Israeli Conference on Hardware and Side-Channel Attacks (ICHSA) 2021.

Cross-Container Linux Page Cache Attacks

Novak Boškov[1], Trishita Tiwari[2], Ari Trachtenberg[1], and David Starobinski[1]
[1]Boston University, [2]Cornell University

{boskov, trachten, staro}@bu.edu, tt544@cornell.edu

Container-based Virtualization

1. Package software with all its dependencies
2. Maximize performance
   - Save on boot time, memory usage, amount of storage, etc.
   - Make orchestration easier (service scaling, migration, etc.)

How do containers share storage?

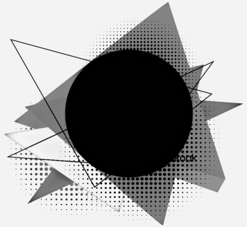

Mainly using layered (CoW) file systems.

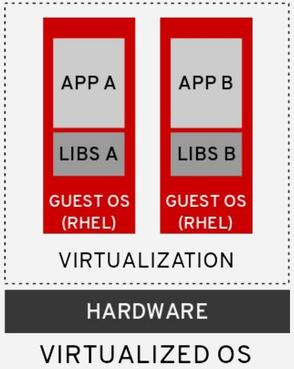

- **Multiple** OS kernel instances
- **Multiple instances** of shared libraries
- Share hardware but **NOT** OS
- Isolate (mainly) on **hypervisor** and **hardware** levels

- **Single** OS kernel
- (potentially) **Single** instance of shared libraries
- Share OS **AND** hardware
- Isolate (mainly) on **OS** level

