GOTO Copenhagen 2017

Conference Oct. 1-3, 2017

Secure Substrate: Building the Moby Whale

Diogo Mónica

Security Lead ***@docker***

Follow us @gotocph

- Infrastructure Company
- Moving 2+ billion containers annually
- Infra: Go, some python

- Mobile payments company
- Moving $65 billion annually
- Infra: Java & Ruby, some Go

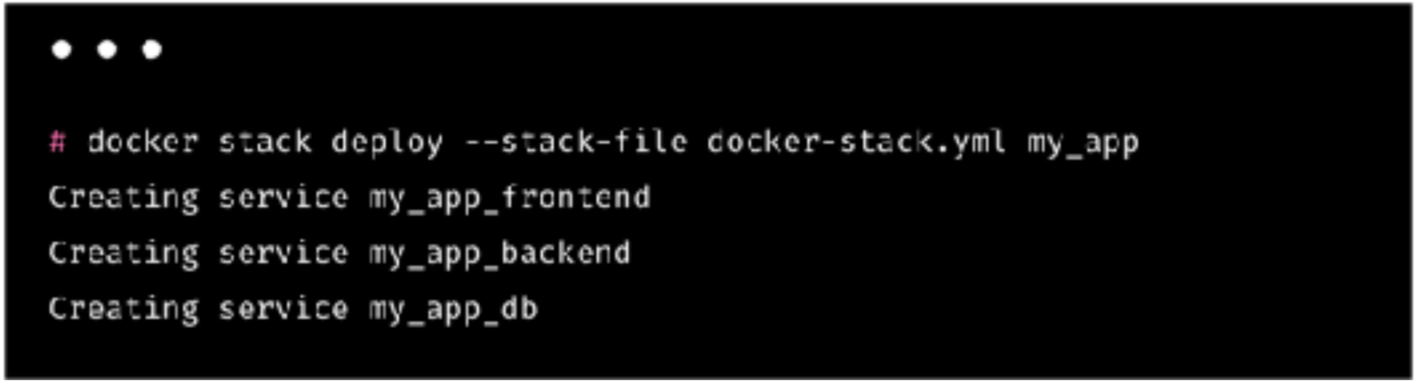

```
# docker stack deploy --stack-file docker-stack.yml my_app
Creating service my_app_frontend
Creating service my_app_backend
Creating service my_app_db
```

