
access

attach

VFS

security_file_open
(function)



Userspace

my-fancy-service

open file
/tmp/token

Kernel

open
(syscall)

Beesting
eBPF Program

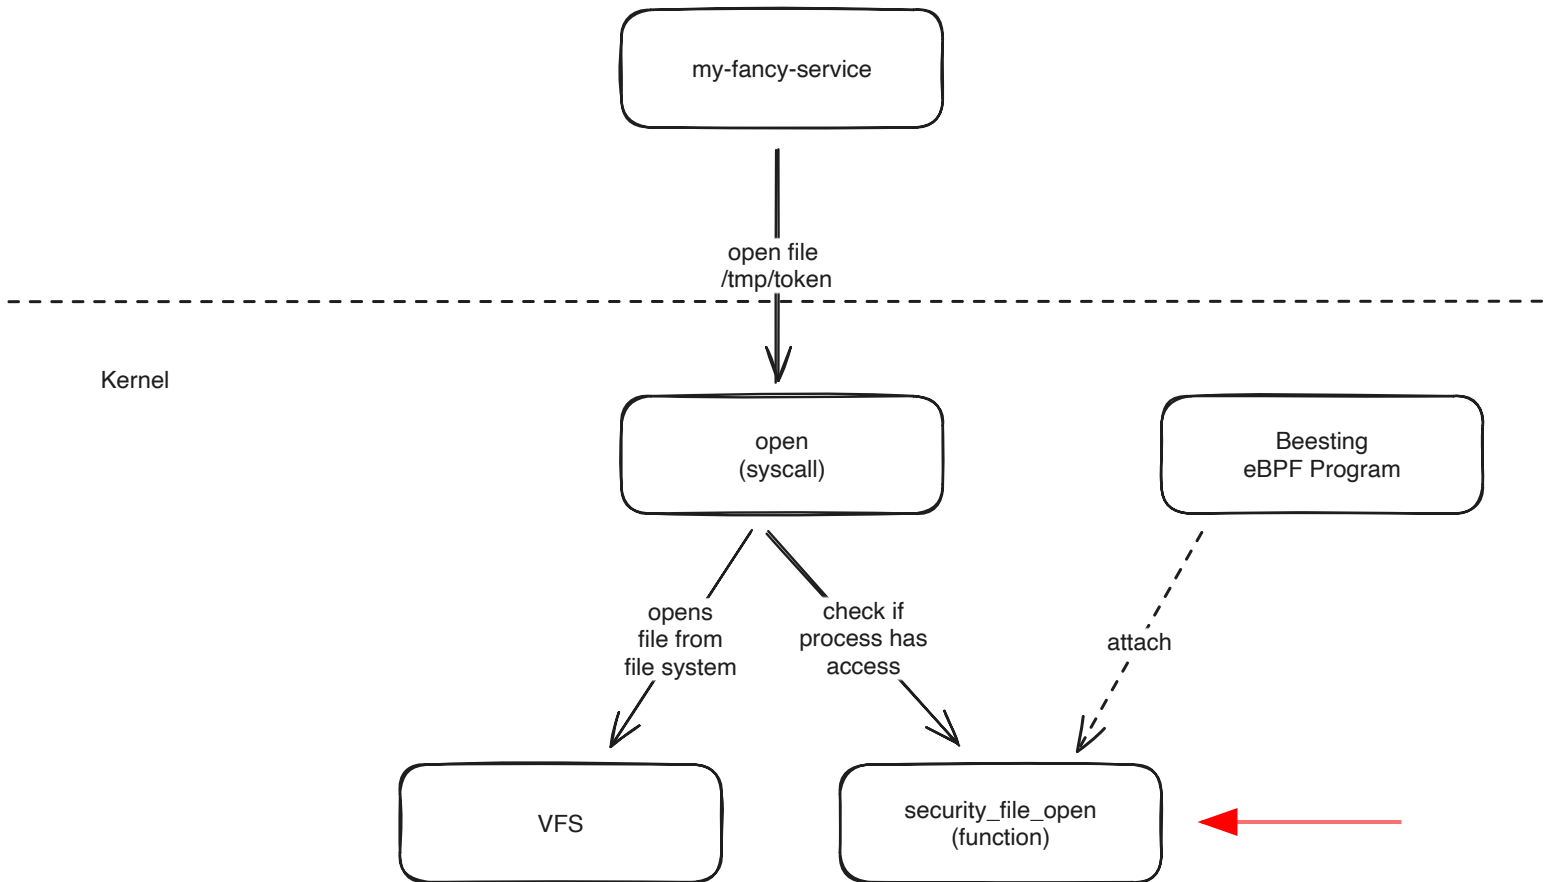
opens
file from
file system

check if
process has
access

attach

VFS

security_file_open
(function)



Userspace

my-fancy-service

open file
/tmp/token

Kernel

open
(syscall)