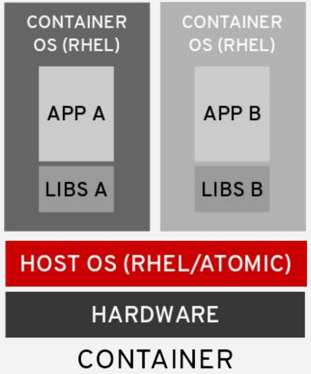

https://www.redhat.com/en/blog/virtual-machines-or-containers-maybe-both

CoW FS (e.g., **overlayfs**) as Root FS

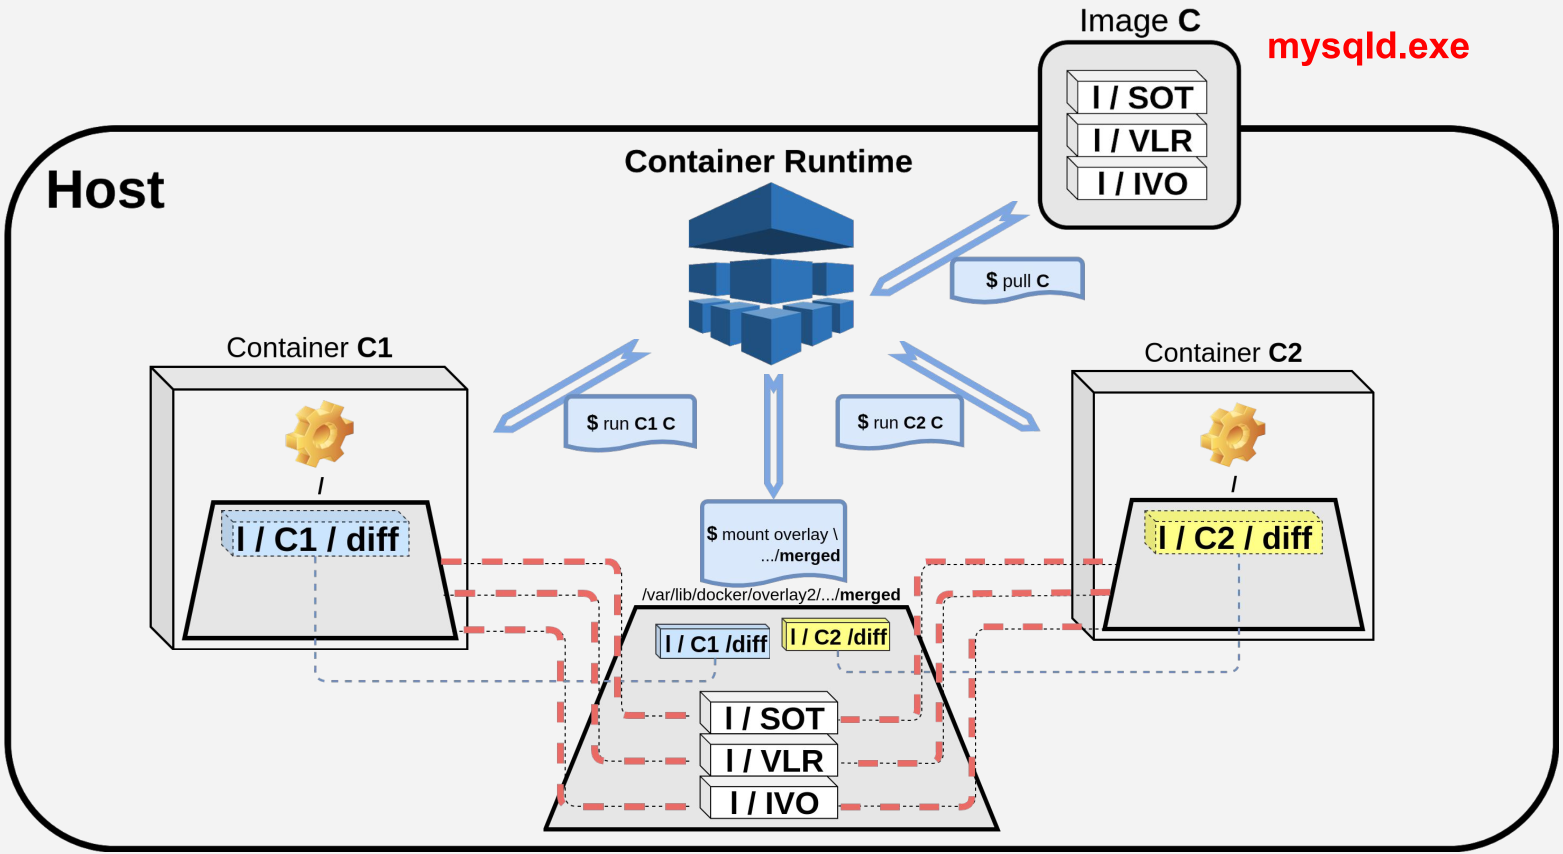

Containers Reuse The Page Cache For Shared Files

1. Page Cache is indexed using **inode** objects
2. Both containers point to the same **inode**, despite using different mount targets

```
C1$ ls -i /.../overlay2/abc/merged/mysqld
2497575

C2$ ls -i /.../overlay2/1b3/merged/mysqld
2497575
```