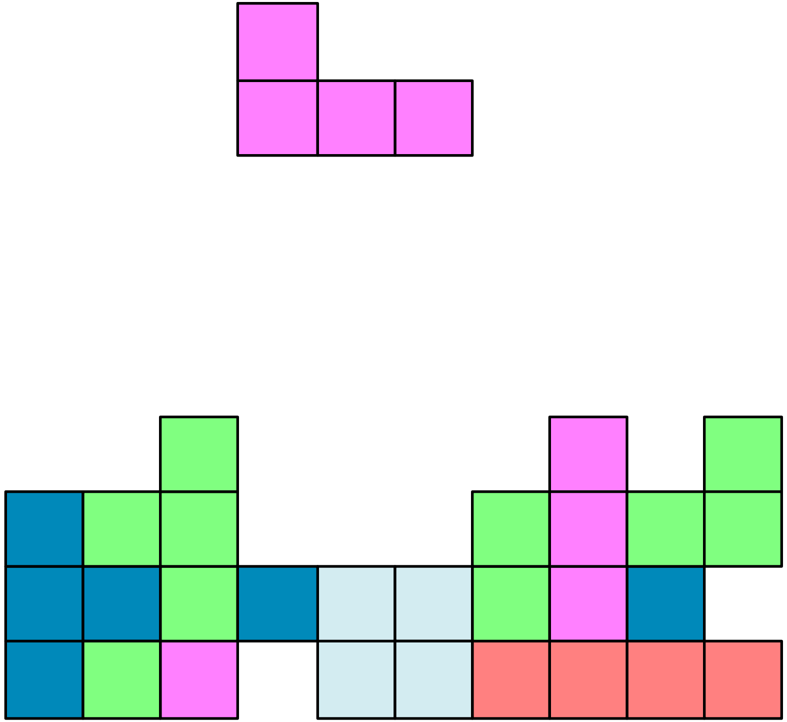

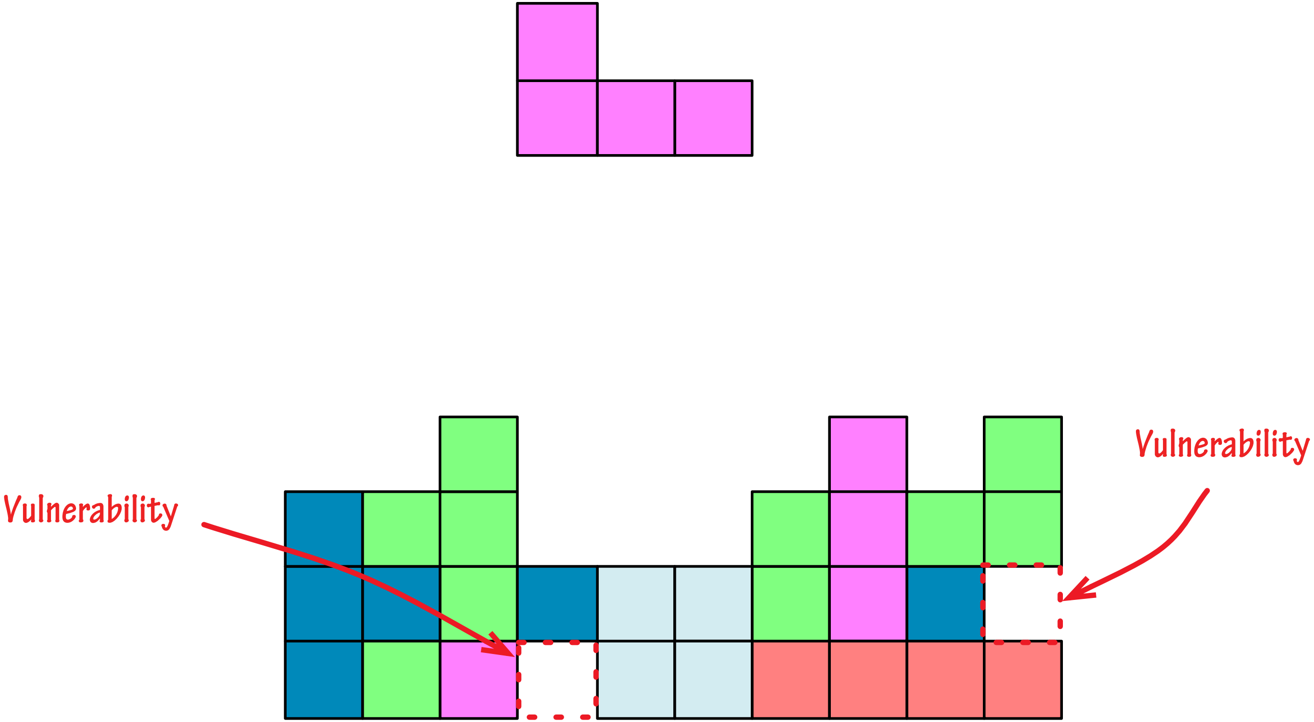

Vulnerability

Vulnerability

Security Tetrominos

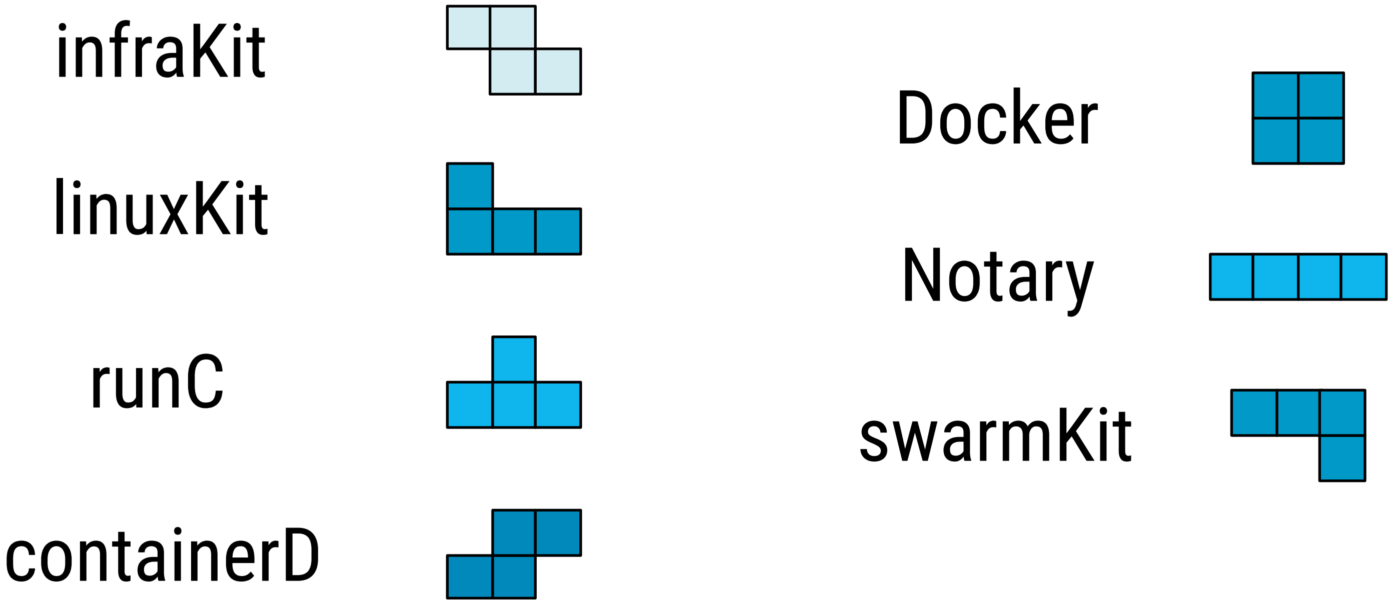

infraKit

linuxKit

runC

containerD

Docker

Notary

swarmKit

infraKit

Infrastructure independent machine management

infraKit

Platform Agnostic

infraKit

Declarative Updates

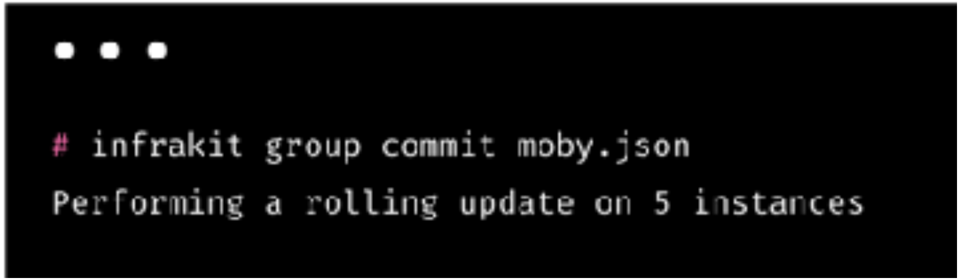

```
# infrakit group commit moby.json
Performing a rolling update on 5 instances
```

infraKit

Reverse Uptime

10:00

infraKit

Reverse Uptime

Reverse Uptime

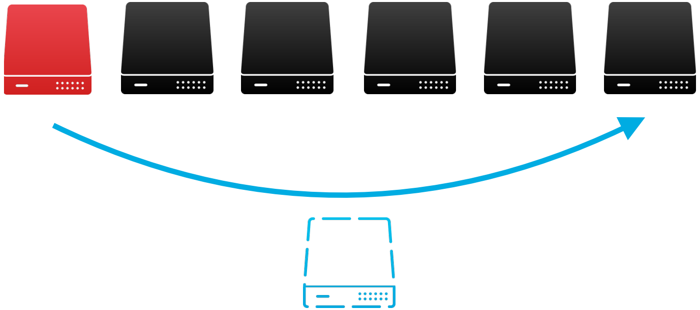

Golden Image Freshness

infraKit

Rolling Deploys

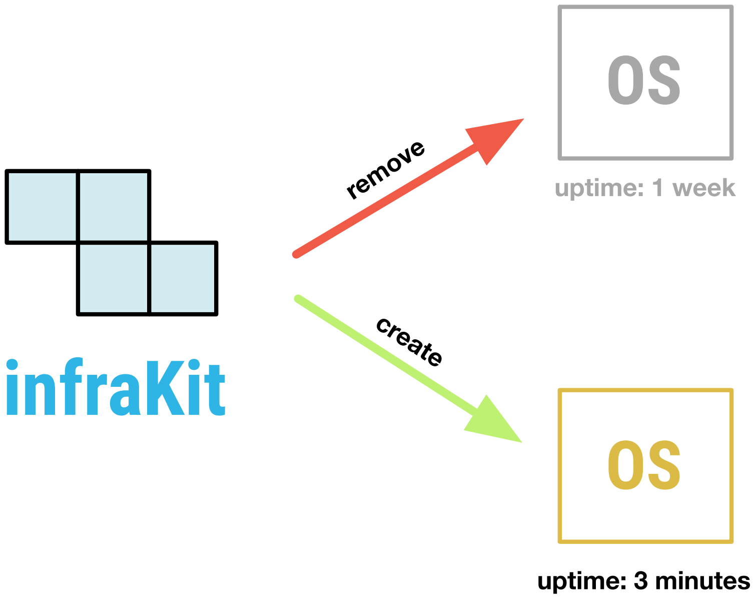

linuxKit

The most secure OS builder for your containers

linuxKit

Immutable Linux
OS builder

```
kernel:
  image: "linuxkit/kernel:4.9.x"
  cmdline: "console=ttyS0 page_poison=1"
init:
  - linuxkit/init:1b8a7e394d…
onboot:
  - name: dhcpcd
    image: “linuxkit/dhcpcd:7d2b8aaaf…
    command: ["/sbin/dhcpcd", “--nobackground" ]
trust:
  org:
    - linuxkit
```

linuxKit

Minimal Base

linuxKit

Immutable Infrastructure

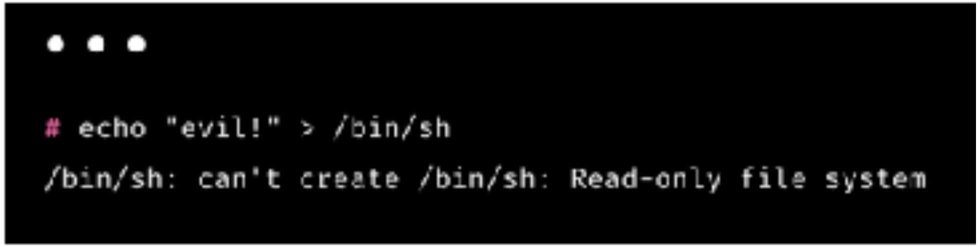

```
# echo "evil!" > /bin/sh
/bin/sh: can't create /bin/sh: Read-only file system
```