Beesting
eBPF Program

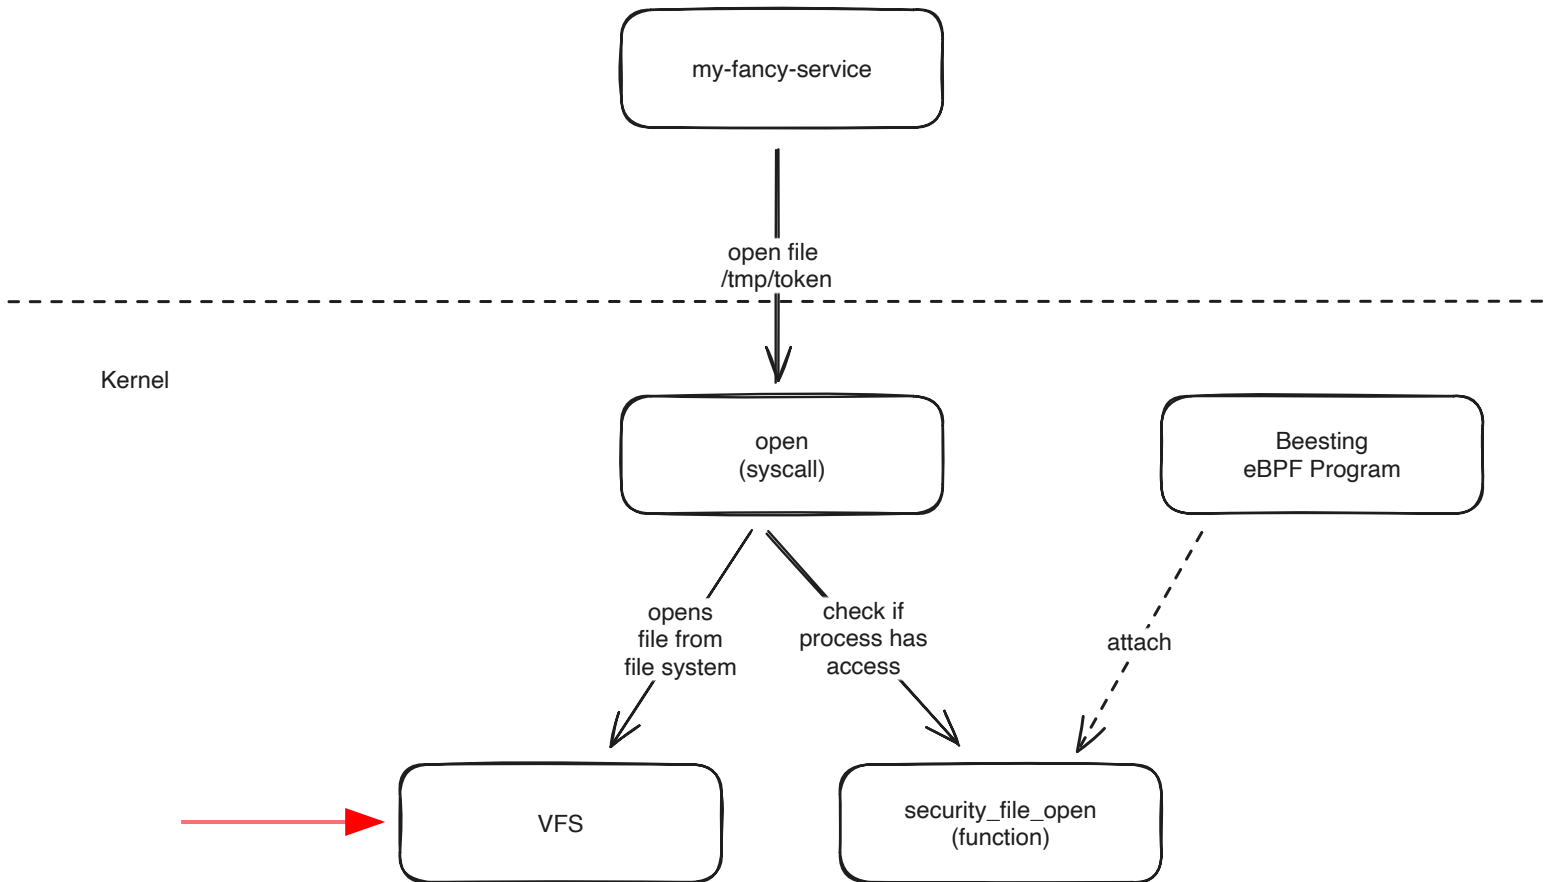
opens
file from
file system

check if
process has
access

attach

VFS

security_file_open
(function)



Userspace

my-fancy-service

open file
/tmp/token

Kernel

open
(syscall)

