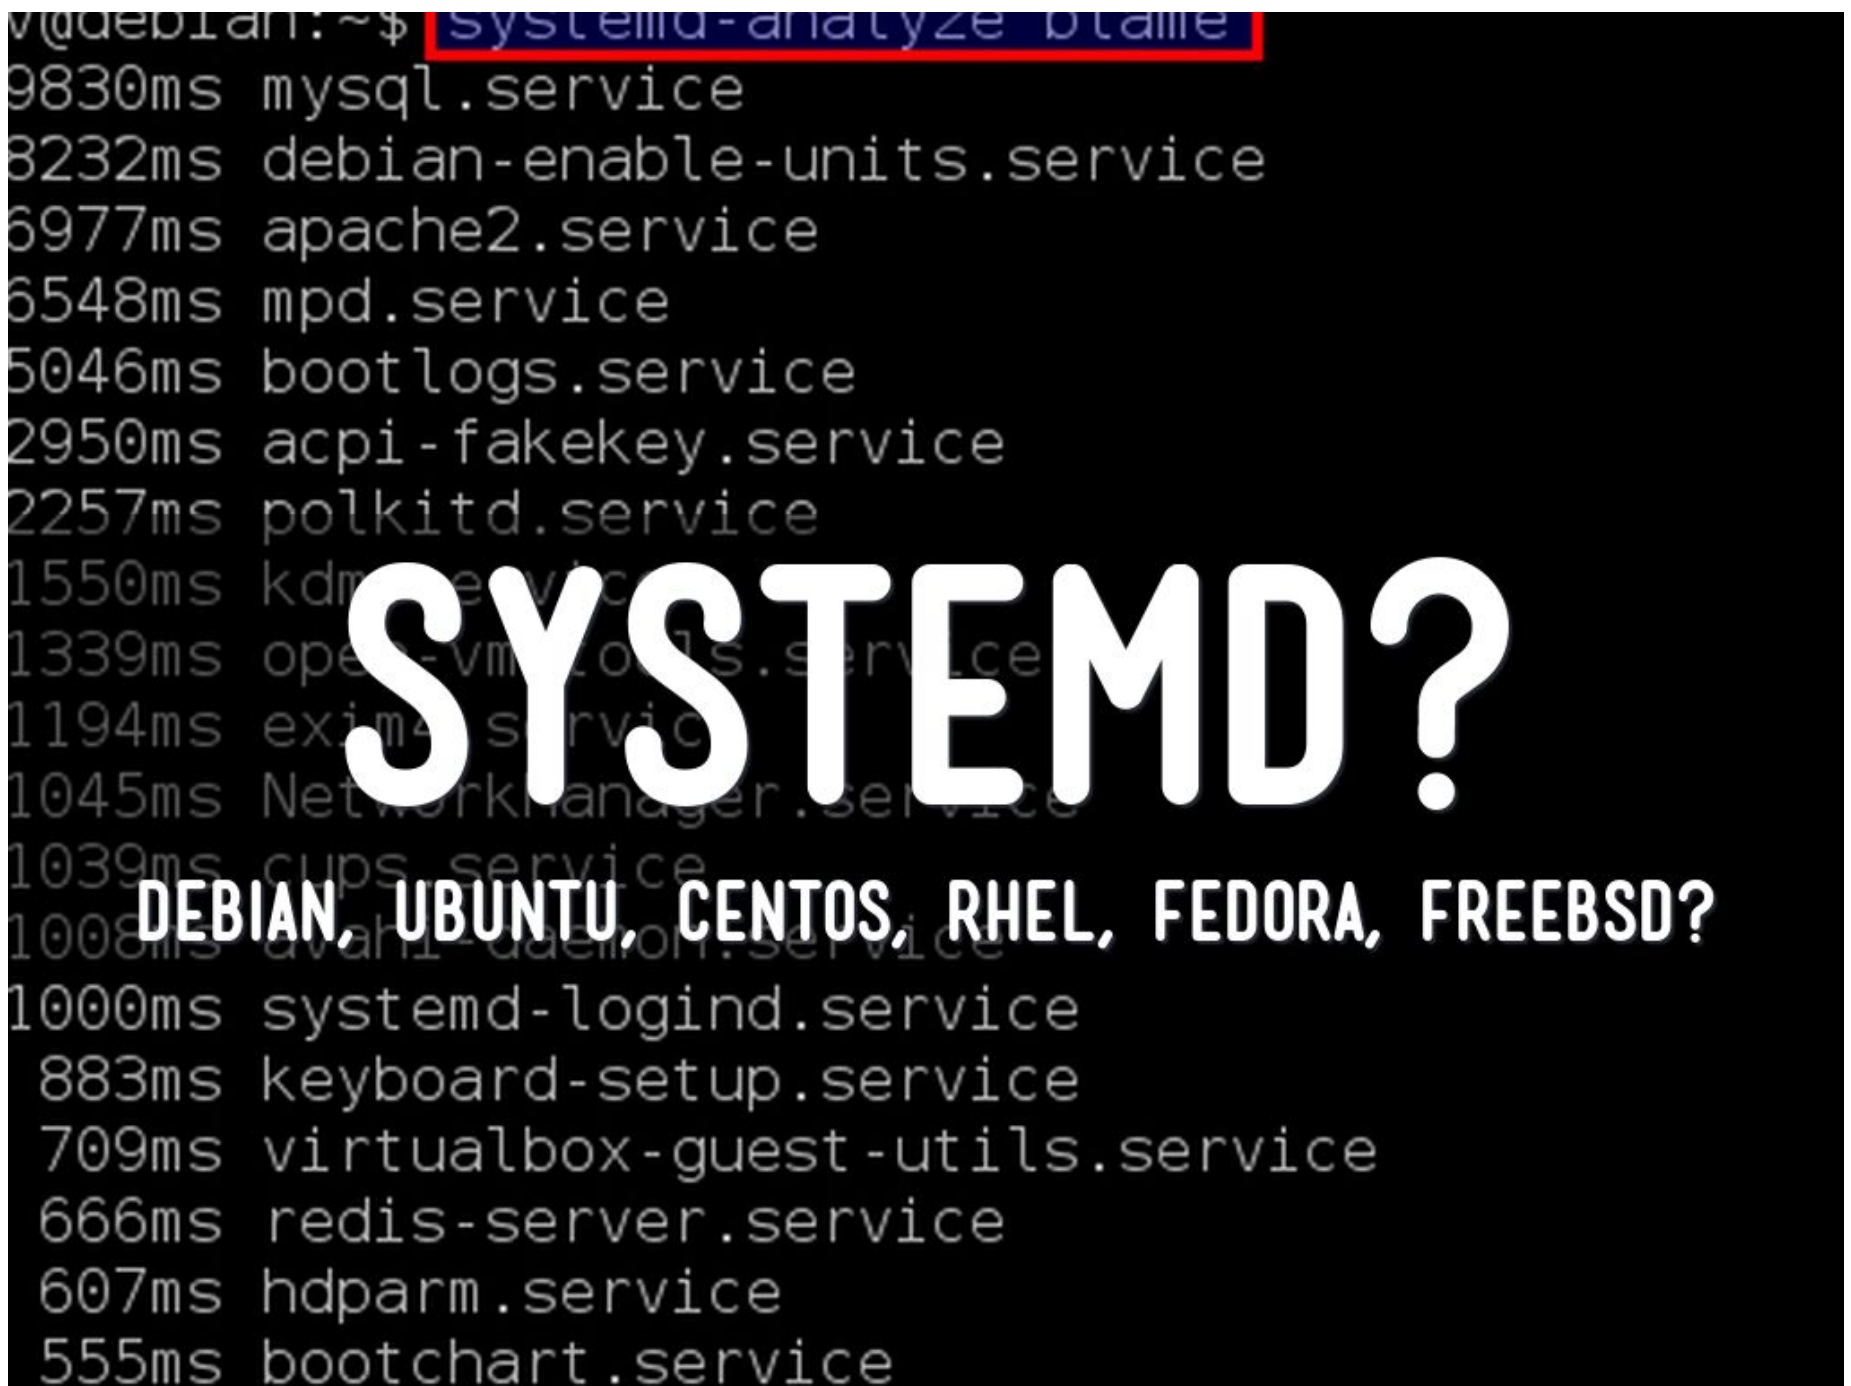






WHY CONTAINERS OR SERVICES IN PARTICULAR?







Also, not security from container
but for it!