linuxKit

Incubating the Future

Landlock LSM

WireGuard

FAST, MODERN, SECURE VPN TUNNEL

Hewlett Packard Enterprise

ARM

intel

Microsoft

linuxKit

Type-safe
System
Daemons

OCaml

linuxKit

Already in use by millions of users

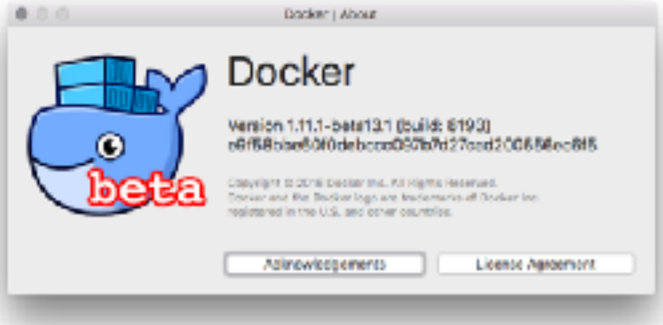

Docker

Version 1.11.1-beta13.1 (build: 8193)
[illegible]

Copyright © 2016 Docker Inc. All Rights Reserved.
Docker and the Docker logo are trademarks of Docker Inc.
registered in the U.S. and other countries.

runC

Lightweight universal container runtime

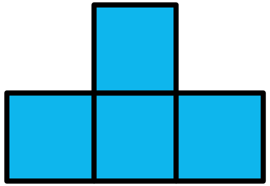

runC

- Namespace Isolation
- Cgroups

Namespaces

PID MNT IPC NET …

Cgroups

CPU BLKIO MEM PIDS …

containerD

Container runtime supervisor

containerD

Image Pulls

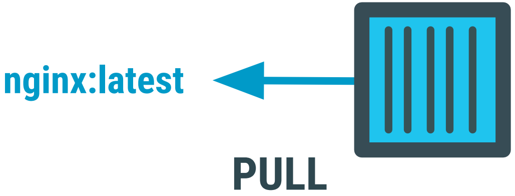

containerD

Image Pulls

nginx:sha256:29d234…

PULL

containerD

Content Addressable
Image Pulls

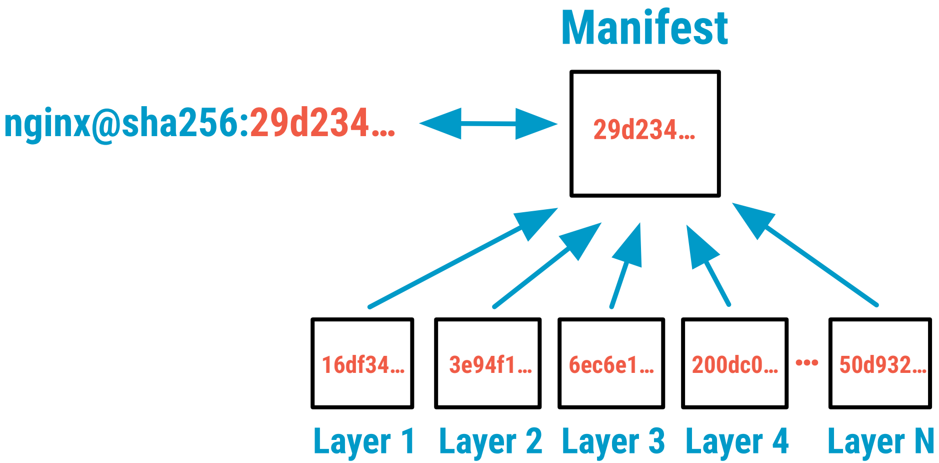

Notary

Trusted software delivery

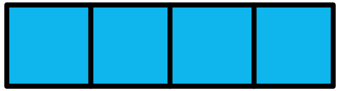

Notary

Image Pulls

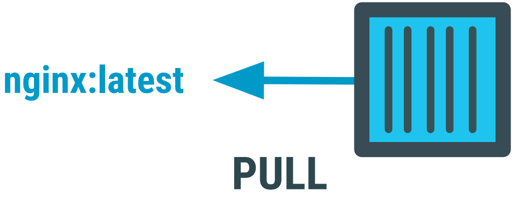

Notary

Cryptographic
Name Resolution

latest	↔	da4f25c…
stable	↔	1b33e92…
edge	↔	9dfe47d…

Notary

- Threshold Signing
- Survival Key Compromise

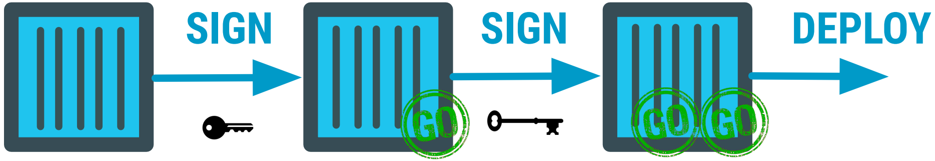

Docker

Secure-by-default software container platform

Docker

- SELinux & AppArmor
- Capability Whitelist
- Syscall Whitelist
- Notary integration