Beesting
eBPF Program

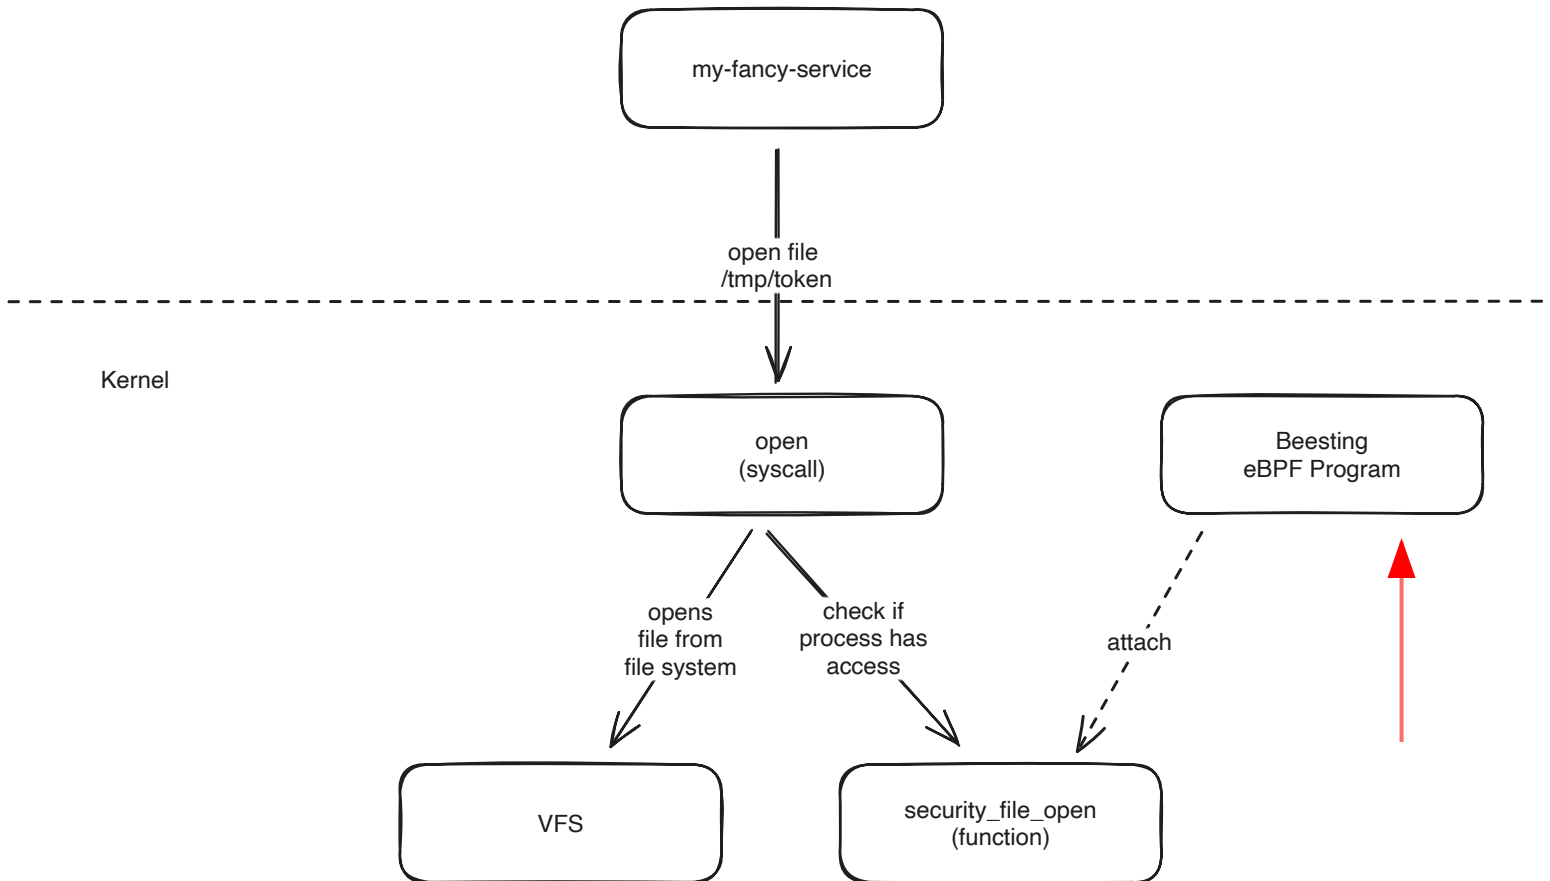
opens
file from
file system

check if
process has
access

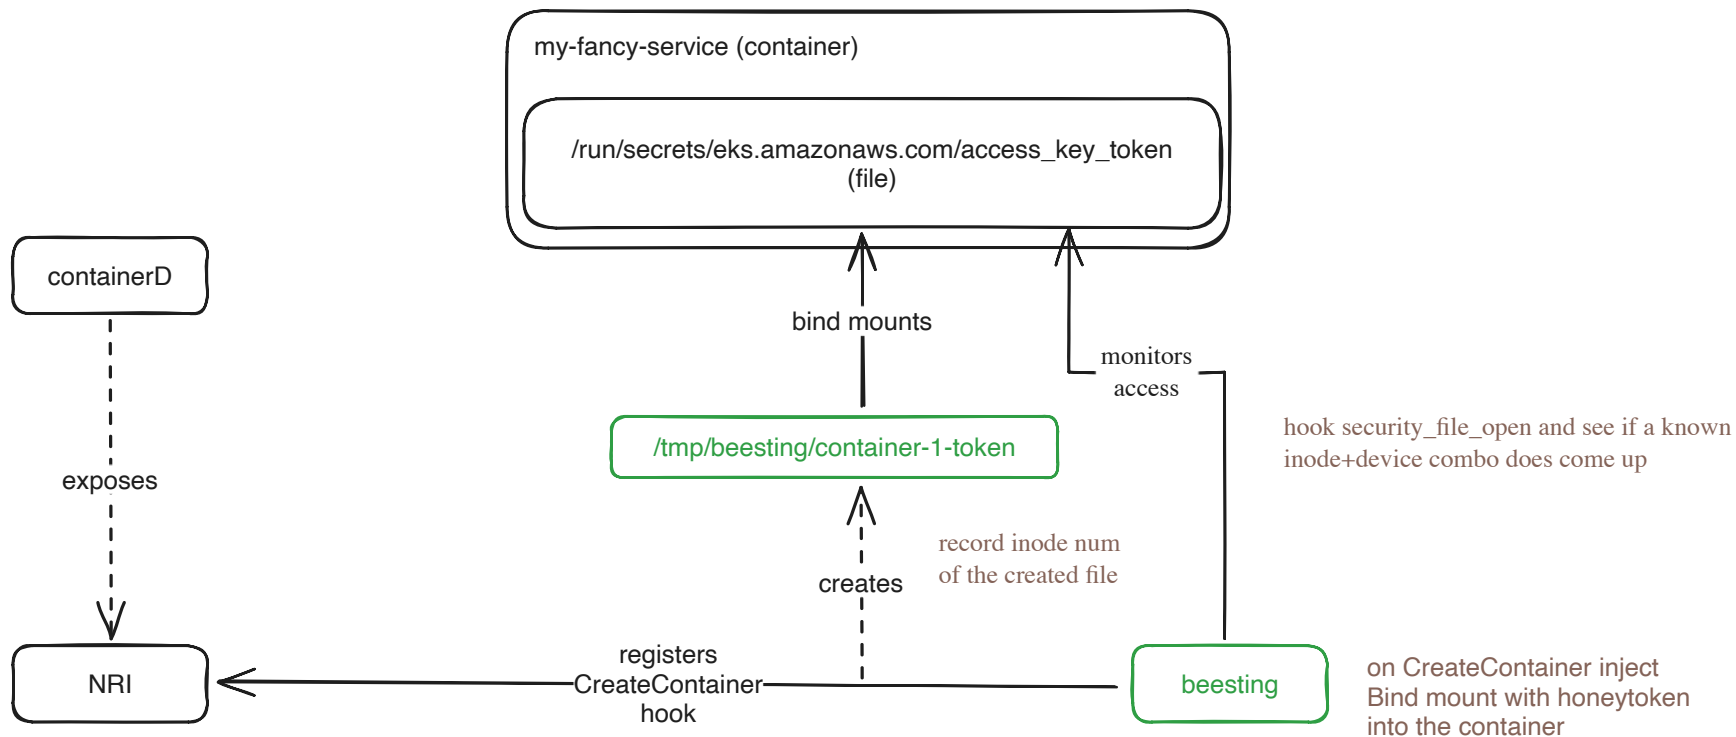