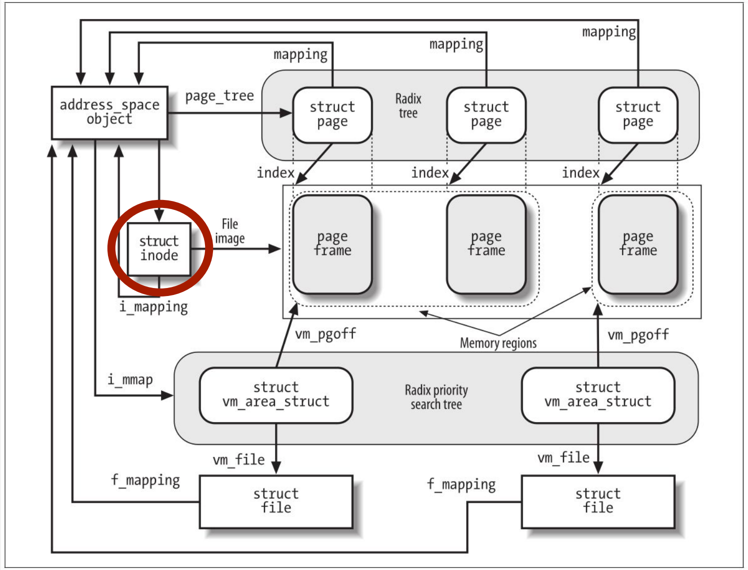

Figure 16-2. Data structures for file memory mapping

"Understanding the Linux Kernel", Bovet & Cesati

Our Attack

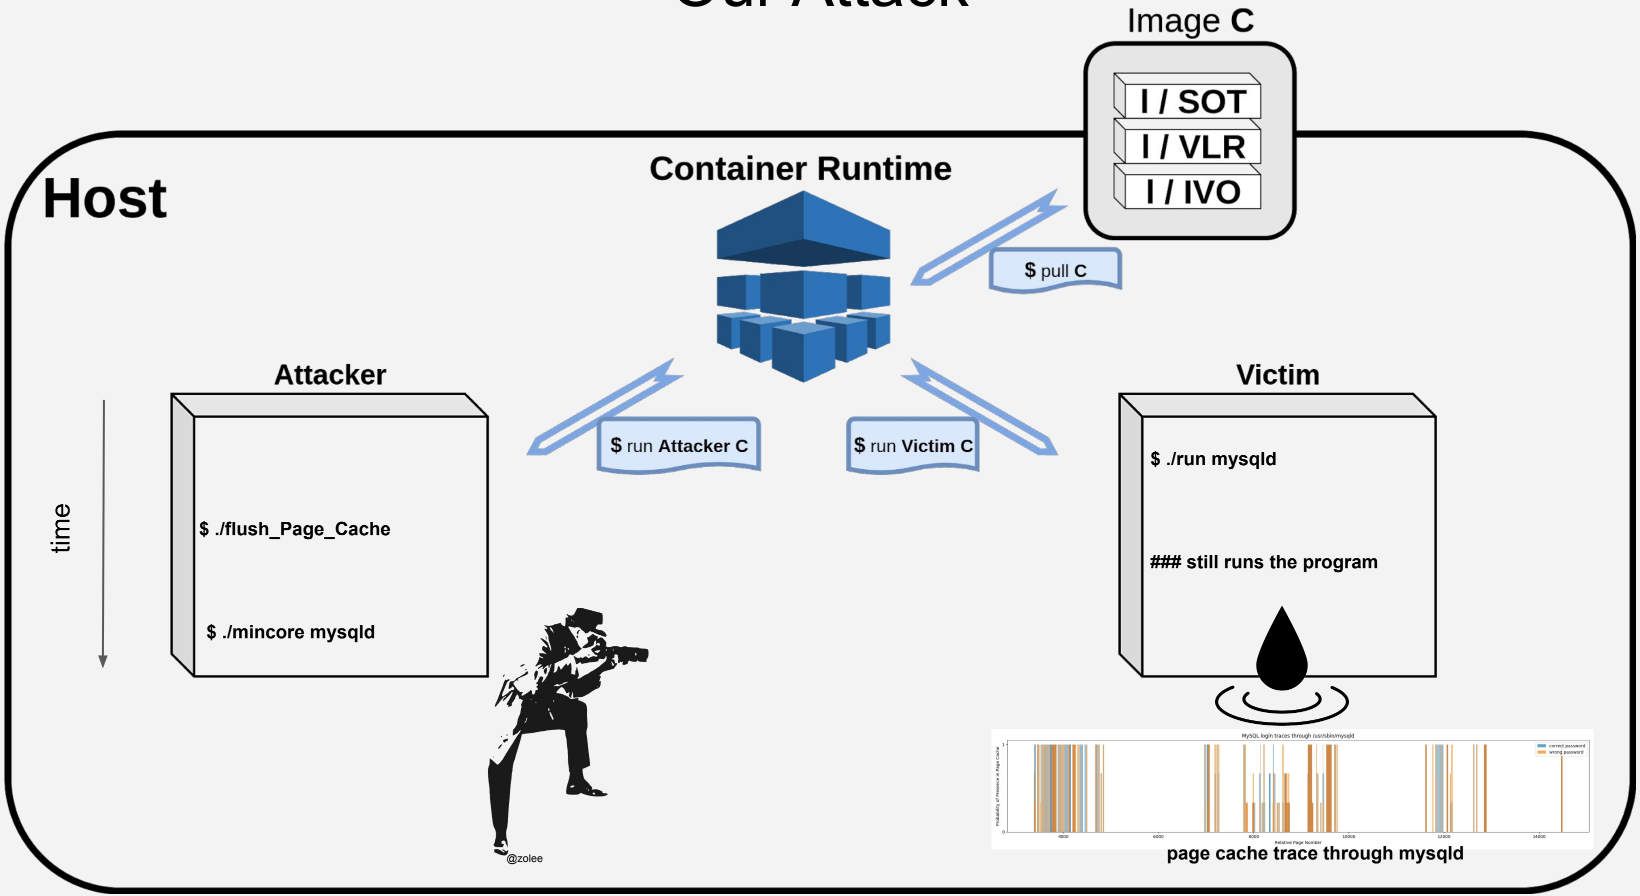

Image **C**

I / SOT
I / VLR
I / IVO

Host

Container Runtime

```
$ pull C
```

Attacker

```
$ run Attacker C
```

Victim

```
$ run Victim C
```

time

```
$ ./flush_Page_Cache

$ ./mincore mysqld
```

```
$ ./run mysqld

### still runs the program
```

page cache trace through mysqld

What Traces Look Like?

MySQL login with the correct password **VS** MySQL failed login due to wrong password