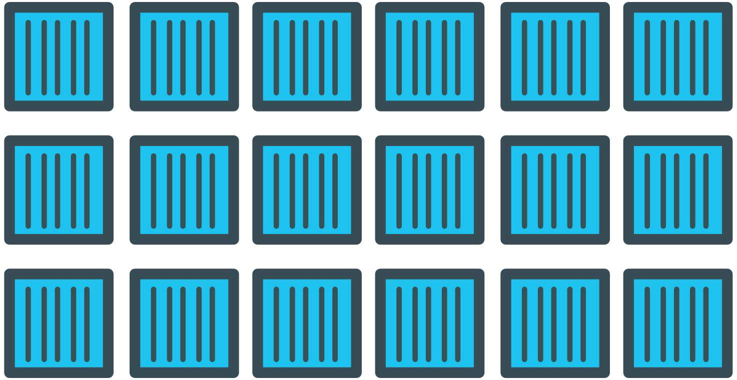

Least-privilege

container orchestrator

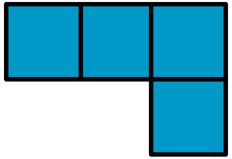

swarmKit

Secure Node Introduction

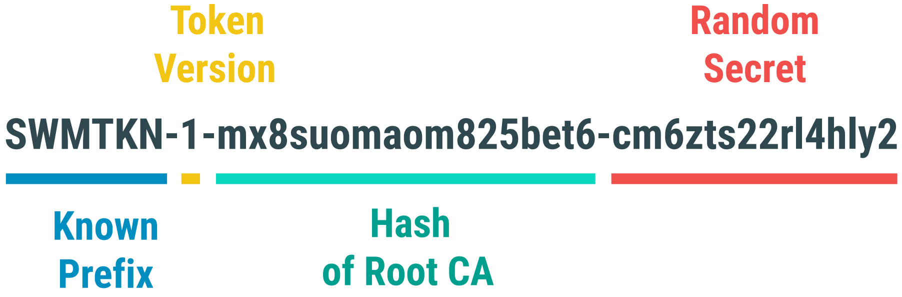

Token Version

Random Secret

SWMTKN-1-mx8suomaom825bet6-cm6zts22rl4hly2

Known Prefix

Hash of Root CA

swarmKit

Cryptographic
Node Identity

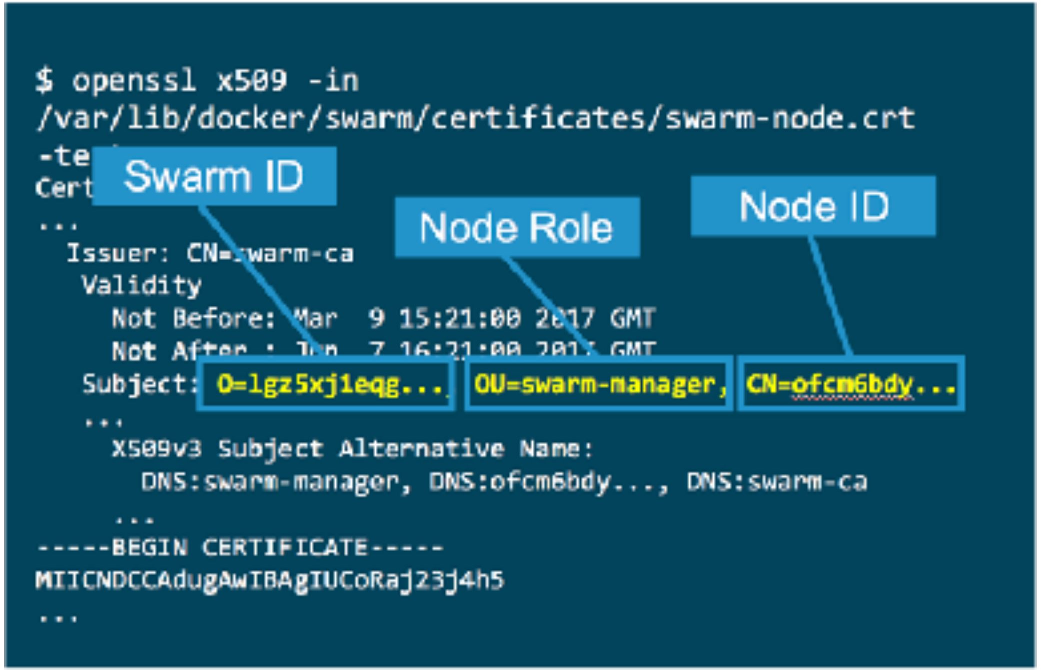

```
$ openssl x509 -in
/var/lib/docker/swarm/certificates/swarm-node.crt
-te
Cert
...
  Issuer: CN=swarm-ca
   Validity
     Not Before: Mar  9 15:21:00 2017 GMT
     Not After : Jun  7 16:21:00 2017 GMT
   Subject: O=lgz5xj1eqg..., OU=swarm-manager, CN=ofcm6bdy...
   ...
     X509v3 Subject Alternative Name:
       DNS:swarm-manager, DNS:ofcm6bdy..., DNS:swarm-ca
       ...
-----BEGIN CERTIFICATE-----
MIICNDCCAdugAwIBAgIUCoRaj23j4h5
...
```