attach

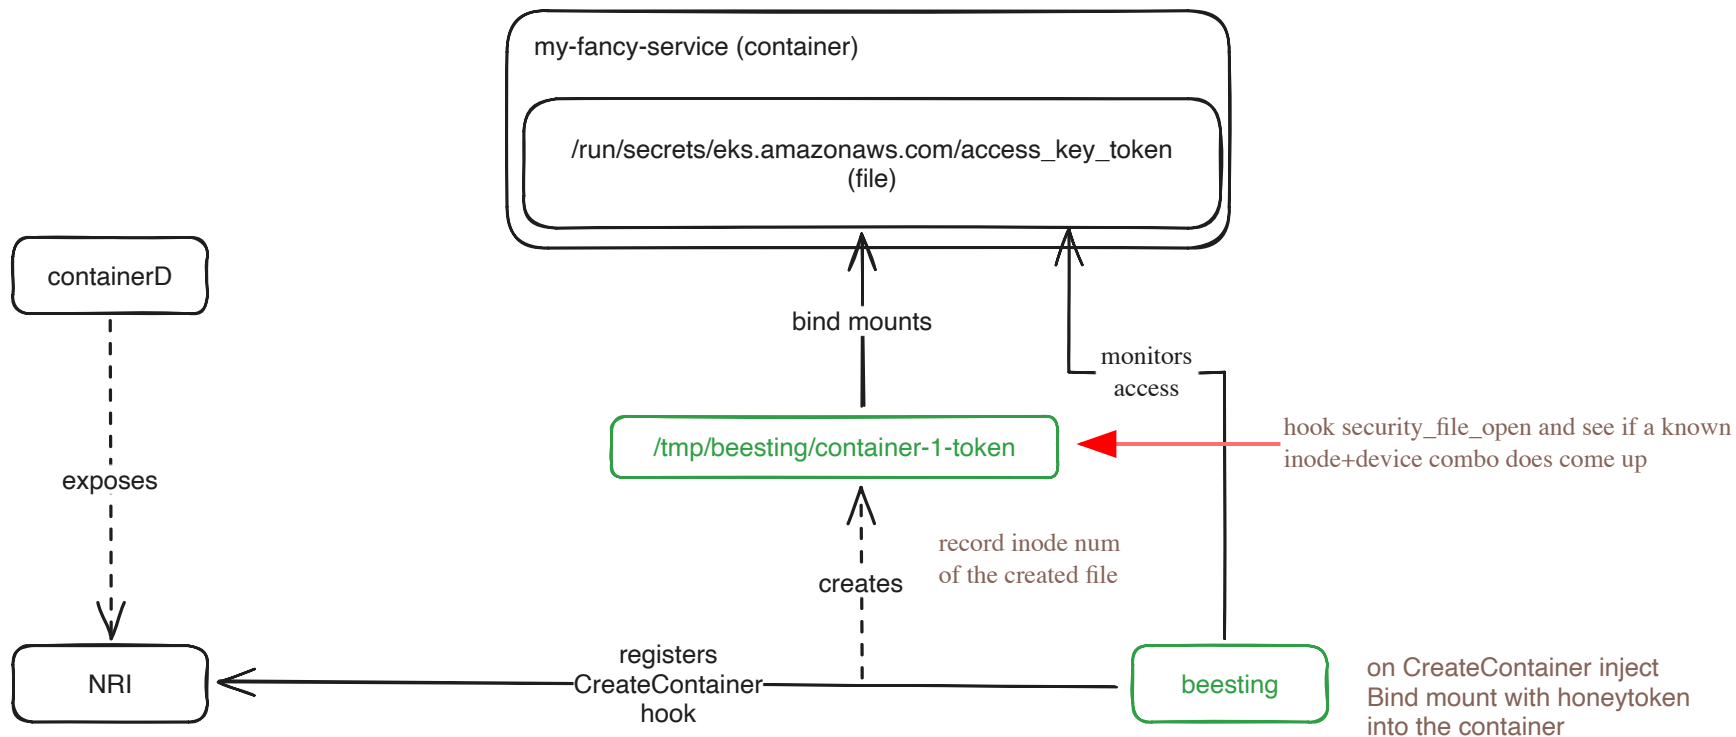
VFS

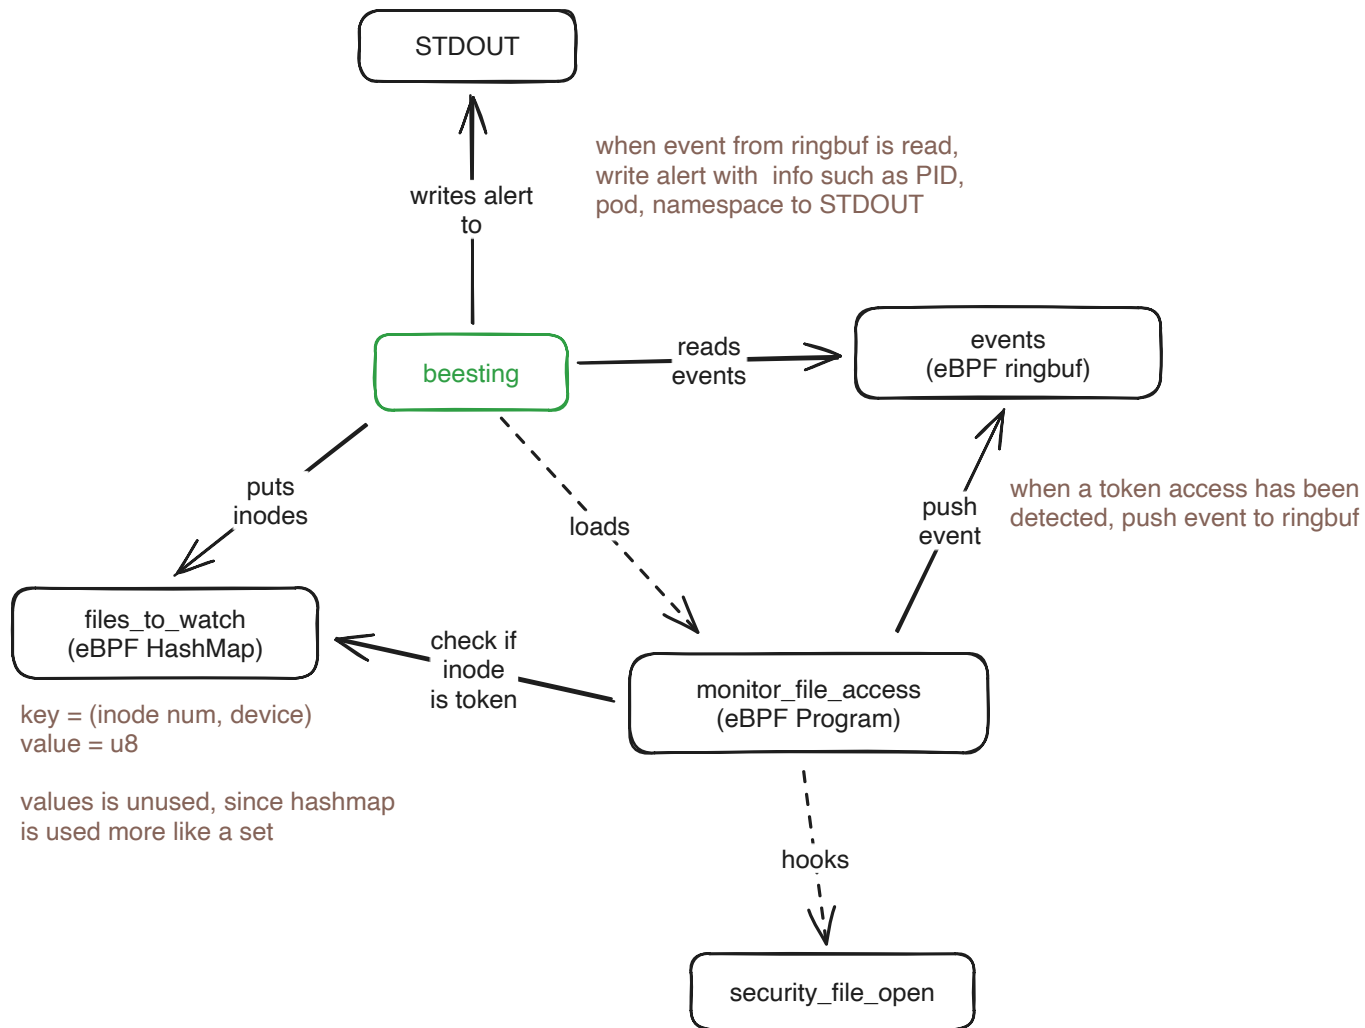
security_file_open
(function)

