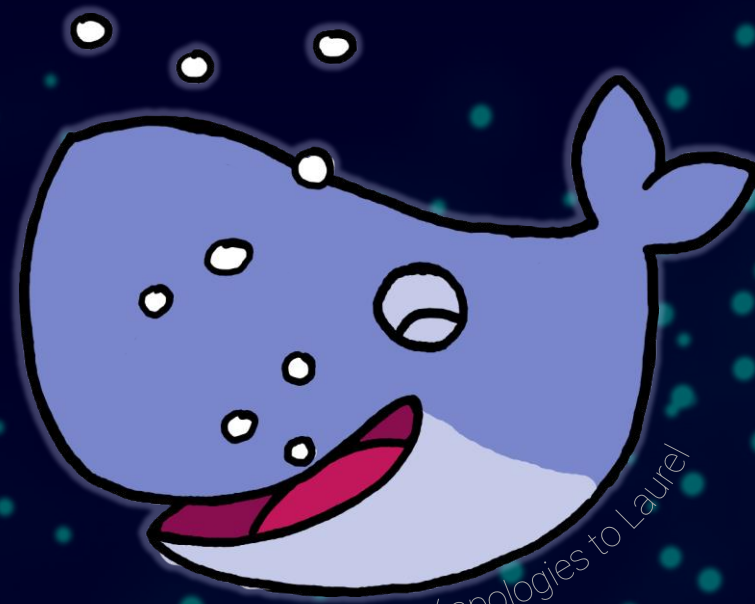
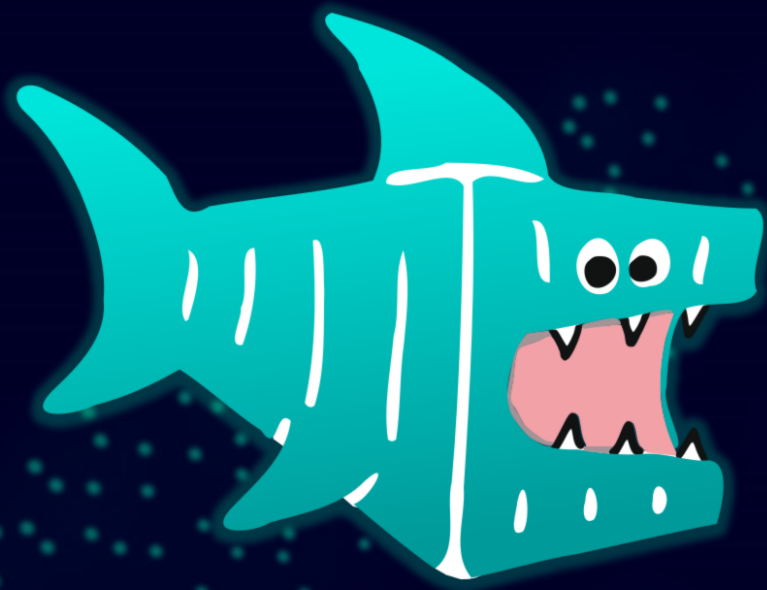


Edgeshark

Capture live container traffic from your comfy desktop Wireshark



w/ apologies to Laurel

\$ who am i



Principal Key Expert in Siemens Digital Industries

- virtual (real-time) communication for containers and VMs
- container technology, especially some of its building blocks
- OT container security

Open Source (GitHub)



- [@siemens/edgeshark](#) – *this presentation*
- [@onsi/gomega/gleak](#) – *leak testing for Go routines*
- [@google/nftables](#) – *xtables match/target payload support*
- [@thediveo/lxkns](#) – *Linux kernel namespace discovery*
- [@thediveo/morbid](#) – *ephemeral containers in unit testing*
- ...

Edgeshark?

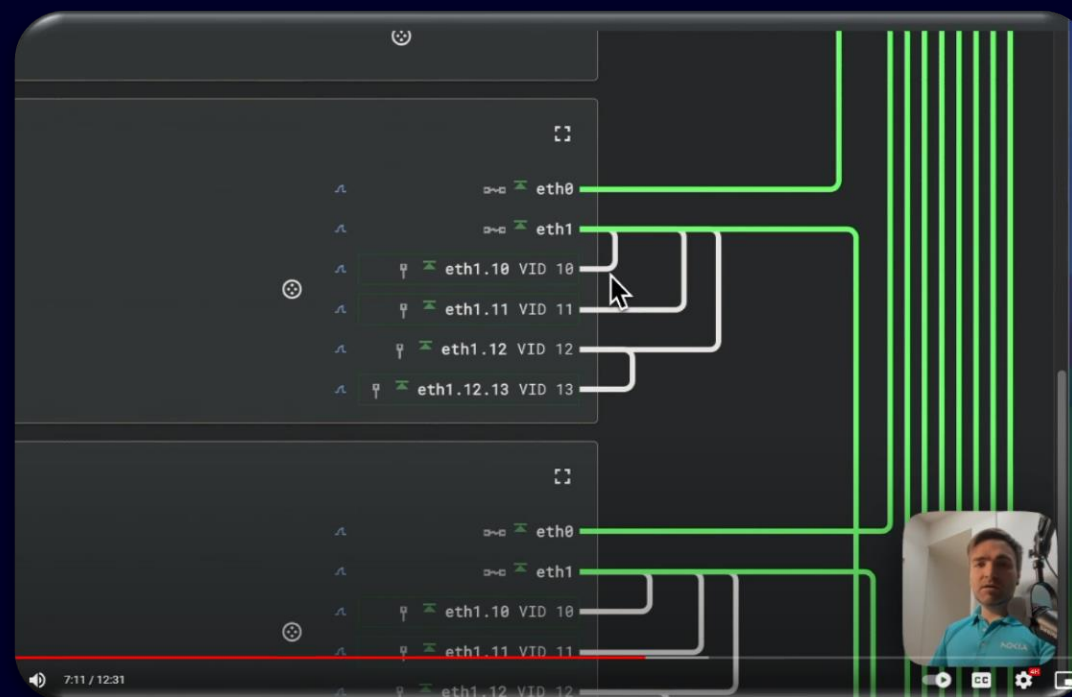


...erm, noooo!

Edgeshark!

*In the words of Roman Dodin,
Product Line Manager at Nokia, Github: @hellt*

“ It allows to do magic ”



www.youtube.com/watch?v=iY90a_Gn5W0

Do YOU Need Magic?

YOUR
container **communication**?



tap into container's and host's...