swarmKit

MTLS Between
All Nodes

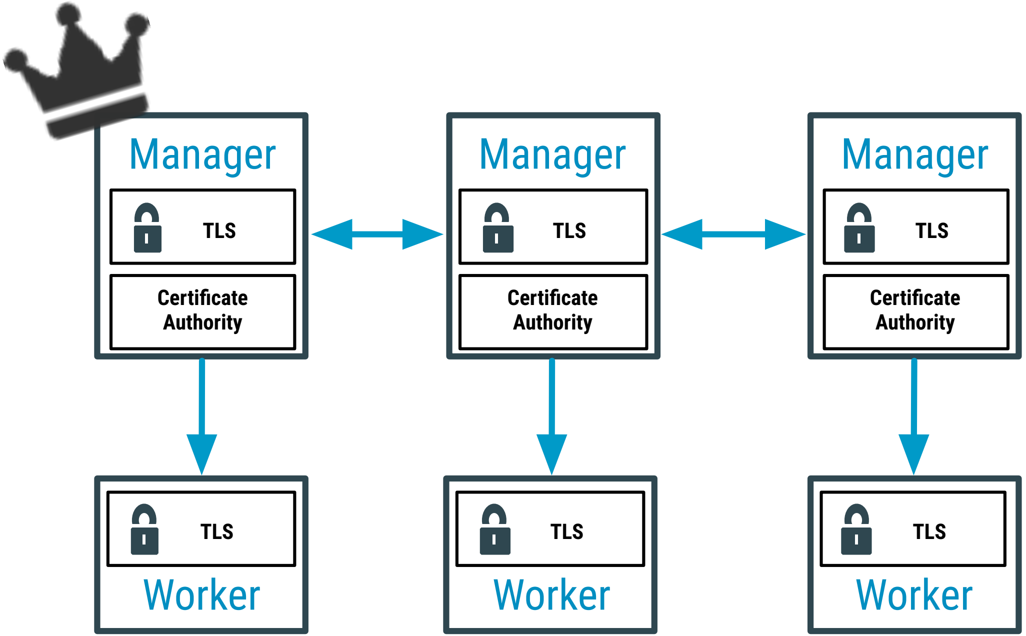

swarmKit

Secure Secret Distribution

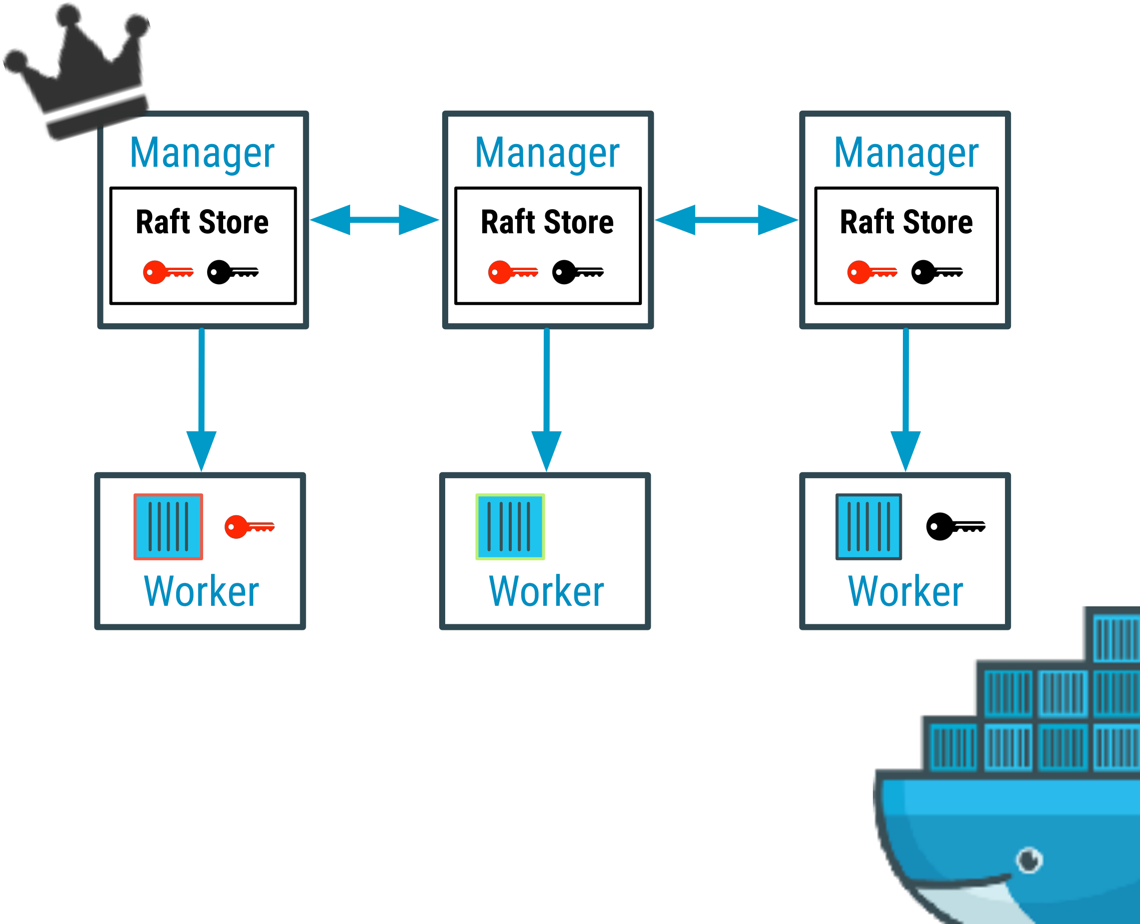

swarmKit

Transparent Root Rotation

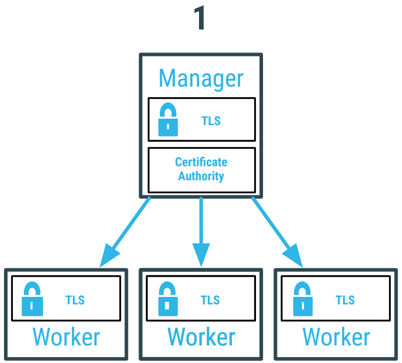

swarmKit

Transparent Root Rotation

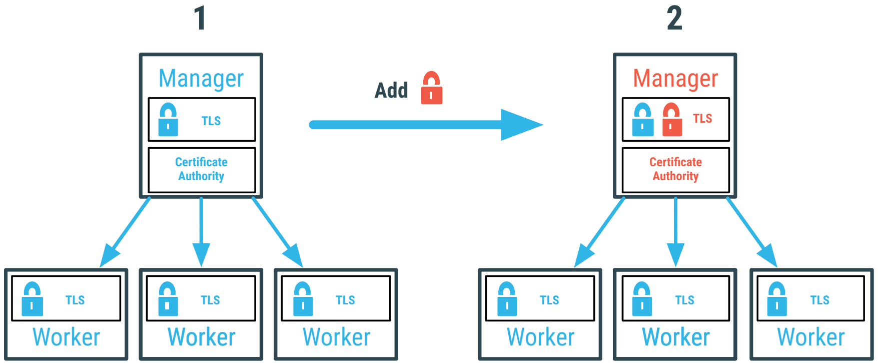

swarmKit

Transparent Root Rotation

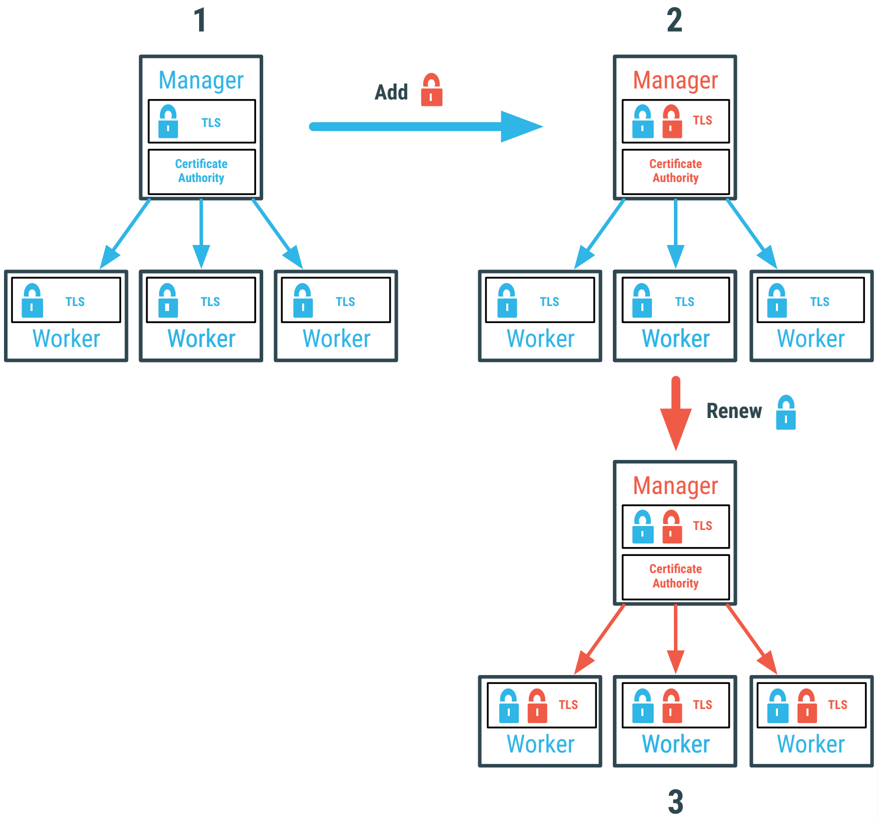

swarmKit

Transparent Root Rotation

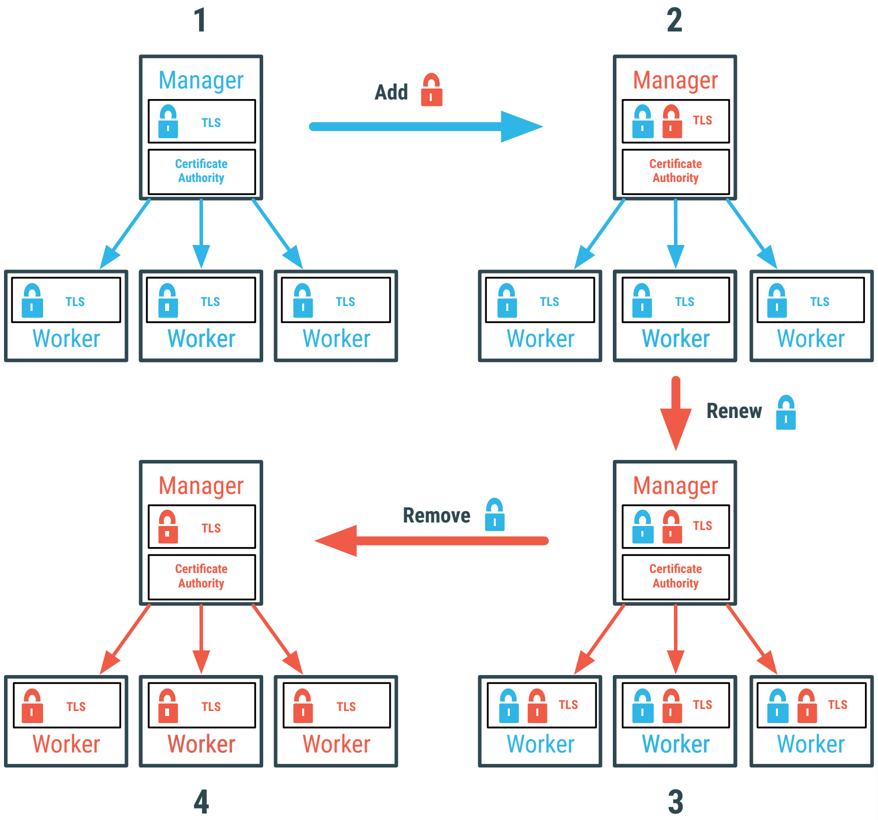

Bringing it all together

Notary for Docker image name resolution

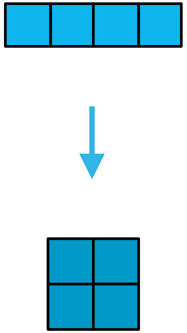

Notary for Docker image name resolution

Cryptographically Verified Pulls

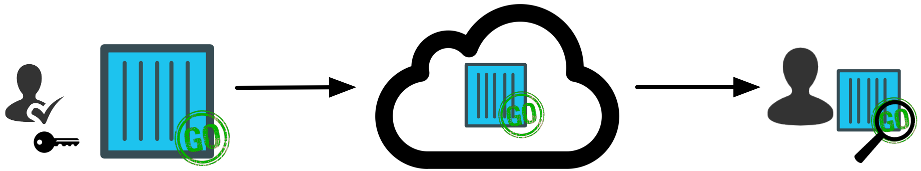

swarmKit delivered Docker containers

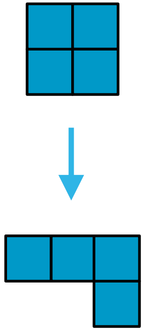

swarmKit delivered Docker containers

