



Serge Hallyn

Cisco

September 14, 2017

# Problem

- Cgroups impose resource limits on services
- Services query resource availability through /proc
- Then scale per the reported availability
- /proc is system-level

# Problem

- Cgroups impose resource limits on services
- Services query resource availability through /proc
- Then scale per the reported availability
- /proc is system-level
- Solutions:
  - ▶ Virtualize /proc etc in kernel
  - ▶ Update user-space to consider cgroup limits etc
  - ▶ Libresource (<http://github.com/lxc/libresource>)
  - ▶ LXCFS (<http://github.com/hallyn/lxcfs>)

- Overmounts certain /proc files:
  - ▶ /proc/cpuinfo
  - ▶ /proc/diskstats
  - ▶ /proc/meminfo
  - ▶ /proc/stat
  - ▶ /proc/swaps
  - ▶ /proc/uptime
- Used to virtualize cgroupfs (before cgroup namespaces)

- Overmounts certain /proc files:
  - ▶ /proc/cpuinfo
  - ▶ /proc/diskstats
  - ▶ /proc/meminfo
  - ▶ /proc/stat
  - ▶ /proc/swaps
  - ▶ /proc/uptime
- Used to virtualize cgroupfs (before cgroup namespaces)
- Unimplemented features:
  - ▶ btime
  - ▶ loadavg
  - ▶ /sys/devices/system/cpu Open question: If container has cpus 0 and 3, should we show cpu0 and cpu1, or cpu0 and cpu3?
  - ▶ sysinfo (out of scope)

# Questions/Comments?

- Join! :)
- [shallyn@cisco.com](mailto:shallyn@cisco.com)
- [serge@hallyn.com](mailto:serge@hallyn.com)