- "wiring"
- IP configuration
- DNS configuration
- open and forwarded ports
- traffic

...all put into context

HOWTO Magic?

Before: Shell Acrobatics

```
$ # capture remote container traffic
$ echo "YESIWANTMYSUDOPASSWORDINMYSHELLHISTORY" | \
  ssh me@containerhost 'sudo -S nsenter -t
    $(docker inspect --format "{{.State.Pid}}"
    edgeshark-gostwire-1)
    -n dumpcap -q -i eth0 -w -' | \
  wireshark -k -i -
```

```
$ # needs IP tooling inside cont
$ ssh me@containerhost \
  docker exec mycontainer ip route show
```

```
$ # needs cooperative shell inside container
$ ssh me@containerhost \
  docker exec mycontainer cat /etc/resolv.conf
```

SORRY!
no SSH for you

Edgeshark: All in a Single Place

The screenshot displays the Edgeshark interface with several sections:

- transport**: A table showing network traffic details.
- routing**: A table showing IP routes.
- network interfaces**: A list of network interfaces.

St	Proto	Socket	Address	Port	Service	Remote	Port	Service	Group	Container	Proce
↻	TCP		0.0.0.0	:5001						edgeshark-edgeshark-1	
↻	TCP		::	:5001						edgeshark-edgeshark-1	
↻	TCP		172.18.0.3	:5001		172.18.0.1	:41136			edgeshark-edgeshark-1	
↻	TCP		172.18.0.3	:5001		192.168.46.1	:55568			edgeshark-edgeshark-1	
↻	TCP		172.18.0.3	:5001		172.18.0.1	:41136			edgeshark-edgeshark-1	
↻	TCP		172.18.0.3	:5001		192.168.46.1	:55568			edgeshark-edgeshark-1	
↻	TCP		127.0.0.11	:38659					systemd(1)	dockerc	
↻	TCP		172.18.0.3	:41280		172.18.0.2	:5000			edgeshark-edgeshark-1	
↻	UDP		127.0.0.11	:47149					systemd(1)	dockerc	

IP	Destination	Gateway	Metric
0.0.0.0/0	172.18.0.1	eth0	metric 0
172.18.0.0/16		eth0	metric 0

network interfaces

- eth0 (~mcv1an)
- veth7e13b61
- br-875bfa57f100 (~ghost-in-da-edge)
- systemd(1)

edgeshark

- [services.linuxkit/01-docker]edgeshark-edgeshark-1 []
- [services.linuxkit/01-docker]edgeshark-gostwire-1 []

Instant Magic

github.com/siemens/edgeshark ▶ [README](#) ▶ [Quick Start](#) ▶ [Docker Host](#)



service



extcap

WIRESHARK

`wget ... | docker compose -f - up`
(or App deployment in case of Industrial Edge)

github.com/siemens/cshargextcap

[wireshark.org](https://www.wireshark.org)

Container Host
(Docker, ...)

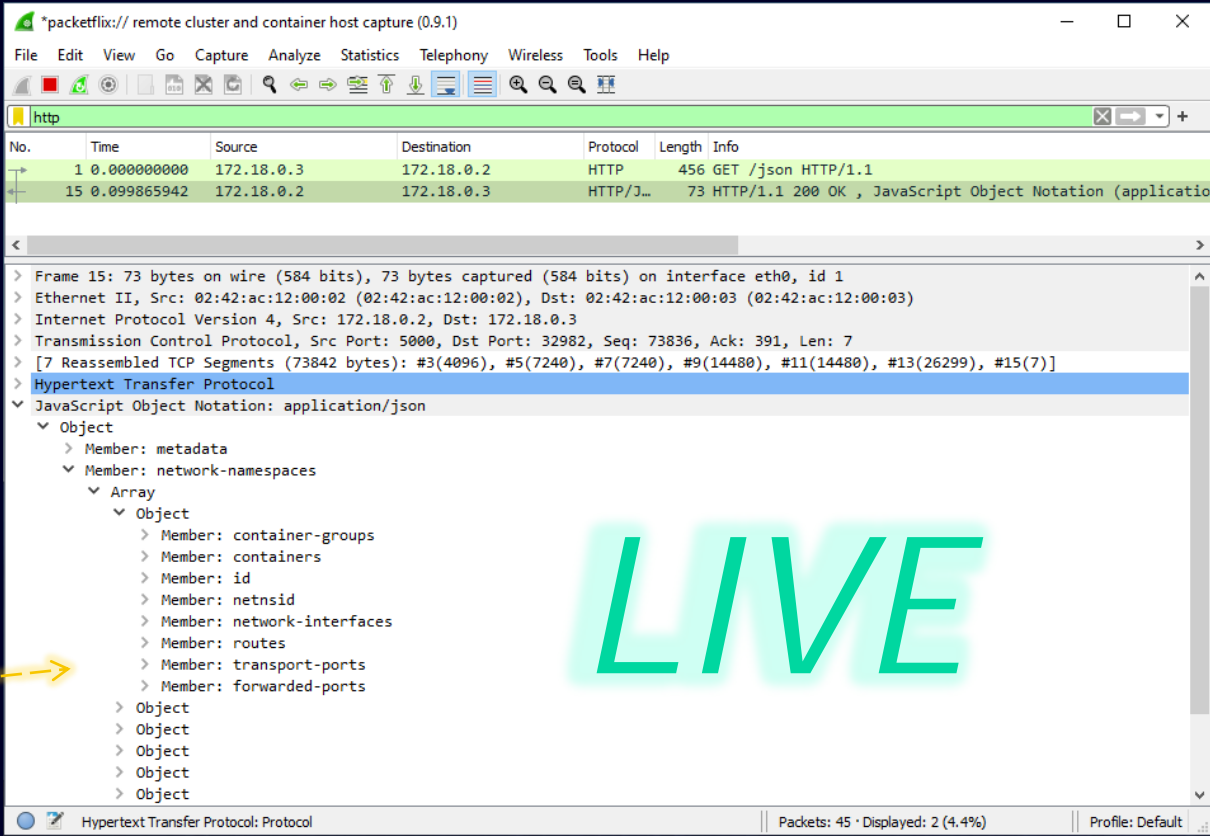
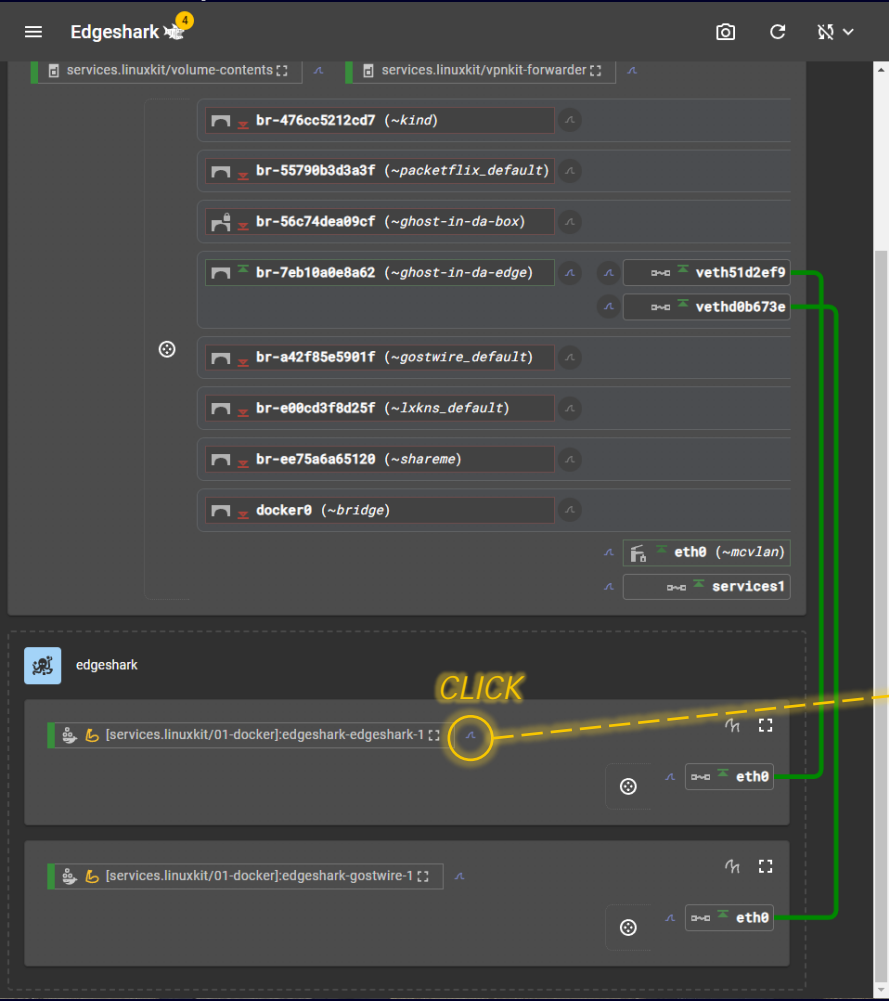
- vanilla Linux
- WSL2
- Docker Desktop (macos, Windows)
- Industrial Edge
- x86 64bit
- ARM 64bit