Authorized, Authenticated, Encrypted delivery of Resources

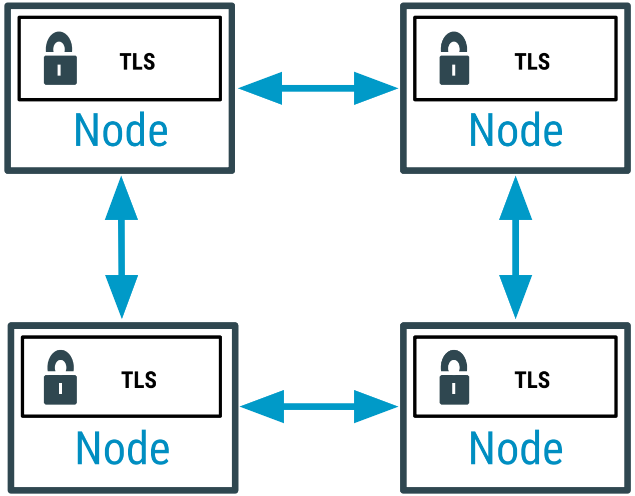

infraKit for swarmKit Bootstrap

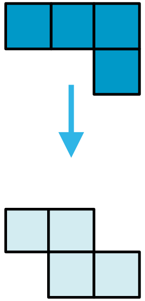

infraKit for swarmKit Bootstrap

Secure Node Cluster Introduction

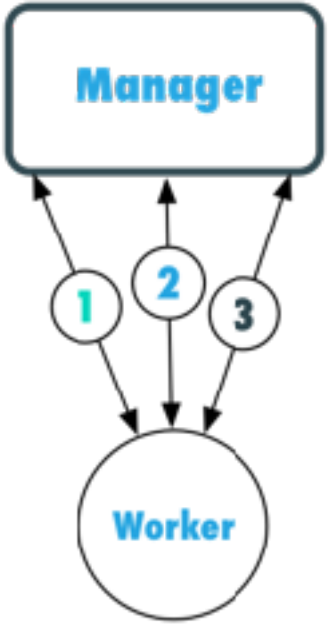

1. Retrieve and validate Root CA Public key material.
2. Submit new CSR along with secret token.
3. Retrieve the signed certificate.

linuxKit as the base OS builder

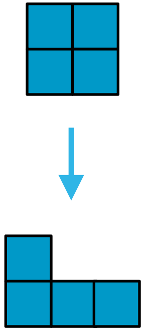

linuxKit as the base OS builder

Hardened Configuration

Notary for secure dependency resolution

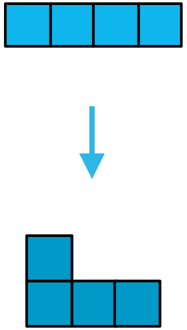

Notary for secure dependency resolution

Cryptographically Verified Build

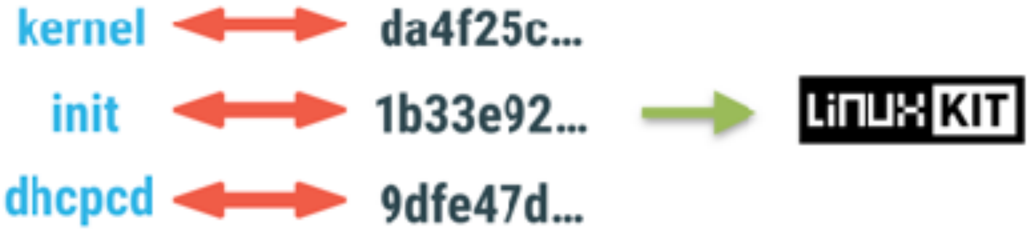

kernel ⟷ da4f25c...

init ⟷ 1b33e92...

dhcpcd ⟷ 9dfe47d...

