


Isolating Processes using Docker User Namespaces and Seccomp

4 October 2016

Paul Novarese
Technical Account Manager
Docker, Inc.
pvn@docker.com
@pvn



Agenda

- Preliminaries
- Container Security Considerations
- Containment
- Namespaces
- What is Seccomp?
- Demos?


"CONTAINER"



"SECURITY"

memegenerator.net

The Iceberg



Your code

Your vendor's code

Containment

- namespaces -> what you can see
- cgroups -> what you can use
- seccomp -> what you can do



Containment

...applications deployed in containers are **more secure** than applications deployed on the bare OS because **even if a container is cracked** they **greatly limit the damage** of a successful compromise...

<https://www.gartner.com/doc/3375717/secure-docker-containers-operation>

Namespaces



Namespaces

```
/ # ps -eo pid,user,args
```

```
PID    USER    COMMAND
```

```
 1 root   registry serve /etc/docker/registry/config.yml  
 9 root   sh  
13 root   ps -eo pid,user,args
```

```
/home/docker> ps -eo pid,user,args | egrep "registry|PID"
```

```
  PID USER    COMMAND
```

```
1385 root   registry serve /etc/docker/registry/config.yml
```


Namespaces

```
/ # ps -eo pid,user,args
```

PID	USER	COMMAND
1	root	registry serve /etc/docker/registry/config.yml
25	root	sh
29	root	ps -eo pid,user,args

```
/home/docker> ps -eo pid,user,args | egrep "registry|PID"
```

PID	USER	COMMAND
748	165536	registry serve /etc/docker/registry/config.yml

Enabling userns remapping

```
/home/docker> ps ax | grep [d]ockerd  
1425 pts/0    Sl      0:00 dockerd --userns-remap=dockremap
```

```
/home/docker> grep dockremap /etc/passwd  
dockremap:x:100:101:Linux User,,,:/home/dockremap:/bin/false
```

```
/home/docker> cat /etc/subuid  
dockremap:165536:65536
```

seccomp




seccomp profiles

```
{
  "defaultAction": "SCMP_ACT_ERRNO",
  "architectures": [
    "SCMP_ARCH_X86_64",
    "SCMP_ARCH_X86",
    "SCMP_ARCH_X32"
  ],
  "syscalls": [
    {
      "name": "accept",
      "action": "SCMP_ACT_ALLOW",
      "args": []
    },
    {
      "name": "accept4",
      "action": "SCMP_ACT_ALLOW",
      "args": []
    },
  ],
}
```

How do I get it?

- You already have it!
- Default profile has been applied to containers since engine 1.10
- For custom profiles, pass **--security-opt** option on the command line.

The Iceberg (again)



STOP TALKING



DEMO TIME

memegenerator.net

GET EXCITED



IT'S DEMO TIME

memegenerator.net

TIME FOR A LIVE

DEMO



COULD GO WRONG?

memegenerator.net

Demo?

- A DIY demo is available
- <https://twitter.com/pvn> (it will be the pinned tweet)
- If you're reading this in the distant future and I've unpinned the tweet, try this URL instead:
<https://github.com/pvnovarese/2016-08-ContainerCon-Berlin/blob/master/README.md>

Further Reading, References, etc

- The definitive presentation on userns support:
https://events.linuxfoundation.org/sites/events/files/slides/User%20Namespaces%20-%20ContainerCon%202015%20-%202016-9-final_0.pdf
- Default seccomp profile:
<https://github.com/docker/docker/blob/master/profiles/seccomp/default.json>
- Seccomp docs:
<https://github.com/docker/docker/blob/master/docs/security/seccomp.md>
- Security non-events:
<https://docs.docker.com/engine/security/non-events/>
- Gartner Report: How to Secure Docker Containers in Operation
<https://www.gartner.com/doc/3375717/secure-docker-containers-operation>
- Your Software is Safer in Docker Containers:
<https://blog.docker.com/2016/08/software-security-docker-containers/>

Booth D38 @ LinuxCon + ContainerCon

Tues Oct 4th

- Build Distributed Systems without Docker, using Docker Plumbing Projects - Patrick Chanezon, David Chung and Captain Phil Estes
- Getting Started with Docker Services - Mike Goelzer
- Swarmkit: Docker's Simplified Model for Complex Orchestration - Stephen Day
- User Namespace and Seccomp Support in Docker Engine - Paul Novarese
- Build Efficient Parallel Testing Systems with Docker - Docker Captain Laura Frank

Wed Oct 5th

- How Secure is your Container? A Docker Engine Security Update - Phil Estes
- Docker Orchestration: Beyond the Basics - Aaron Lehmann
- When the Going gets Tough, get TUF Going - Riyaz Faizullahoy and Lily Guo

Thurs Oct 6th

- Orchestrating Linux Containers while Tolerating Failures - Drew Erny
- Unikernels: When you Should and When you Shouldn't - Amir Chaudhry
- **Berlin Docker Meetup**

Friday Oct 7th

- Tutorial: Comparing Container Orchestration Tools - Neependra Khare
- Tutorial: Orchestrate Containers in Production at Scale with Docker Swarm - Jerome Petazzoni



docker



docker

Photo credits (all creative commons licensed)

- Iceberg <http://www.ecoscope.com/iceberg/>
- Horses <https://www.flickr.com/photos/arthurtlabar/4275756092/>
- Catan <https://www.flickr.com/photos/bods/6120445526/>
- Workbench <https://www.flickr.com/photos/toollibrary/14427641289>
- memegenerator.net obv