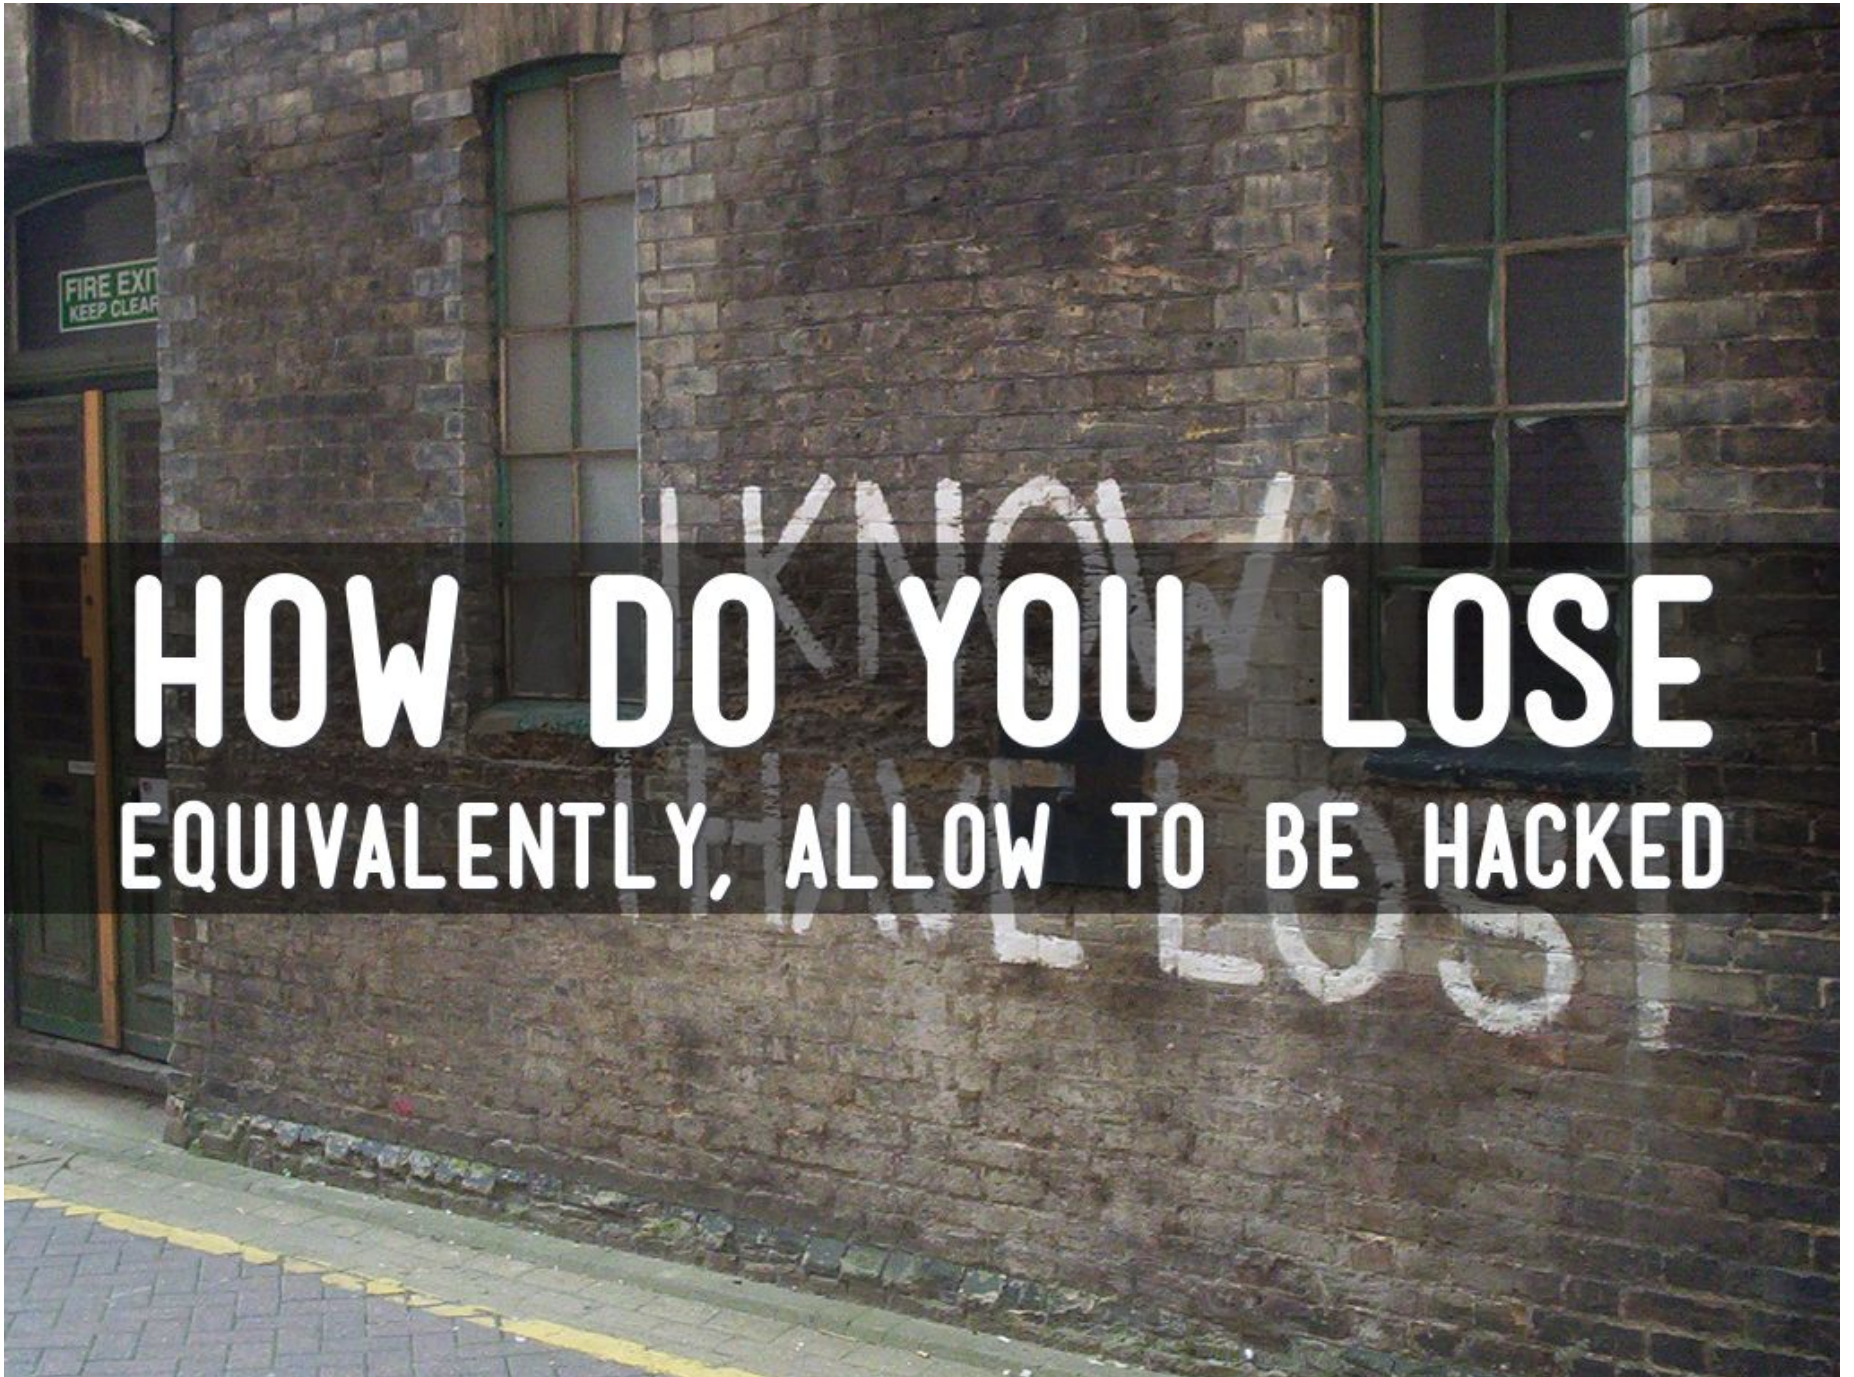
Data is the most valuable entity associated with a system, particularly when it is a sensitive one. Not only are there threats associated with physical access to the box, but also ones where logical access suffices - sql injections etc.

Entropy of data

Deduplication

THINGS TO FOCUS ON

- Malicious Peers/Env
(Break-in)
- Malicious Resident
(Break-out)
- Containment



is a sensitive one. Not only are there threats associated with physical access to the box, but also ones where logical access suffices - sql injections etc.

Vulnerabilities like shellshock and heartbleed have also shown that an exploit in one component can also be used to access others through buffer overflows, memory overruns etc. and/or impact the immunity of system severely.

Cannot always blame databases here.



