AGL Virtualization EG meeting
August 2020

Argo: VirtIO

Christopher Clark & Rich Persaud
Edgeform

Argo: inter-VM data transport

- Xen feature in 4.12, upstreamed from OpenXT in 2019
- Asynchronous authenticated message passing between VMs, performed by the hypervisor
- Prioritizes isolation, access control and mandatory conformance to transfer protocol
- Simple but powerful primitive to build upon
- Host-local connectivity use cases:
  - transport between split device drivers
  - transport for data bus between distributed components: DBUS

Argo: Hypervisor-Mediated Data eXchange

- Isolation
  - No shared memory between VMs
  - Data is copied from the source to the destination
    - Maintains spatial isolation between guest VMs

- Strong mechanism
  - Hypervisor performs the data movement: ensures that memory accesses conform to protocol

- Enforcement of Mandatory Access Control
  - Hypervisor performs the data movement: permissions are enforced

VirtIO: “a de-facto standard for virtual I/O devices”

- Attractive: standardized, widely deployed, documented, well tested, efficient
    - Learned lessons from Xen’s PV split device-driver model: has a familiar structure
    - Has common core transports and data structures that significantly reduce work to implement each new virtual device driver

- Challenge: Isolation
    - Expectation: front-end memory buffers are accessible to back-end drivers

- Challenge: Access Control
    - Data transport performed via shared memory between VMs

VirtIO: a possible path forwards with Argo

- Leverage VirtIO’s transport abstraction + VirtIO’s DMA buffer handling
    - Enables introduction of a virtio-argo transport device driver in front-end
        - Virtqueues and vrings supported
            - Compatibility with the existing VirtIO split device drivers
    - Can co-exist with other existing transports in same system if wanted

- Alternative to virtio-pci and virtual PCI devices
    - Add new device discovery method: via ACPI tables
    - A new I/O path via Argo

- Enables enforcement of Mandatory Access Control over VirtIO data flows
    - Leverages existing hypervisor MAC (XSM) and Argo firewall mechanisms
    - Supports strong isolation between communicating VMs

VirtIO: a possible path forwards with Argo

- Simple transition for VMs using VirtIO
    - Adds a single driver for the guest
        - primary interfaces are within kernel to VirtIO, and ACPI to platform
        - candidate for upstreaming to mainline Linux

- Back-end driver and userspace support:
        - toolstack
        - QEMU virtio-argo driver
        - libargo
        - Linux Argo driver

Argo: Pointers

- Argo and Hypervisor-Mediated Data eXchange (HMX)
    - Xen Design & Developer Summit 2019
        - https://static.sched.com/hosted_files/xensummit19/92/Argo%20and%20HMX%20-%20OpenXT%20-%20Christopher%20Clark%20-%20Xen%20Summit%202019.pdf
    - Platform Security Summit 2018
        - https://www.platformsecuritysummit.com/2018/speaker/clark/
- Argo in Xen
    - Design Document: https://xenbits.xen.org/docs/unstable/designs/argo.html
    - Wiki: https://wiki.xenproject.org/wiki/Argo:_Hypervisor-Mediated_Exchange_(HMX)_for_Xen
    - Interface: https://xenbits.xen.org/gitweb/?p=xen.git;a=blob;f=xen/include/public/argo.h
- Argo development in OpenXT
    - https://openxt.atlassian.net/wiki/spaces/DC/pages/737345538/Argo+Hypervisor-Mediated+data+eXchange+Development
- Hypervisor Security: Lessons Learned
    - Ian Pratt, Bromium, at Platform Security Summit 2018
    - https://www.platformsecuritysummit.com/2018/speaker/pratt/
- Mandatory Access Control
    - Linux Security Summit 2019: 26 Years of Flexible MAC
    - https://lssna19.sched.com/event/RHaH/keynote-retrospective-26-years-of-flexible-mac-stephen-smalley-national-security-agency