REGIONAL SDC 25

BY Developers FOR Developers

Regional SDC Denver
April 30, 2025

CXL Ecosystem Innovation Leveraging QEMU-based Emulation

Anisa Su, Adam Manzanares
Samsung

SNIA

What is QEMU?

- Open-source emulator & virtualizer[1]
  - Can emulate peripherals

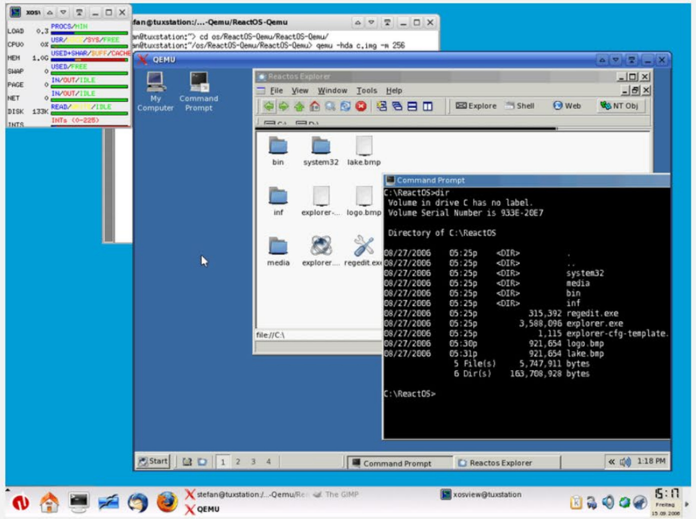

```
an@tuxstation:~> cd os/ReactOS-Qemu/ReactOS-Qemu/
an@tuxstation:~/os/ReactOS-Qemu/ReactOS-Qemu> qemu -hda c.img -m 256
```

```
C:\ReactOS>dir
 Volume in drive C has no label.
 Volume Serial Number is 933E-20E7

 Directory of C:\ReactOS

08/27/2006    05:25p    <DIR>          .
08/27/2006    05:25p    <DIR>          ..
08/27/2006    05:25p    <DIR>          system32
08/27/2006    05:25p    <DIR>          media
08/27/2006    05:25p    <DIR>          bin
08/27/2006    05:25p    <DIR>          inf
08/27/2006    05:25p           315,392 regedit.exe
08/27/2006    05:25p         3,588,096 explorer.exe
08/27/2006    05:25p             1,115 explorer-cfg-template.
08/27/2006    05:30p           921,654 logo.bmp
08/27/2006    05:31p           921,654 lake.bmp
                5 File(s)      5,747,911 bytes
                6 Dir(s)     163,708,928 bytes

C:\ReactOS>
```

Full-system emulation

Run operating systems for any machine, on any supported architecture

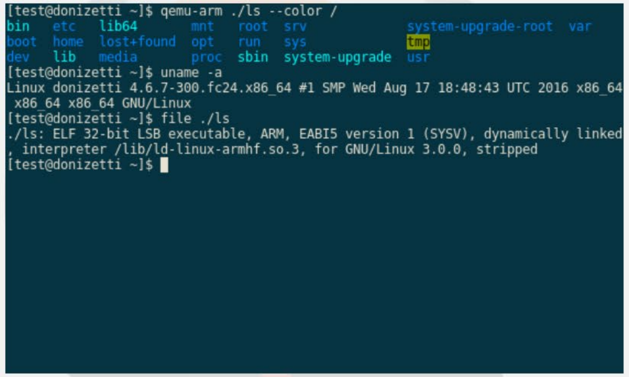

```
[test@donizetti ~]$ qemu-arm ./ls --color /
bin   etc   lib64       mnt   root  srv             system-upgrade-root  var
boot  home  lost+found  opt   run   sys             tmp
dev   lib   media       proc  sbin  system-upgrade  usr
[test@donizetti ~]$ uname -a
Linux donizetti 4.6.7-300.fc24.x86_64 #1 SMP Wed Aug 17 18:48:43 UTC 2016 x86_64
 x86_64 x86_64 GNU/Linux
[test@donizetti ~]$ file ./ls
./ls: ELF 32-bit LSB executable, ARM, EABI5 version 1 (SYSV), dynamically linked
, interpreter /lib/ld-linux-armhf.so.3, for GNU/Linux 3.0.0, stripped
[test@donizetti ~]$ 
```

User-mode emulation

Run programs for another Linux/BSD target, on any supported architecture

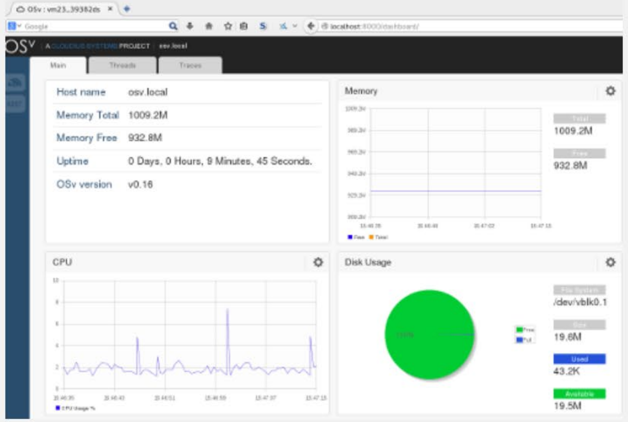

Host name osv.local
Memory Total 1009.2M
Memory Free 932.8M
Uptime 0 Days, 0 Hours, 9 Minutes, 45 Seconds.
OSv version v0.16

Virtualization

Run KVM and Xen virtual machines with near native performance

[1] https://www.qemu.org

Why is QEMU Useful?