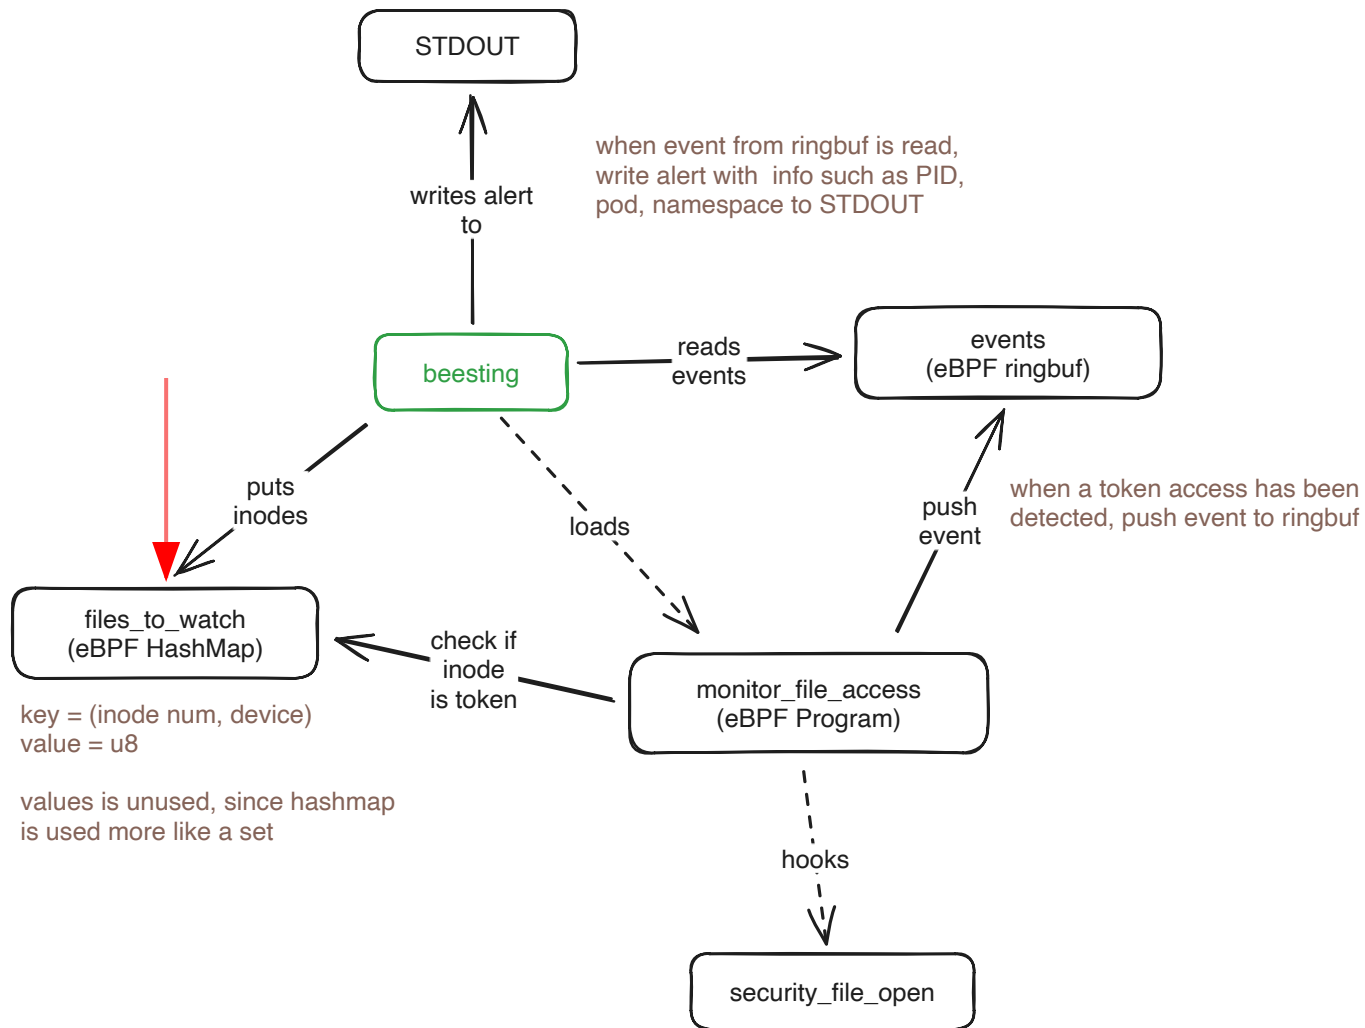


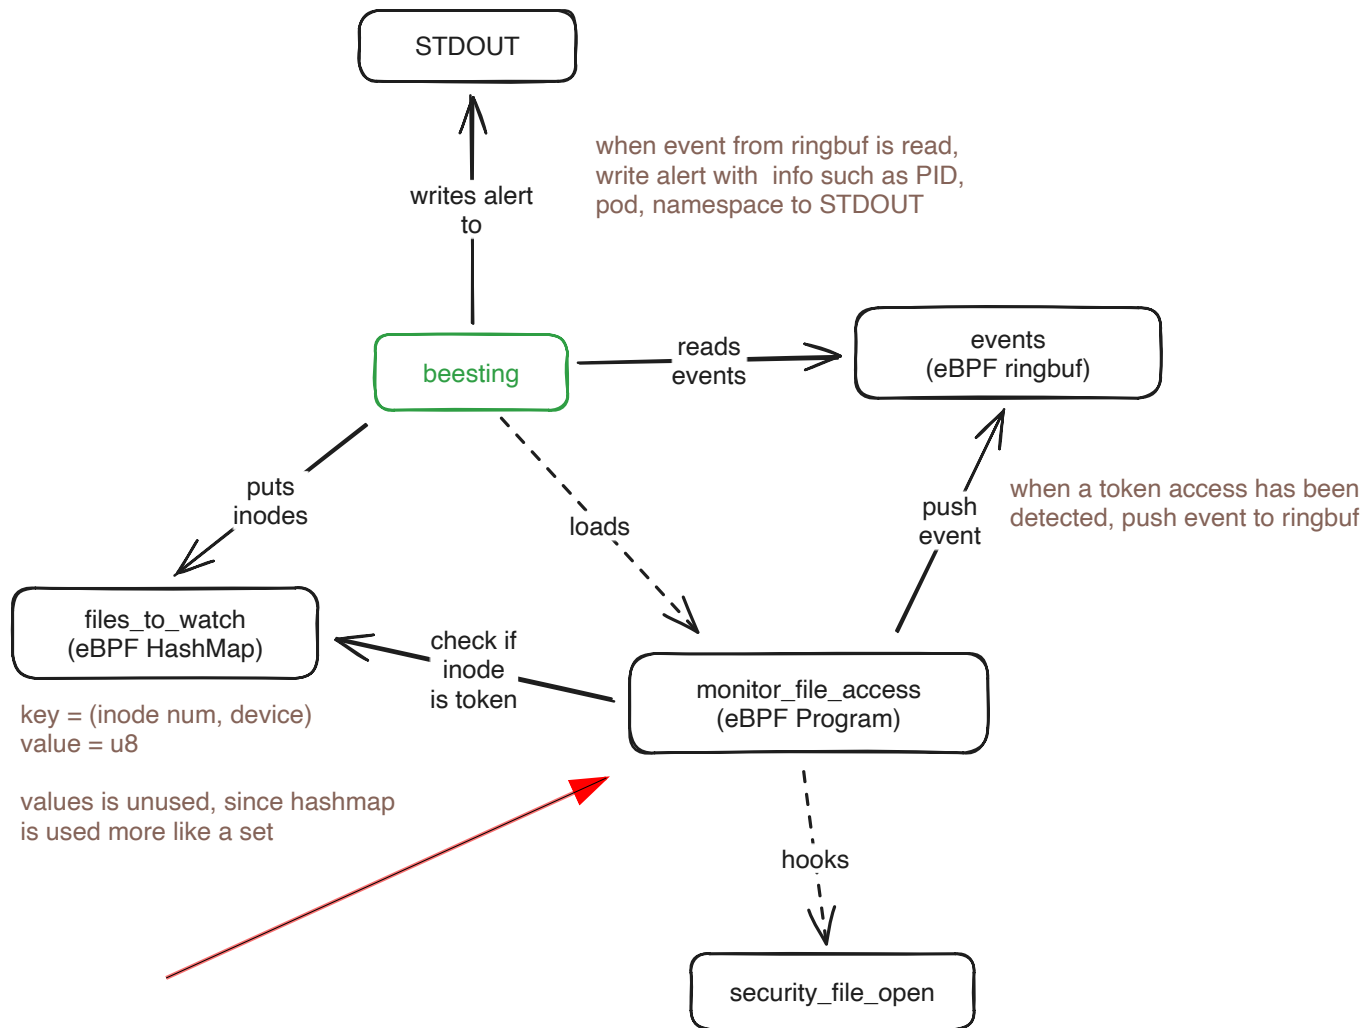
PoCv3 is based on PoCv2

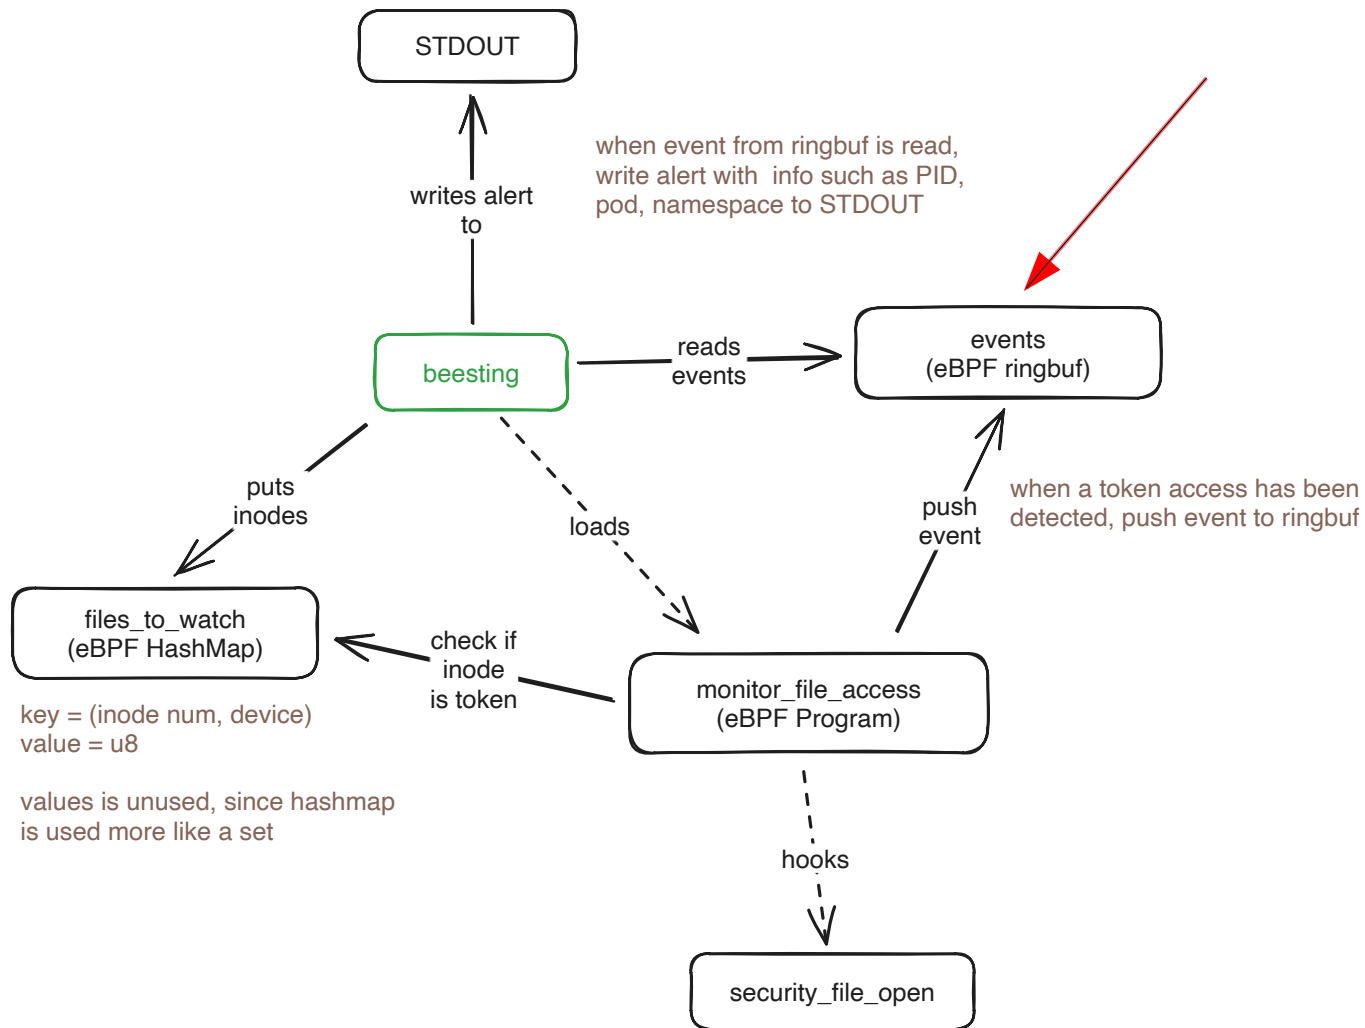


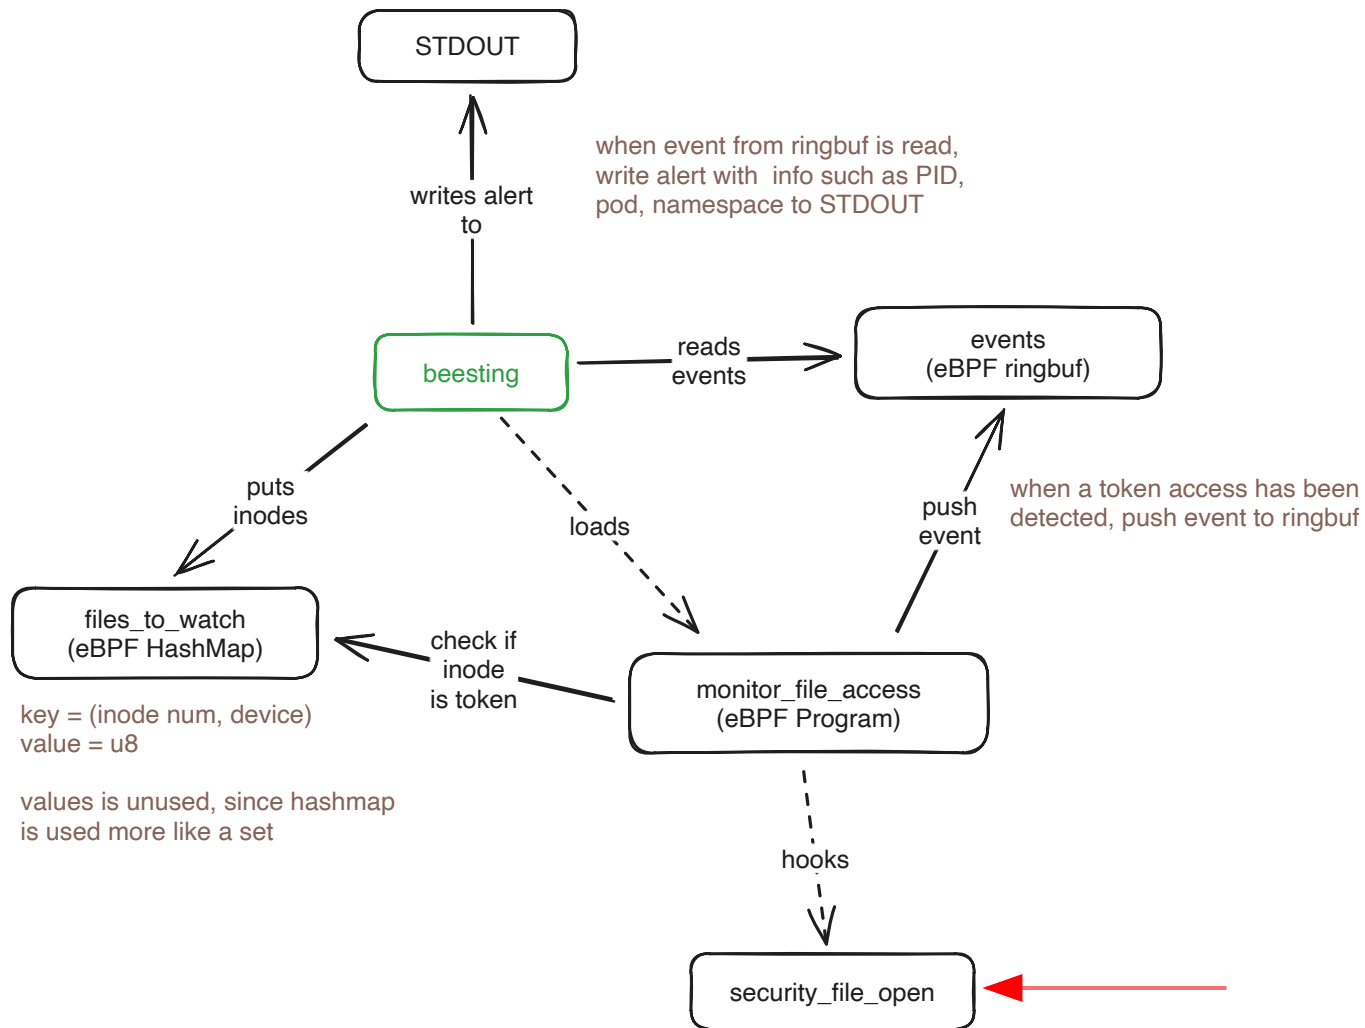












```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.103810ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-svpzl	1/1	Running	0	20s
dummy-8984df79-kddjh	1/1	Running	0	13s

```
$ skaffold run
```

```
...
```

```
Waiting for deployments to stabilize...
```

```
Deployments stabilized in 5.103810ms
```

```
You can also run [skaffold run --tail] to get the logs