not if they are also corrupted


Maintain logs - journalctl


FSS



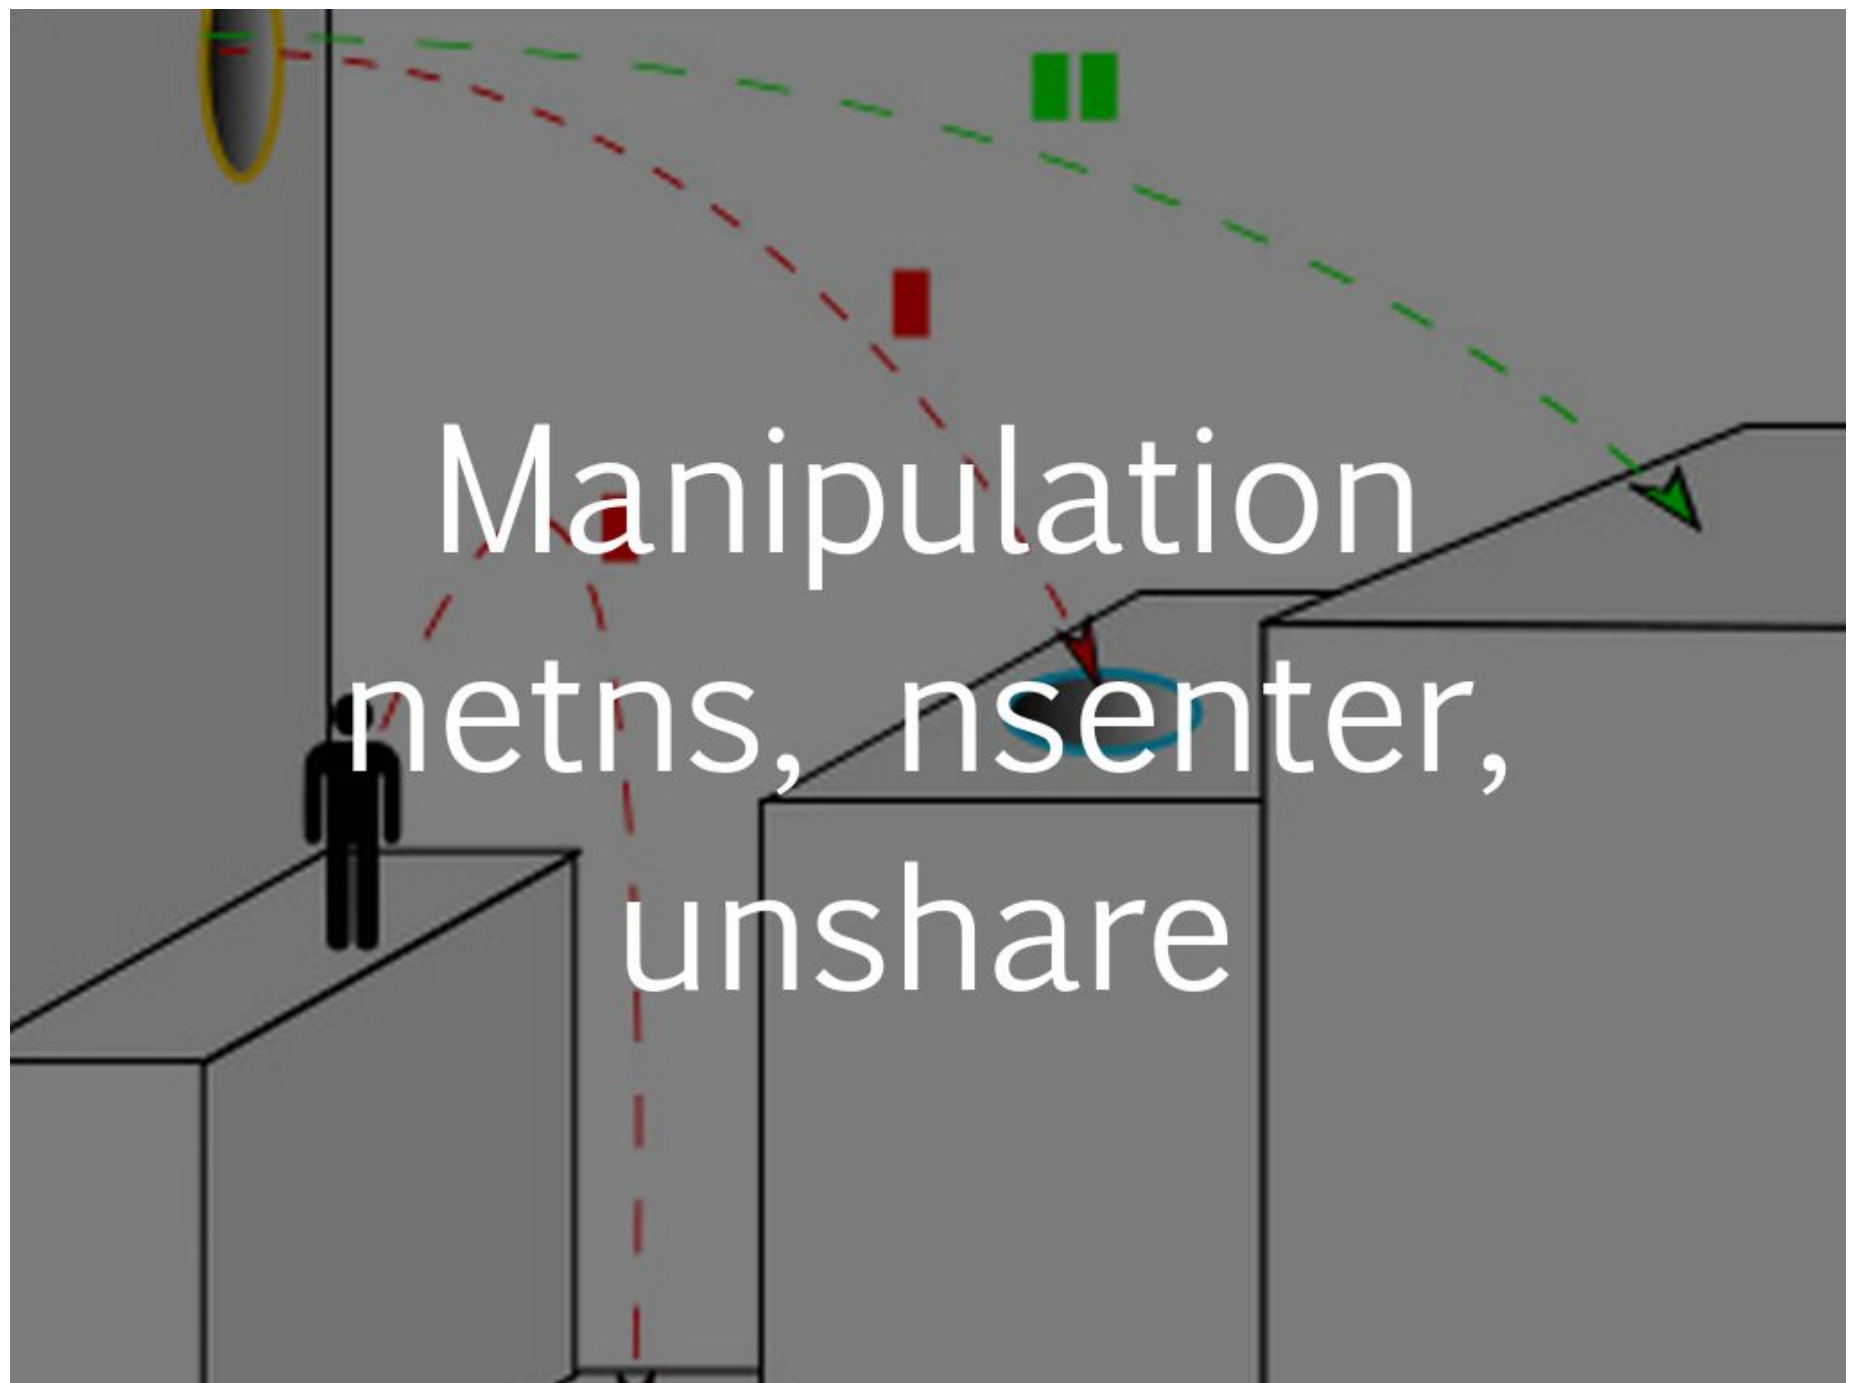
This is where "Principle of least privilege" comes into play. Wikipedia defines it as "a particular abstraction layer of a computing environment, every module (such as a process, a user or a program depending on the subject) must be able to access only the information and resources that are necessary for its legitimate purpose".