- Great for rapidly prototyping end-to-end SW for new hardware features
  - HW/SW developer co-design
  - ex: ZNS, FDP, CXL
  - Create host software that leverages these features
- Samsung Successes – NVMe & CXL Support
  - Testing frameworks can move faster than hw availability
  - Enables the quick adoption of hardware
- Brings people into the ecosystem

Why Emulate CXL Devices?

- CXL HW Availability is limited

- Reproduce success cases for NVMe
  - Build end-to-end SW without waiting for HW samples

- Who benefits from emulation?
  - Operating Systems
    - Driver development, application prototyping
  - Management
    - BMC, Fabric Manager

Features You Can Emulate

- MCTP and Switch CCIs, MHSLD,
- Dynamic Capacity Device
  - dynamically allocate/deallocate memory to a host
- Fabric Management
  - Fabric Manager (FM): logical process that can dynamically configure the system's operational state using the FMAPI
  - FM can be on host machine, BMC, CXL device, CXL switch, etc.

Ex: FM-owned Switch

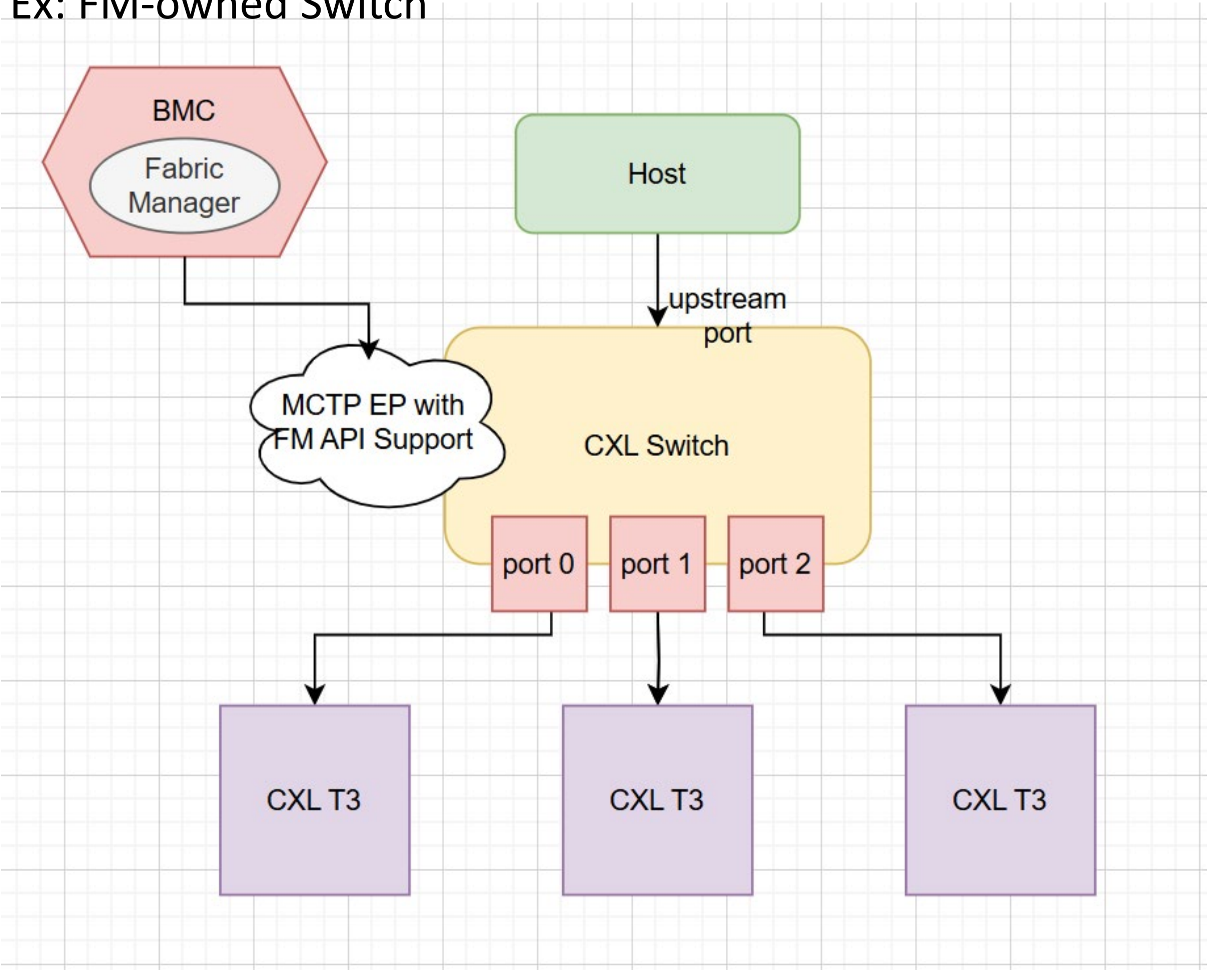

How is memory dynamically added?