```

```
$ k delete -f HACK/dummy.yaml
```

```
deployment.apps/dummy deleted
```

```
$ k apply -f HACK/dummy.yaml
```

```
deployment.apps/dummy created
```

```
$ k get pods
```

NAME	READY	STATUS	RESTARTS	AGE
beesting-agent-svpzl	1/1	Running	0	20s
dummy-8984df79-kddjh	1/1	Running	0	13s

```
$ k exec deploy/dummy -- ls -alh /var/run/secrets/eks.amazonaws.com/
total 12K
drwxr-xr-x    2 root    root    4.0K Jan 11 09:14 .
drwxr-xr-x    4 root    root    4.0K Jan 11 09:14 ..
-rw-r--r--    1 root    root    16 Jan 11 09:14 access_key_token
```

```
$ k exec deploy/dummy -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root    root    4.0K Jan 11 09:14 .  
drwxr-xr-x    4 root    root    4.0K Jan 11 09:14 ..  
-rw-r--r--    1 root    root    16 Jan 11 09:14 access_key_token
```

```
$ k exec deploy/dummy -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root    root      4.0K Jan 11 09:14 .  
drwxr-xr-x    4 root    root      4.0K Jan 11 09:14 ..  
-rw-r--r--    1 root    root       16 Jan 11 09:14 access_key_token
```

```
$ k exec deploy/dummy -- cat /var/run/secrets/eks.amazonaws.com/access_key_token  
2yWeFNuHzb7wUw==
```