Desktop
(Browser, Wireshark)

- Linux (x86 64bit, ARM 64bit)
 - .apk (Alpine)
 - .deb (Debian/Ubuntu)
 - .rpm (Fedora/...)
 - PKGBUILD (Arch)
- Windows (x86 64bit only)
- macos ARM64/x86 64bit (.zip)

Tap into Virtual Networking Topology and Traffic

http://localhost:5001



LIVE

Tap into Container Communication Details

container-internal forwarded ports
(here: Docker's embedded DNS)

open or connected ports

IP routing and address(es)

The screenshot displays the network namespace ID: 4026533134. It shows a list of containers under 'neighborhood services (host-internal)'. Below this, there is a 'port forwarding' section with a table:

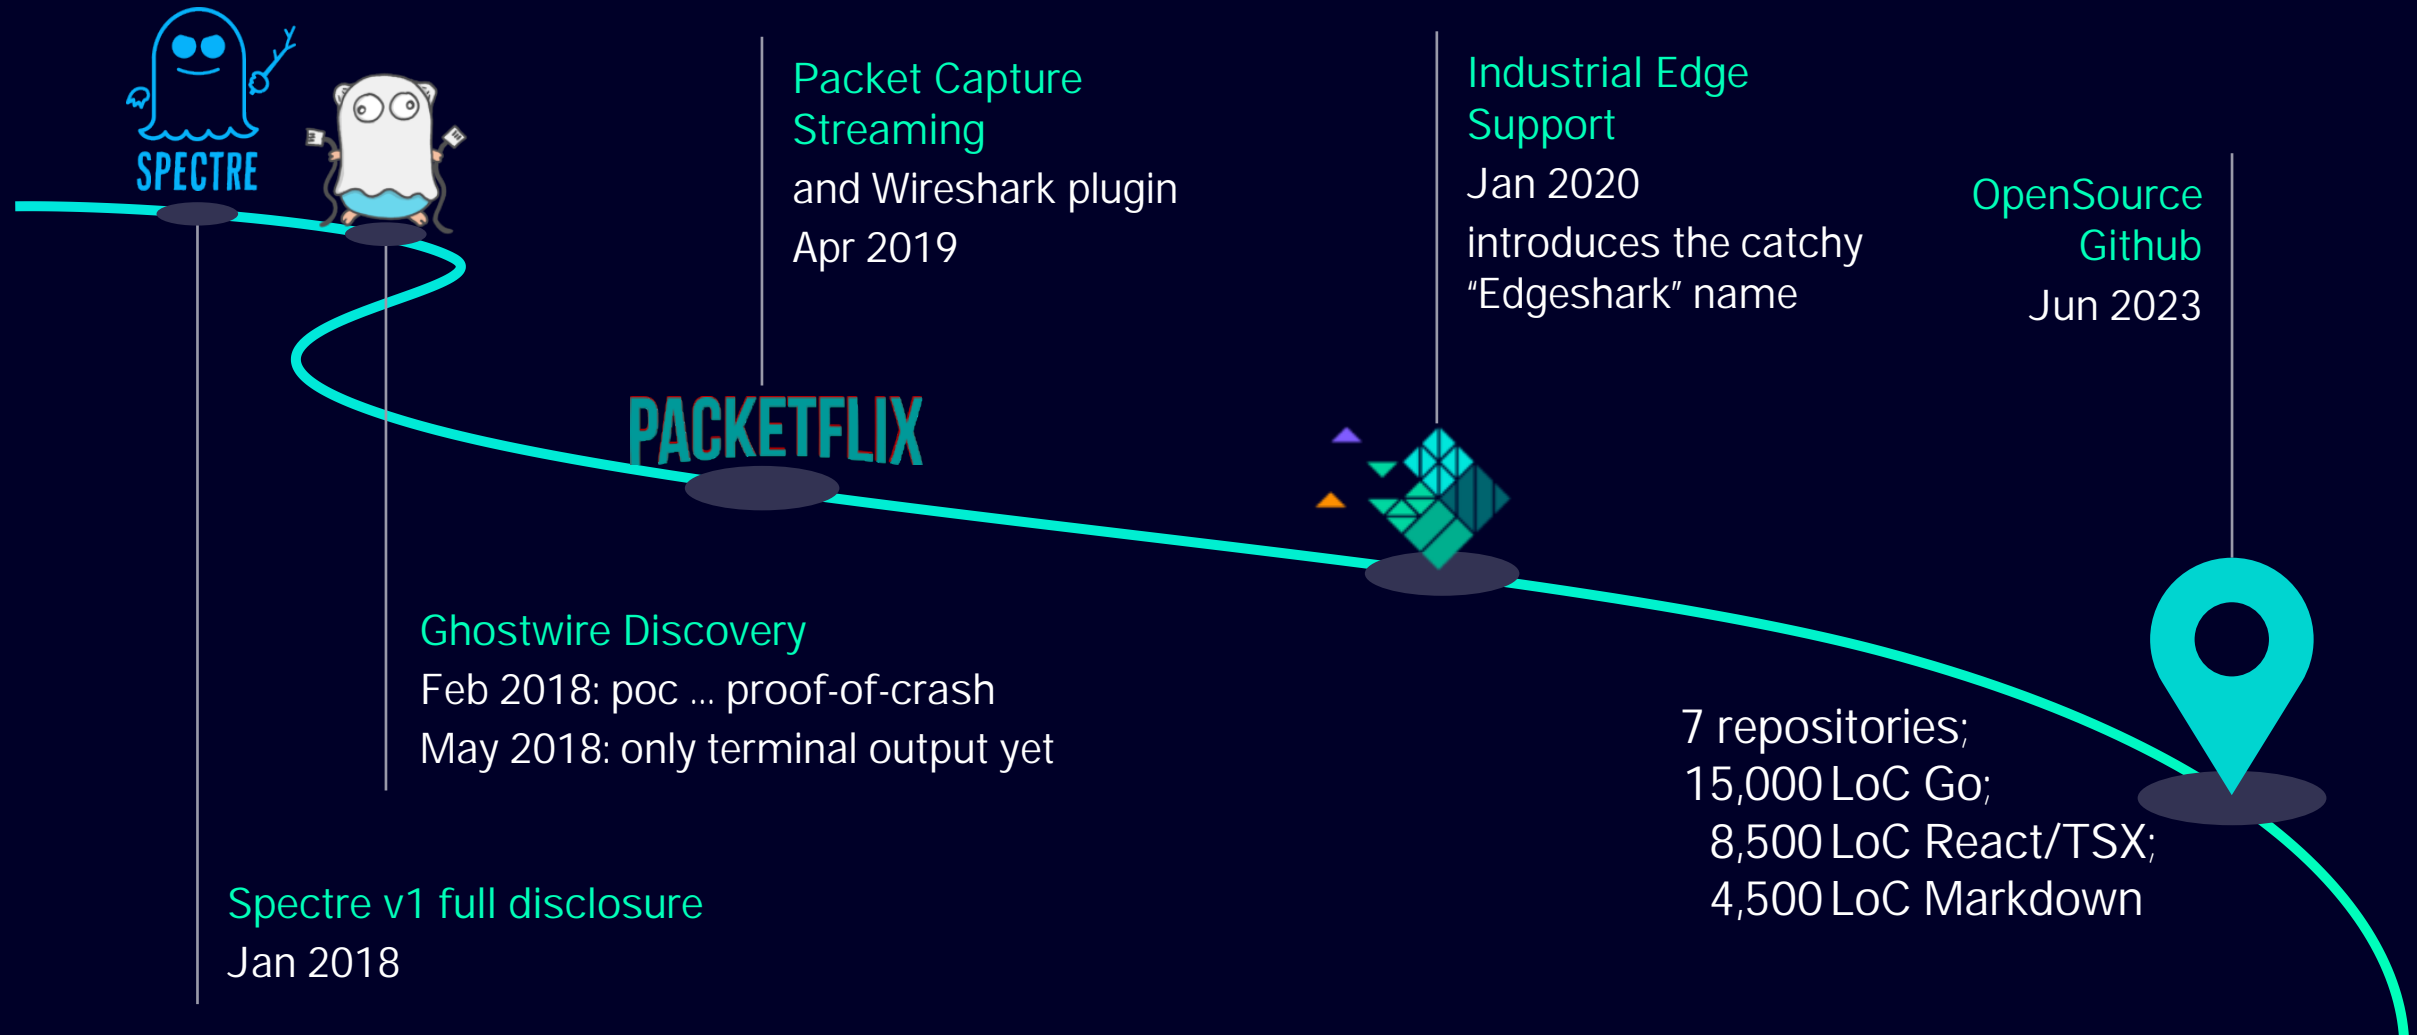
Proto	Address	Port	Service	Forwarded to	Port	Service	Group · Container · Process
TCP	127.0.0.11	:53	domain	127.0.0.11	:38659		systemd(1) · dockerd -H fd:// -contain
UDP	127.0.0.11	:53	domain	127.0.0.11	:47149		systemd(1) · dockerd -H fd:// -contain