→ LINUX KIT

infraKit plus Notary for trusted OS Provisioning

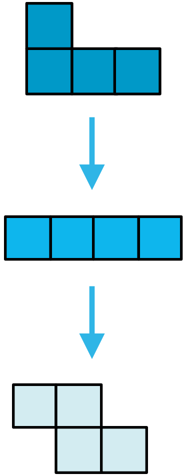

infraKit plus Notary for trusted OS Provisioning

Cryptographically Verified Boot

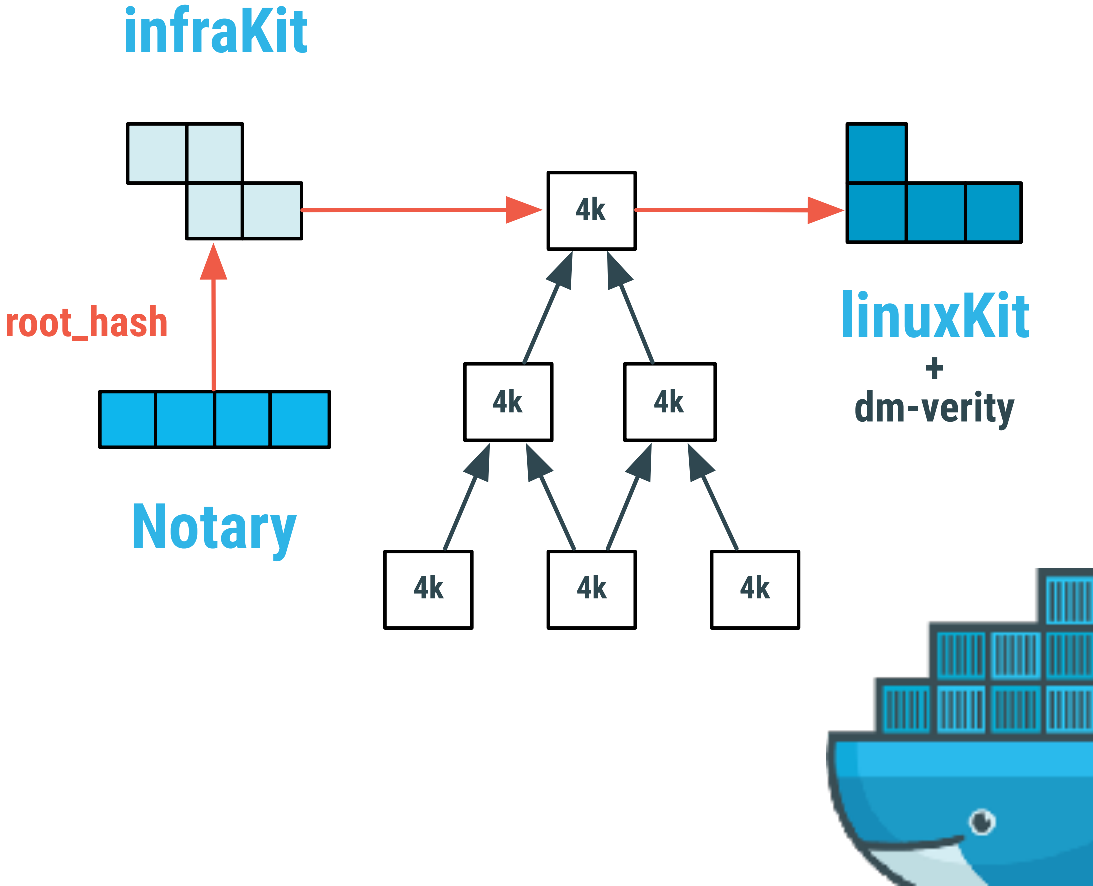

Layered runC, containerD, Docker Runtime

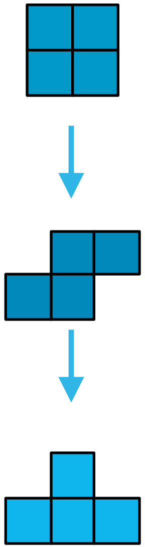

Layered runC, containerD, Docker Runtime

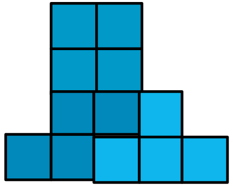

Secure-by-default Container Execution

Secure-by-default Container Execution

runC, containerD, Docker, swarmKit, Notary

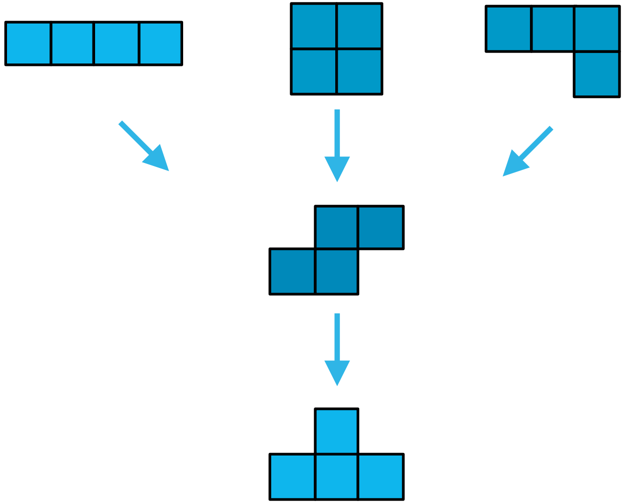

runC, containerD, Docker, swarmKit, Notary

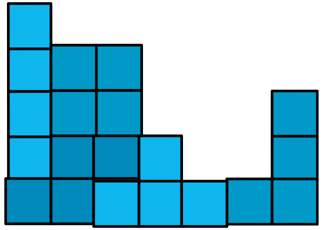

Secure-by-default Container Platform

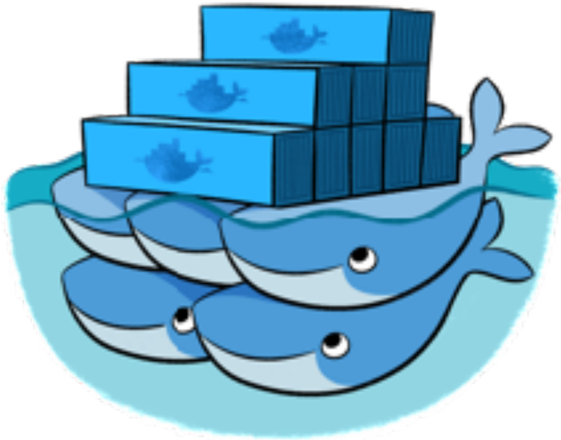

runC, containerD, Docker, swarmKit, Notary, infraKit, linuxKit