```
$ k exec deploy/dummy -- ls -alh /var/run/secrets/eks.amazonaws.com/  
total 12K  
drwxr-xr-x    2 root    root      4.0K Jan 11 09:14 .  
drwxr-xr-x    4 root    root      4.0K Jan 11 09:14 ..  
-rw-r--r--    1 root    root      16 Jan 11 09:14 access_key_token
```

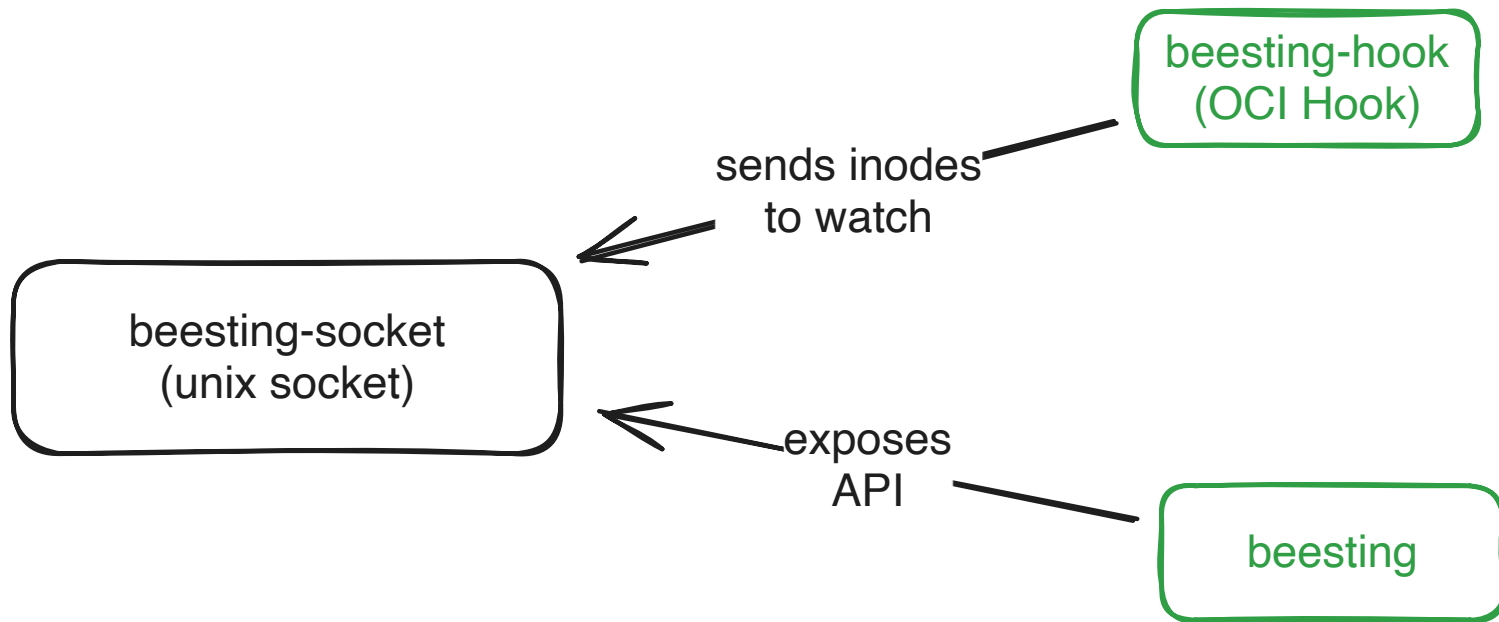
```
$ k exec deploy/dummy -- cat /var/run/secrets/eks.amazonaws.com/access_key_token  
2yWeFNuHzb7wUw==
```

```
$ k logs ds/beesting-agent --tail=2  
time=2025-01-11T09:14:37.455Z level=DEBUG msg="watch inode" token.Inode=92 token.Dev=78  
2025-01-11T09:23:17Z 🌟 Honey Token Access Detected!  
  Pod: default/dummy-8984df79-kddjh  
  Container: dummy-pod,  
  PID: 379512,  
  StartTime: 584352740276296  
  Comm: cat
```

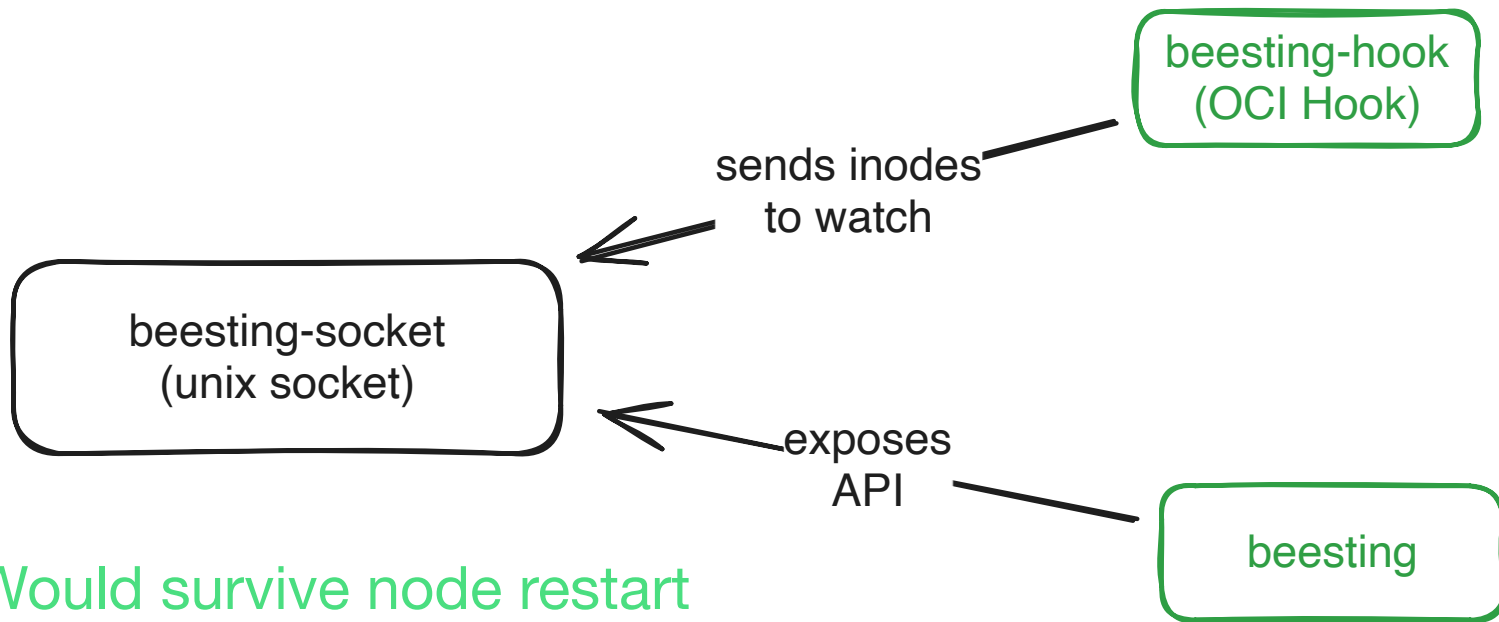



Further research

File injection based on PoCv1



monitoring is the same as PoCv3



Would survive node restart

How to secure it?

monitoring is the same as PoCv3

Further research

Hook read operations and scan for
pattern

Very flexible

<https://github.com/patrickpichler/beesting/>

<https://patrickpichler.dev>

Home

Posts 📡

Figuring out which helpers are available in what kernel version in eBPF

eBPF helpers are a vital part of any eBPF program. It is often not easy to figure out, which helper you have available for a certain program type at a given Linux Kernel Version. The goal of this blog post is,...

November 10, 2024 · 6 min · Patrick Pichler

Hello Blog

Hello world from patrickpichler.dev I am planning to write about different kubernetes/cloud/container related topics. So stay tuned!

September 26, 2023 · 1 min · Patrick Pichler