Below the port forwarding table is a 'transport' section with a table showing open or connected ports:

St	Proto	Socket	Address	Port	Service	Remote	Port	Service	Group · Container · Proce
🔒	TCP		0.0.0.0	:5001					edgeshark-edgeshark-1
🔒	TCP		::	:5001					edgeshark-edgeshark-1
↔	TCP		172.18.0.3	:5001		172.18.0.1	:41136		edgeshark-edgeshark-1
↔	TCP		172.18.0.3	:5001		192.168.46.1	:55568		edgeshark-edgeshark-1
↔	TCP		172.18.0.3	:5001		172.18.0.1	:41136		edgeshark-edgeshark-1
↔	TCP		172.18.0.3	:5001		192.168.46.1	:55568		edgeshark-edgeshark-1
🔒	TCP		127.0.0.11	:38659					systemd(1) · dockerc
↔	TCP		172.18.0.3	:41280		172.18.0.2	:5000		edgeshark-edgeshark-1
🔒	UDP		127.0.0.11	:47149					systemd(1) · dockerc

At the bottom, the 'routing' section shows routes for 0.0.0.0/0 and 172.18.0.0/16, both pointing to eth0 with metric 0. The 'network interfaces' section shows eth0, veth7e13b61, and br-875bfa57f180 (~ghost-in-da-edge).