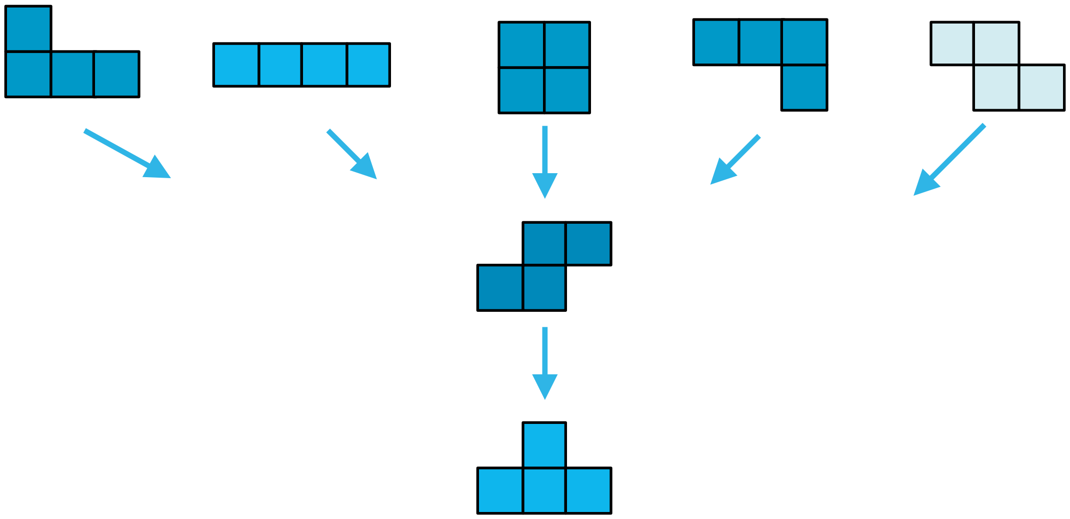

runC, containerD, Docker, swarmKit, Notary, infraKit, linuxKit

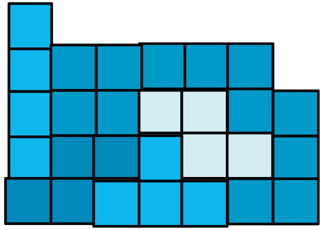

Secure-by-default Infrastructure

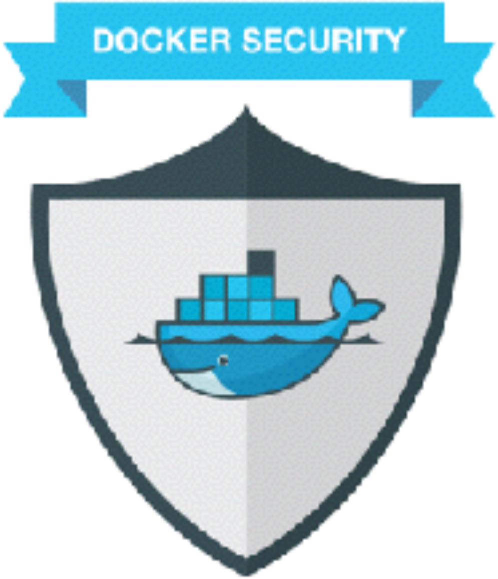

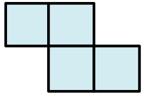

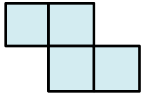

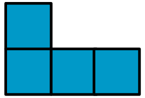

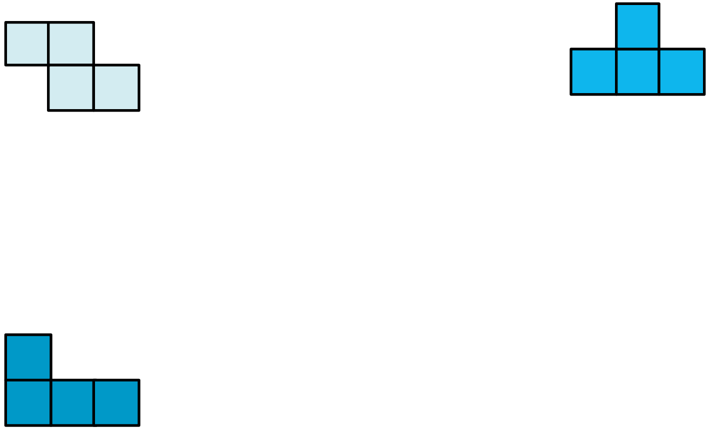

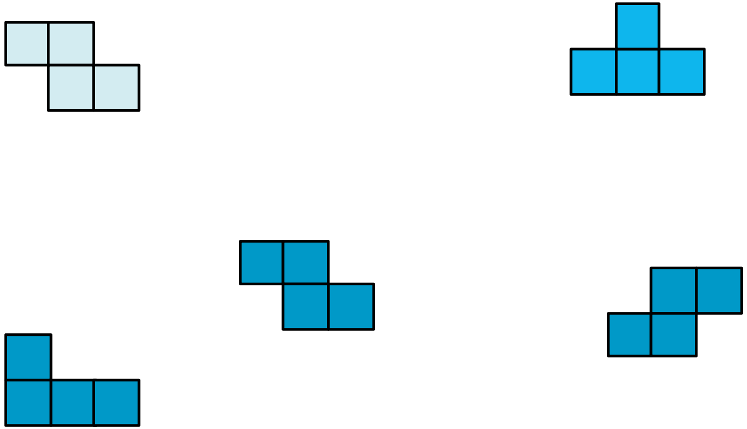

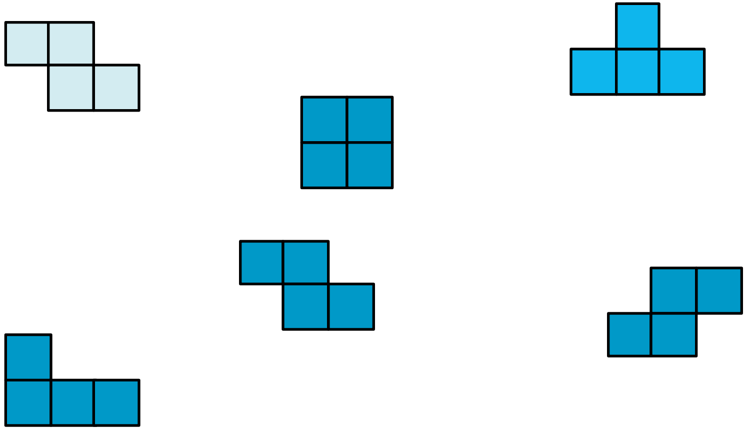

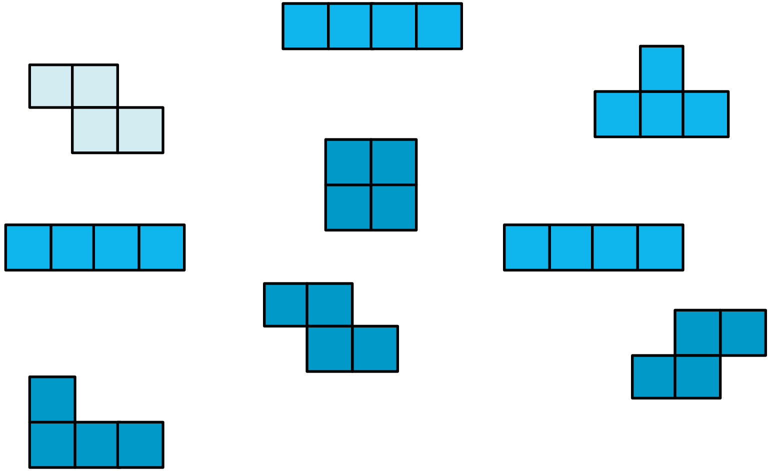

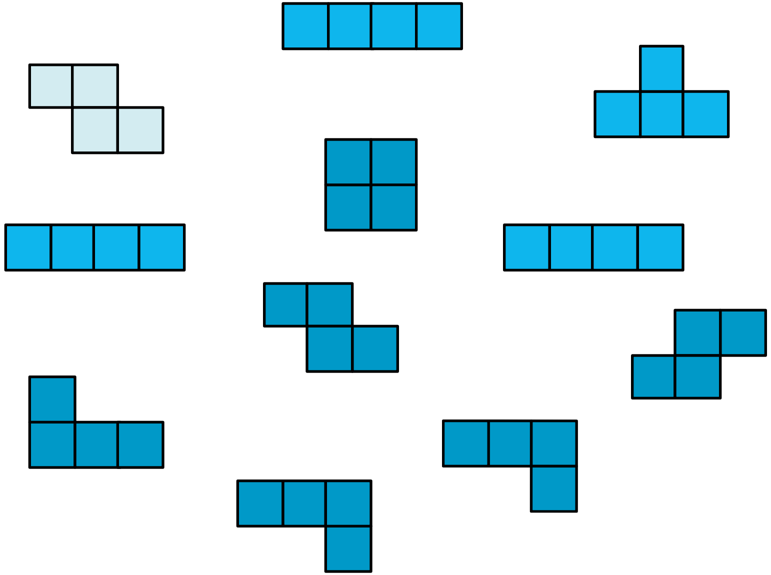

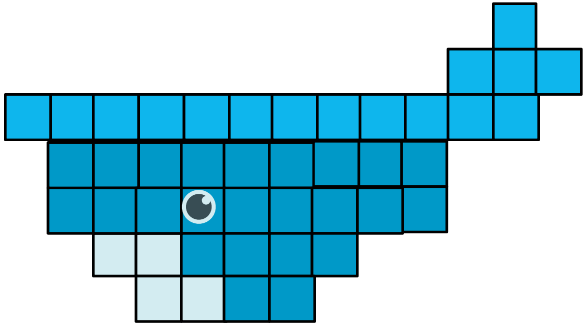

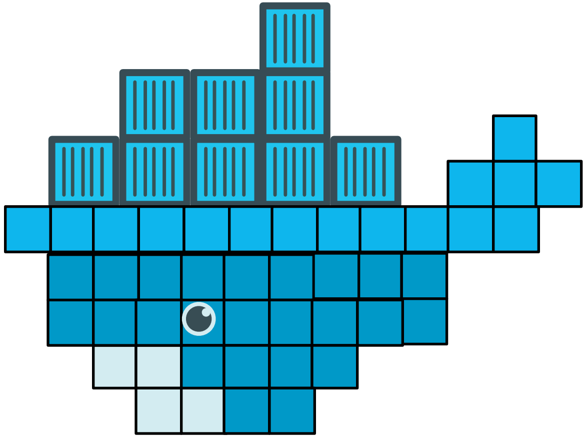

Thank you!

goto; copenhagen

GOTO Copenhagen 2017

Conference Oct. 1-3, 2017

Click 'Rate Session' to rate session and ask questions.

Follow us @gotocph