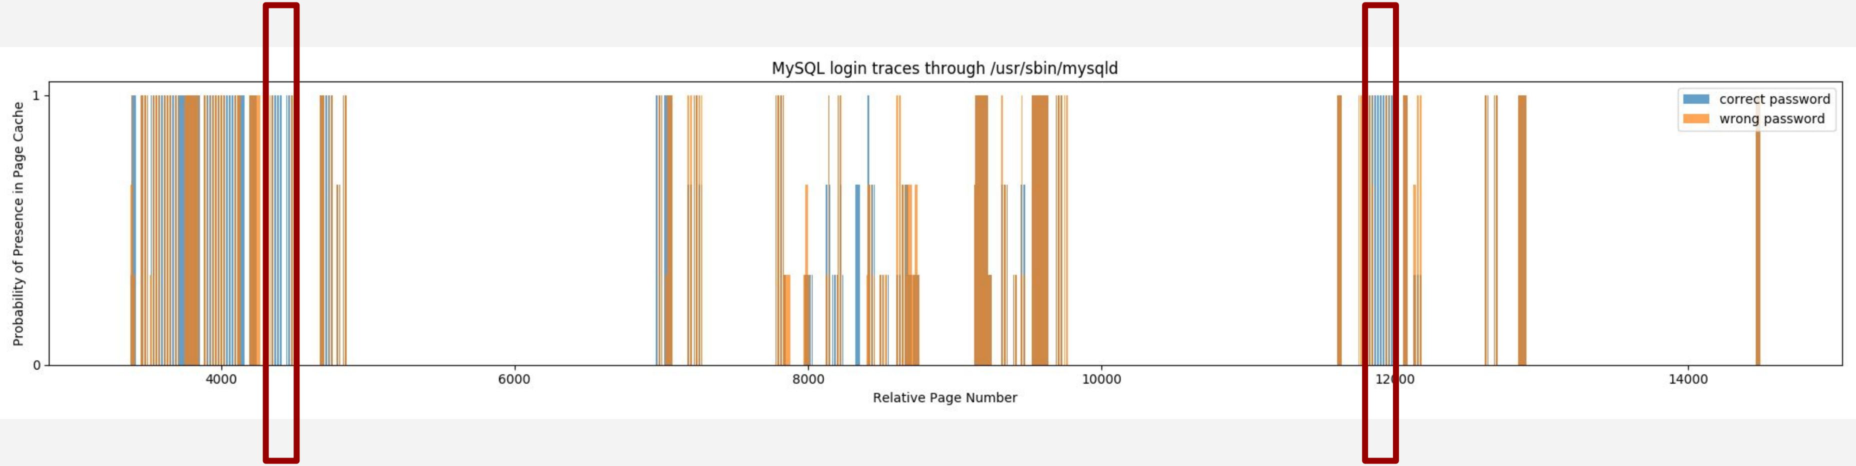

<u>many</u> different pages in the page cache

What Traces Look Like?

Wordpress user changes password **VS** Wordpress user updates content of a post

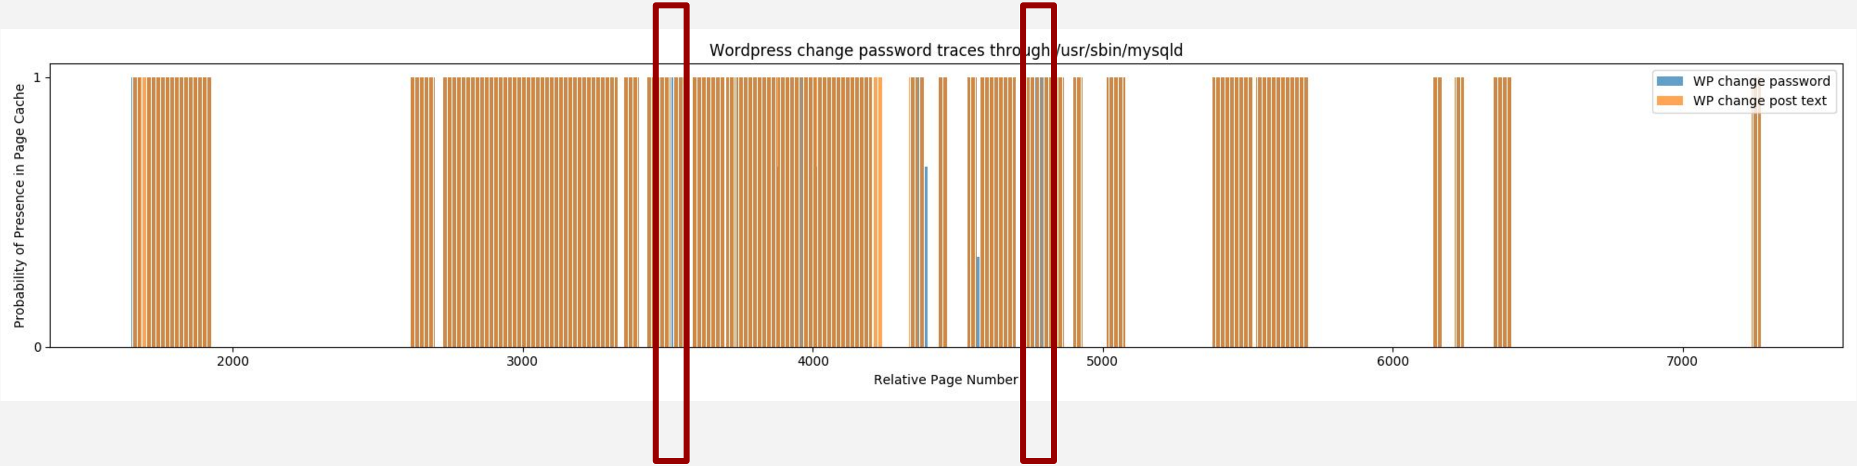

only ***a few*** different pages

Thank you!