DPDK

DATA PLANE DEVELOPMENT KIT

Don't fear uid != 0

Aaron Conole <aconole@redhat.com>

DPDK Userspace - Dublin – 2018

Background - history of root use

- DPDK accesses device memory directly
  - Queues
  - CSRs
  - IVs
- DPDK takes lots of resources
  - PMDs are greedy
  - Lots of hugepage memory

- DPDK developers typically have single user machines
  - Why not just run as root?
  - Security restrictions = annoying

Background - Customers

- Telco / NFVi
  - Different VNF providers share systems
  - Sensitive data - phone calls are **private**
  - Mandates about what services get what privileges

- Non-telco
  - Less widely deployed - definitely full of multi-tenancy
  - Container-based systems will definitely prefer non-root

- **Does DPDK really need to be unrestricted?**

The case for 'root'

- DPDK applications are moving the traffic anyway
  - MitM by default
  - Injection of arbitrary traffic
- DPDK applications are using isolated cores
  - Even if they spin for(;;) they can't kill applications, right?
- Everyone uses UIO anyway, right?

- "Who cares? You need that to get performance!"

Current practice - 'root'

- Current development and test practice is to use root
  - Doesn't translate for users
- PMDs aren't tested with VFIO
  - Setting up the system might be 'hard'
  - Strategies for configuration are more complex
- Developers assume certain deployment models

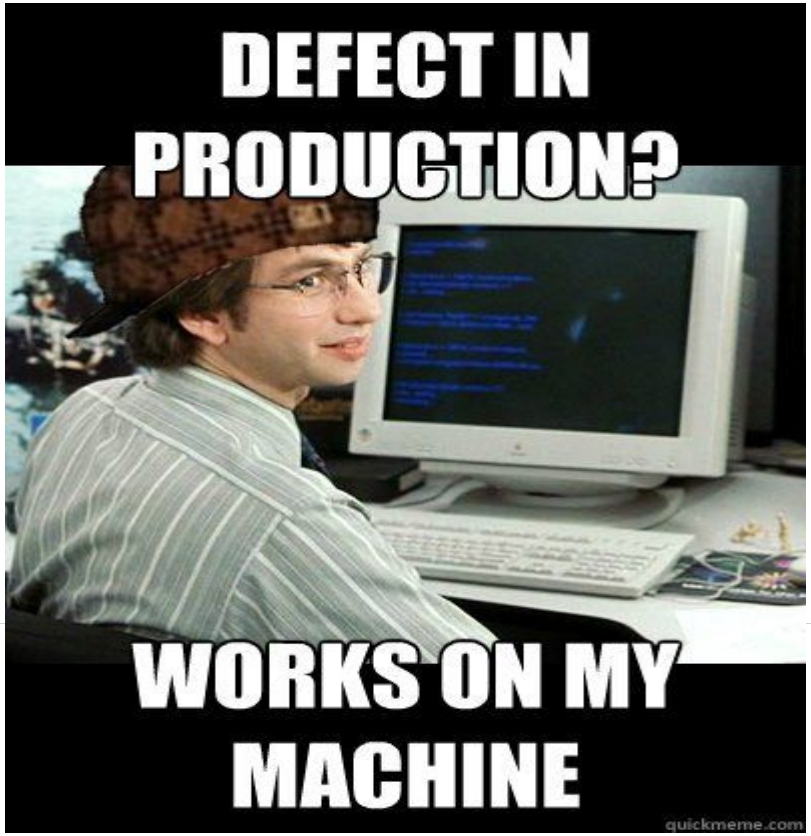

Networks are untrusted by default

- Any leg of the network that touches the internet should be considered untrustworthy
- Some VNF images treat all traffic as untrusted, anyway
- Well known attack space (packets in = scary)
- Root vs. non-root has no impact on this

DPDK has incidental access

- Why should OvS have the ability to write to VM images?
- Or the ability to change configuration apart from OVS related configuration?
- Insert any other application in there - it is the same story

SELinux and DPDK

- They do play nicely, contrary to popular belief
- Make SELinux Enforcing Again!
- SELinux can even help enforce additional fine-grained policies
- No measurable / discernable performance impact

Telco's have mandates

- There are customer requirements to run as non-root
- Don't forget SELinux, too
- This also means that UIO doesn't get used
- VFIO and device permissions (we use hugetlbfs as a group, fyi)

Things I'd like to see - "non-root"

- Defacto set of selinux macros
  - Right now, each application will need their own
  - OvS has a fairly comprehensive suite
- VFIO as the default
  - UIO works but makes no sense with non-root (after all, you can just write to memory)
  - VFIO works with non-root and makes sense
- More documentation

Questions?

Aaron Conole <aconole@redhat.com>