Architect as a Service [AaaS]



Architecture of Magic

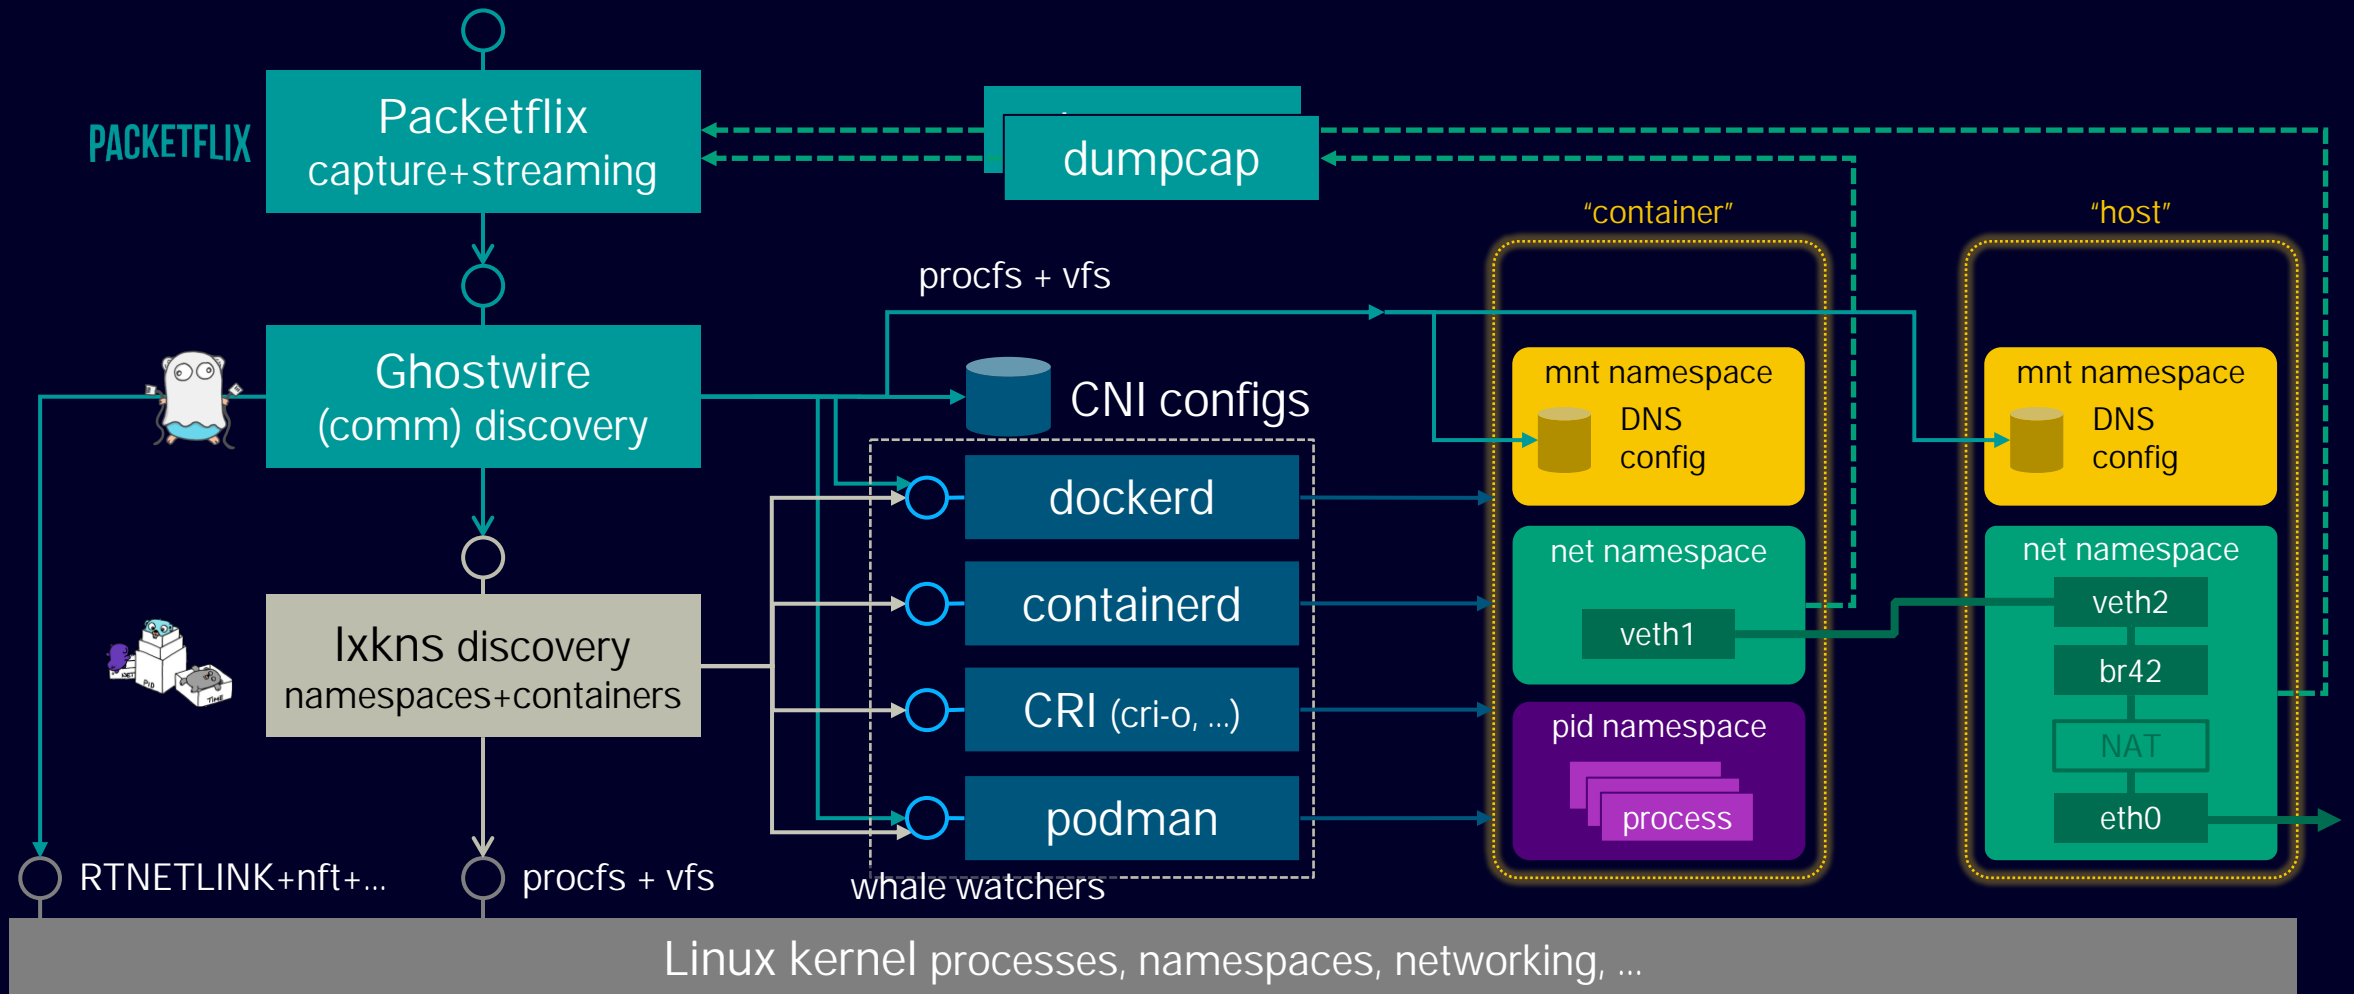
tap into container (host) traffic

tap into container/host configuration data: DNS resolver, CNI, ...

tap into container engine data: container workload, custom network names

tap into kernel networking communication state

Architecture of Magic

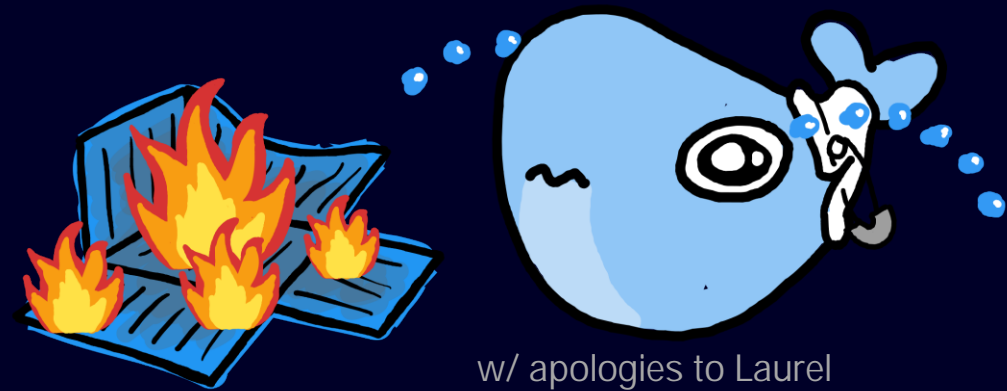


(Where) Do You Mount the Docker Socket?

Any PSS-Os and -Es in the audience?