Convex Hull





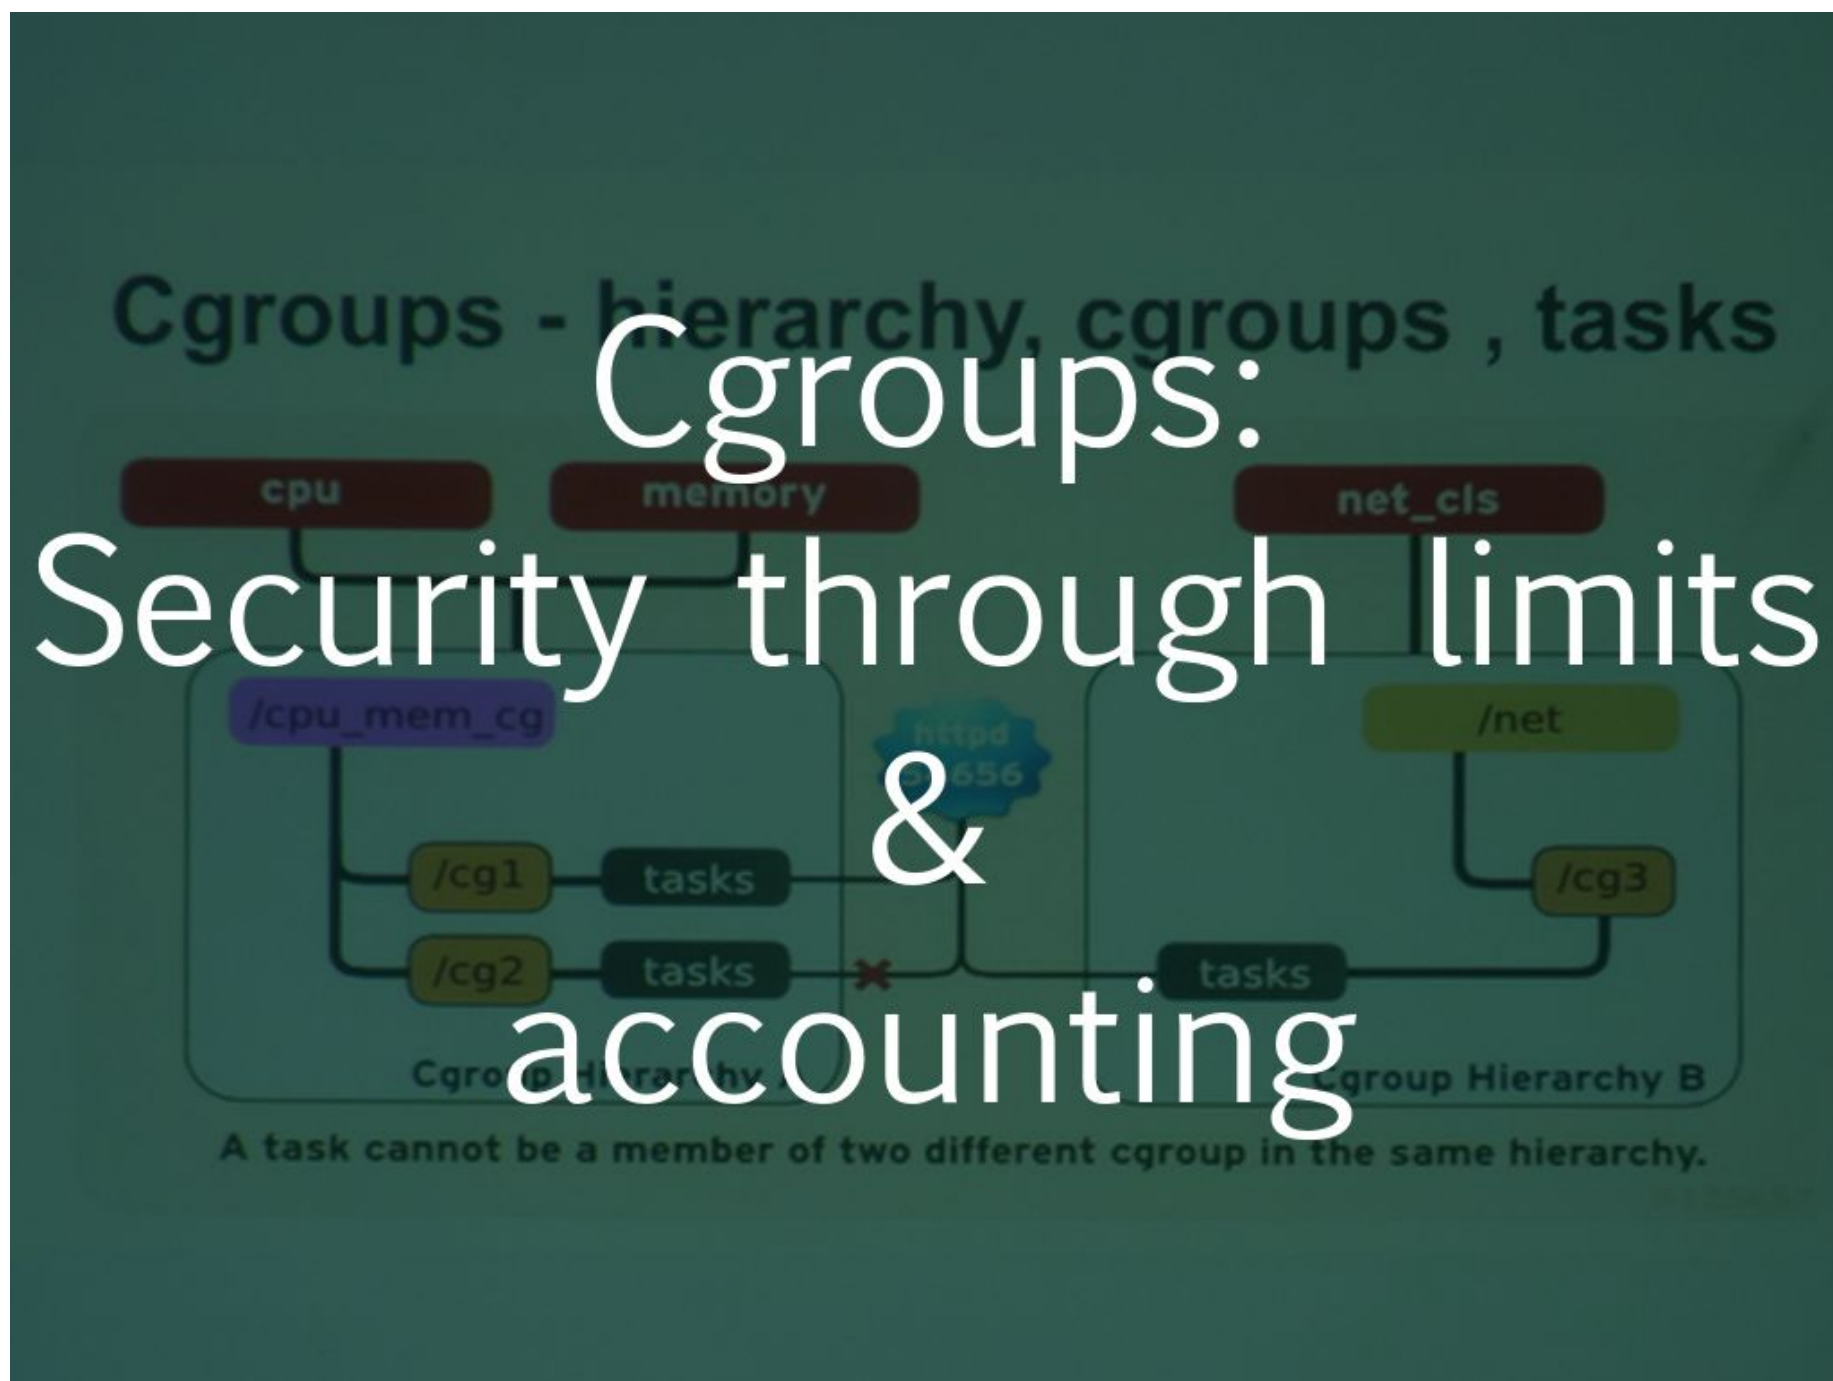
Types PID, Mount, UTS, Network, Device






Linux capabilities have been quite underused since their inception in linux. Barring CAP_SYS_ADMIN, all other capabilities allow for fine grained allocation of privileges to processes. Containers like docker already allow for a container to be run with privileges instead of a fully privileged container

CAP_SYS_ADMIN? Docker/LXC Services Vanilla



Cgroups available: cpu, memory, stats, device



Sealing key is advanced. It should not be used on multiple hosts.

/var/log/journal/e25a4e0b618f43879af033a74902d0af/fss

Please write down the following secret verification key. It should be stored in a safe location and should not be saved locally on disk.

8f41d2-ecf522-3768d7-871ce9/17ec02-35a4e900


Sealing key is automatically changed every 15min.

Keys have been generated for host foobar.example.com/e25a4e0b618f43879af033a74902d0af.


To transfer the verification key to your phone please scan the QR code below:

SECURE LOGGING: FSS

HELPS IN RCA AND RECOVERY

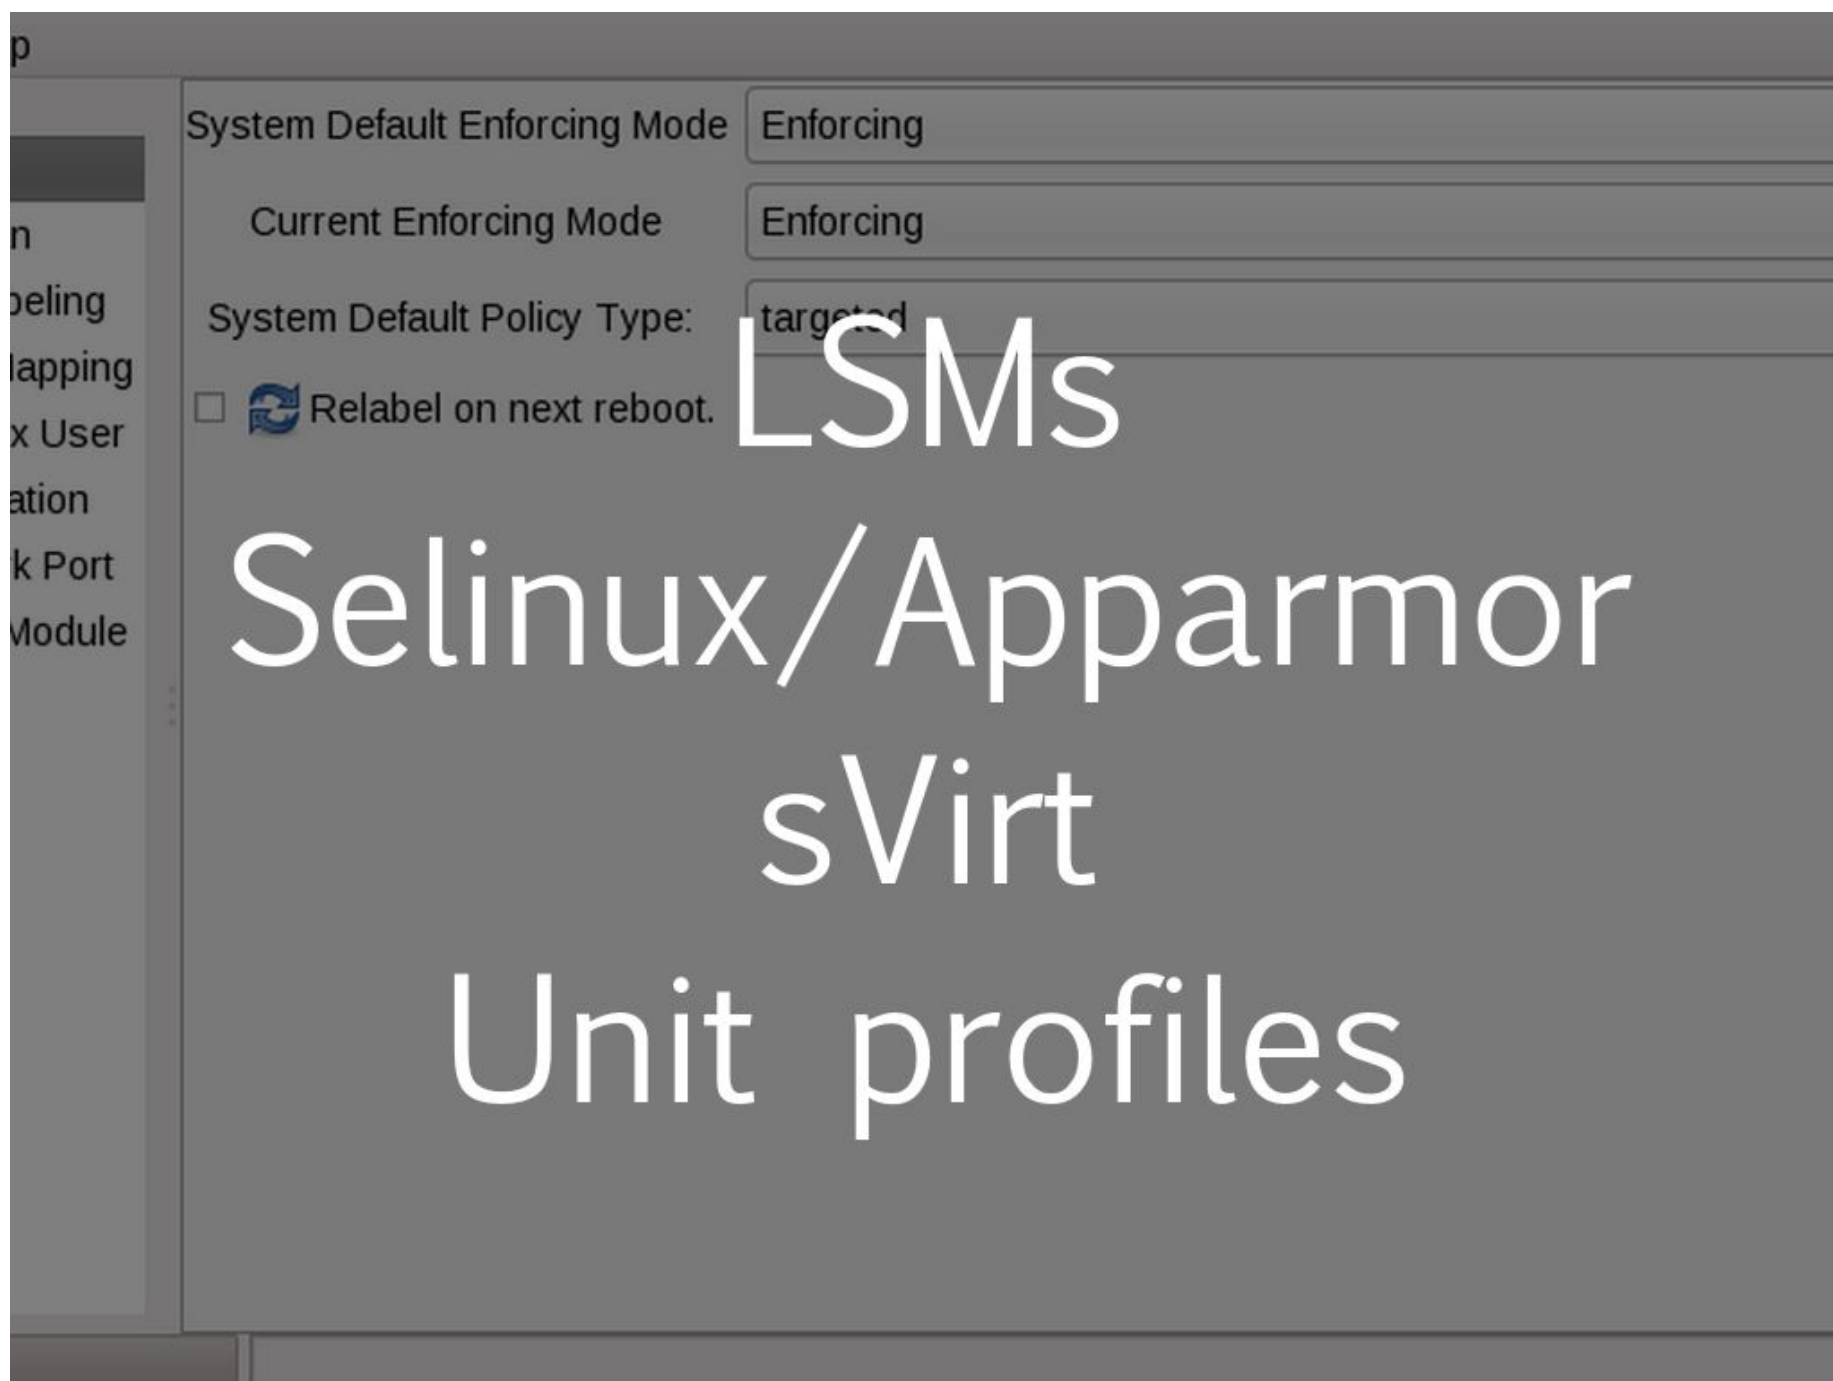






VIRTUALIZATION?

- Container(s) inside a VM?
- Not far fetched for pods
- Also,
- Zerovm - NaCl sandbox
- boot2docker





FIREWALL!

- Iptables
- Some support in docker
- and systemd - service oriented
- Conntrack and netfilter for labels
- traffic control - tc!

OTHERS

- Ulimits
- UID/GID
- ACL
- DBus
- SUID?