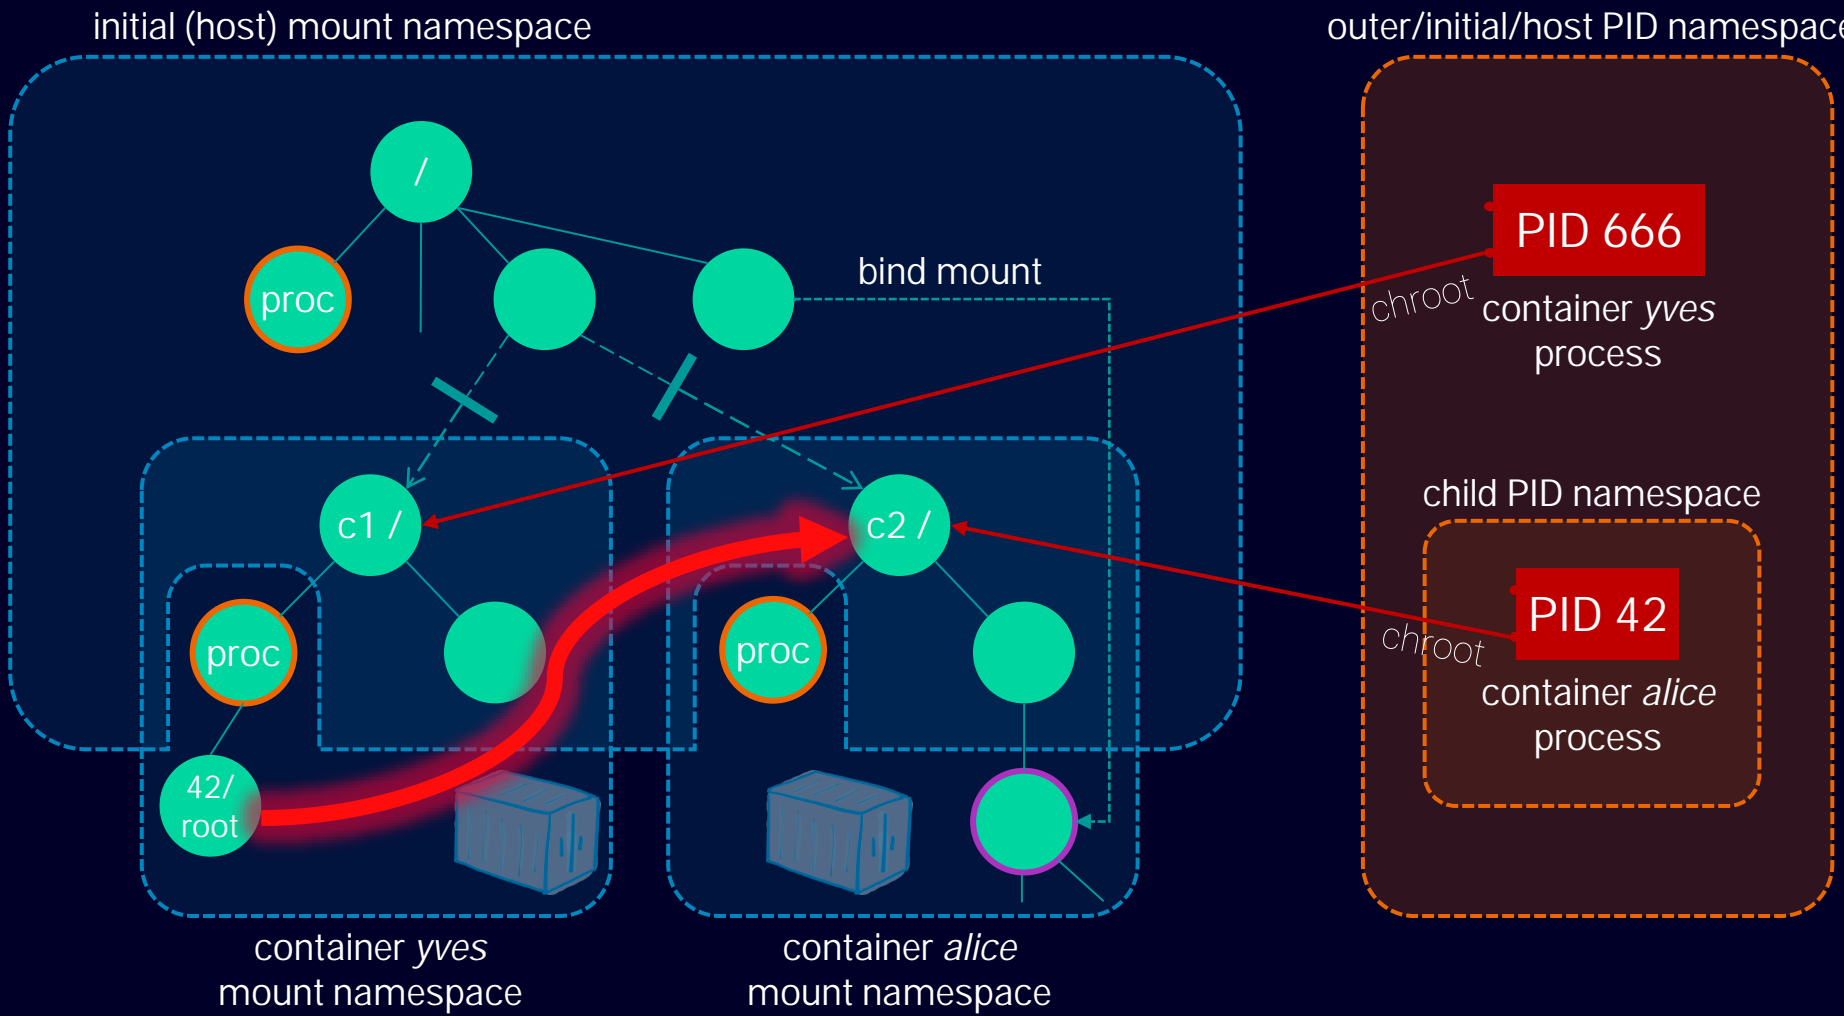
→ *hand out the paper bags*

No, We Don't
[...*Magic*]



w/ apologies to Laurel

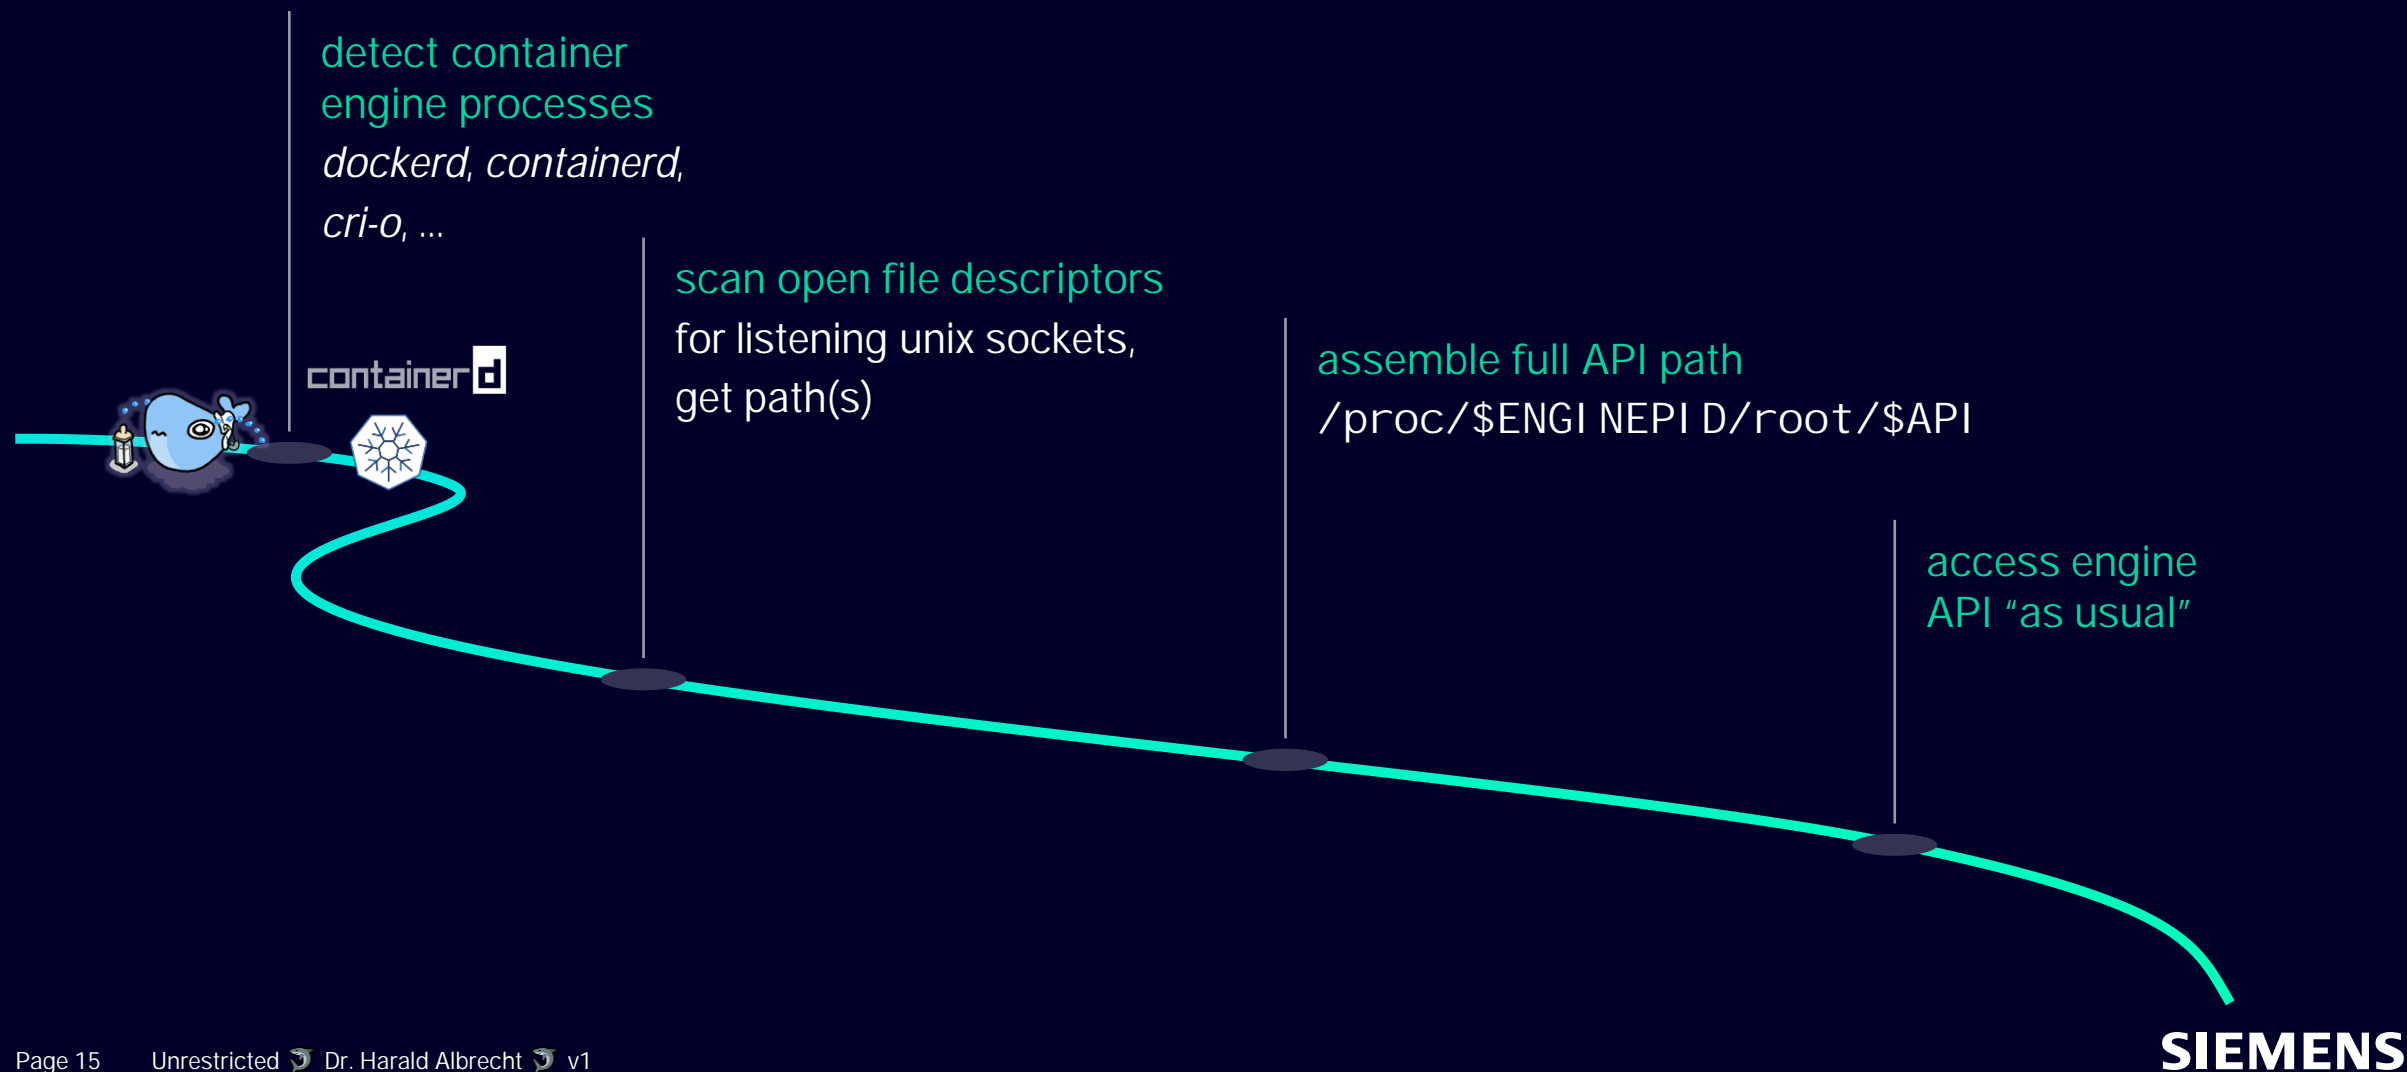
Tapping Into Engines/Containers [Mountineers ... DON'T mount(2)]



- mount namespaces
- separation of mount points
 - NOT: ~~VFS~~ separation.

- proc fs
- /proc/\$PID/root
 - access mediation based on UID/GID or CAP_SYS_PTRACE
 - used by cAdvisor, bcc(!), ...

Whale Whispering: Tapping Into Engines



podman: Detecting the Undetectable



- ↪ podman exits when idle
- ↪ telemetry and diagnosis in prod/development cause *permanent* service operation – but **we cannot rely on podman being active.**

detect
socket activator
processes
...anywhere

scan open fd's for
listening UDS with
well-known suffix
/.../.../podman.sock

activate engine
...and wait for it to
become responsive
as we need its PID

access engine
API "as usual"



Edgeshark

[github.com/siemens/edgeshark]

- live tap into container network traffic directly from Wireshark
 - *downloading* [pcaps] *is so 90's*, everyone's *streaming*!
- ...much more than a “Wireshark Web UI”, Edgeshark...
 - taps into virtual networks topology inside container hosts
 - taps into open and connected ports, forwarded ports (Docker, kube-proxy)
 - taps multiple container engines, *Docker, containerd, cri-o, podman*
 - has engine-in-container auto-detection
 - is KinD-aware (“Kubernetes-in-Docker”)
 - comes with an integrated screenshot feature for documentation purposes
 - ...

| Contact



Dr. Harald Albrecht

Principal Key Expert *Systems Communication and Control Networks*
harald.albrecht@siemens.com

github.com/siemens/edgeshark