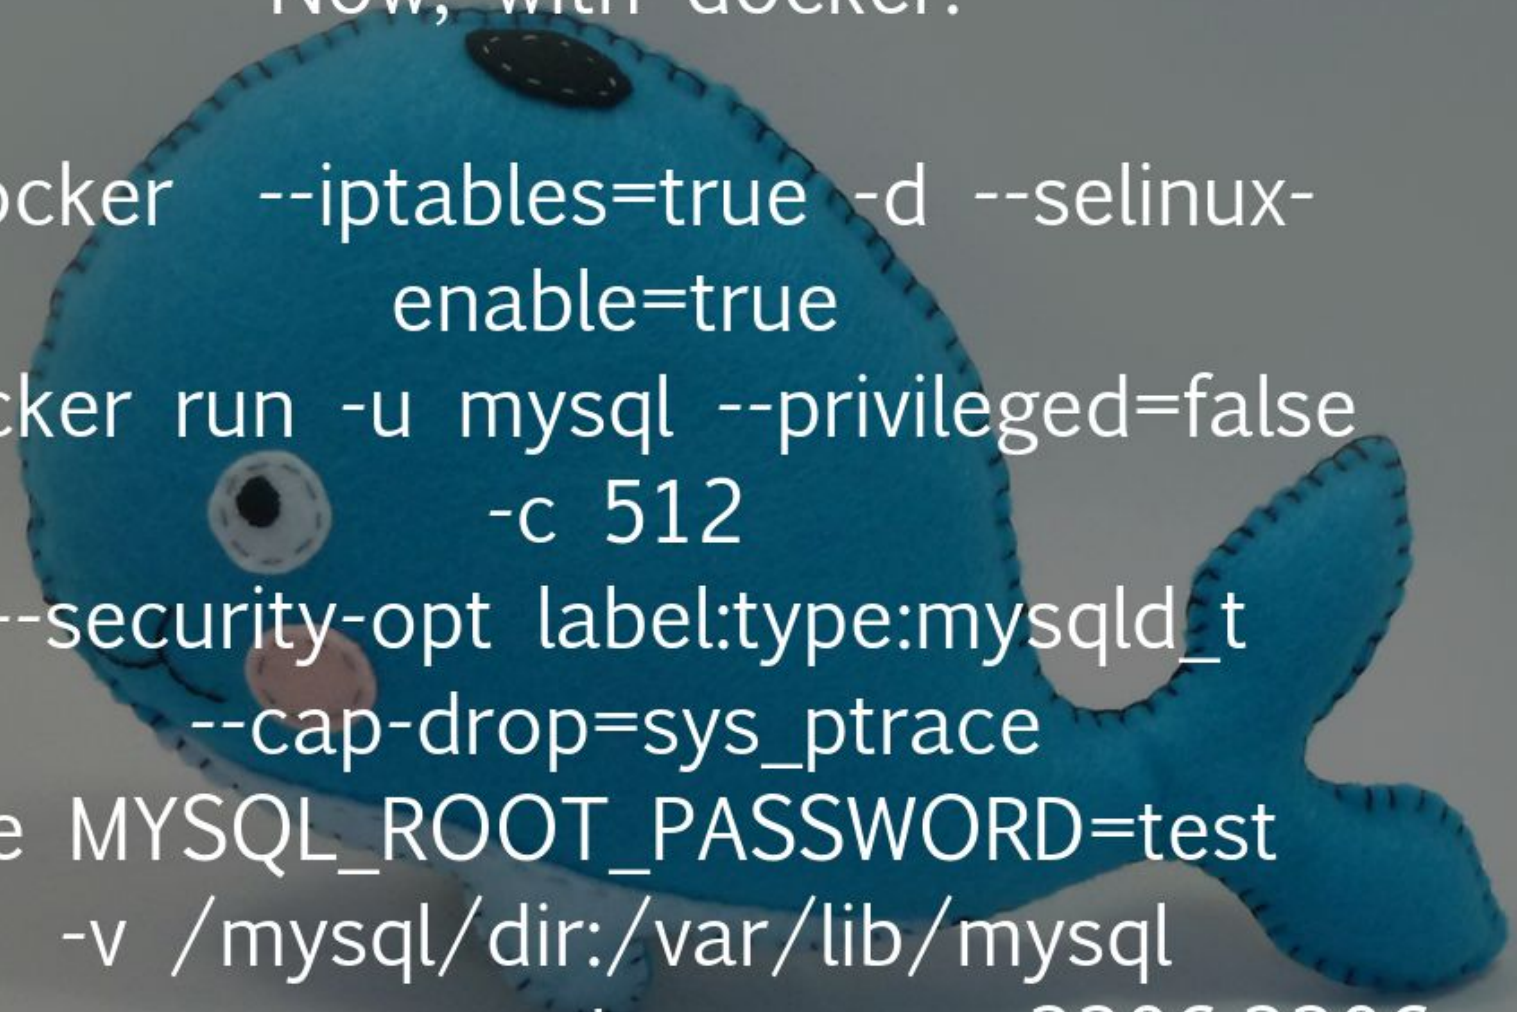



```
[Service]
EnvironmentFile=/etc/sysconfig/mysql
User=mysql
ExecStart=/usr/bin/mysqld --basedir=/usr
PrivateTmp=true
LimitNOFILE=30000
SystemCallFilter=~ioctl
NoNewPrivileges=yes
SELinuxContext=mysql_t
ProtectHome=true
PrivateDevices=true
CapabilityBoundingSet=-CAP_SYS_PTRACE
ProtectSystem=full
```



```
CPUShares=512  
MemoryLimit=2G  
MemorySoftLimit=1.8G  
BlockIOReadBandwidth=/dev/sda 1G  
ControlGroupAttribute=memory.swappiness 70
```

Now, with docker:



```
docker --iptables=true -d --selinux-  
enable=true  
docker run -u mysql --privileged=false  
-c 512  
--security-opt label:type:mysql_d_t  
--cap-drop=sys_ptrace  
-e MYSQL_ROOT_PASSWORD=test  
-v /mysql/dir:/var/lib/mysql  
-m 1g --name=mysql-server -p 3306:3306  
--net=bridge -d mysql:5.6
```




About me:
Raghavendra Prabhu
Product Lead
Percona XtraDB Cluster
@Percona
contact:
wnohang.net / rdprabhu.com

PHOTO CREDITS!

- www.incapsula.com/images/blog/images/2014-mega-vulnerabilities.png
- i1.wp.com/duffy.fedorapeople.org/blog/designs/cgroups/diagram2.png
- noflex.org/wp-content/uploads/2014/09/journaldfss.png
- upload.wikimedia.org/wikipedia/commons/f/f1/Sandbox_with_toys_detail.JPG