1. FM initiates DC Add

2. Device adds info in event log

3.Host receives information in event log; sends device command Memdev DC Add

4. Device allocates the extents for the host

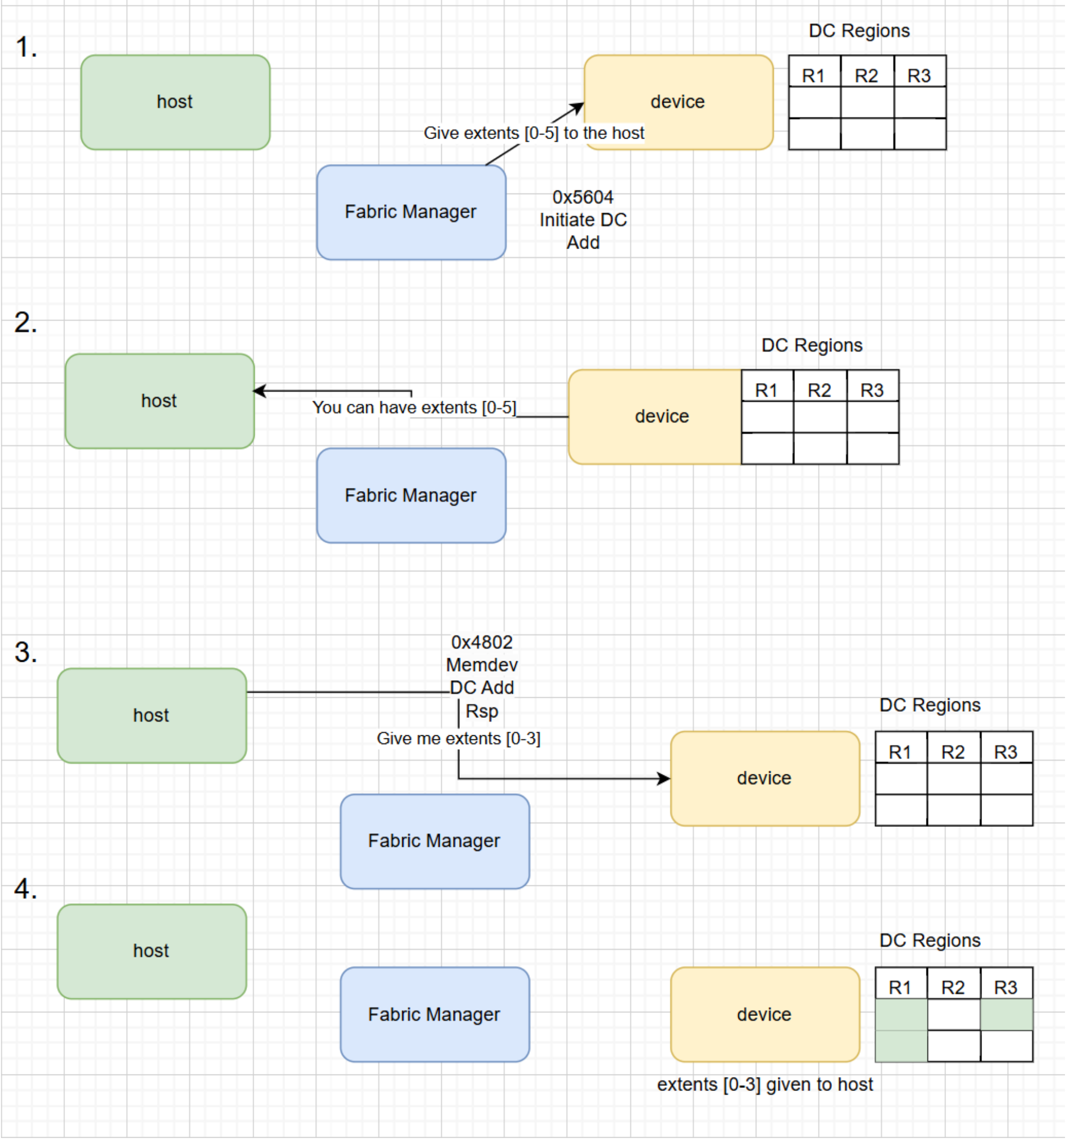

Need for libcxlmi

- Main motivation: send commands to a CXL device
- libnvme utility proven in industry
  - Basis of nvme-cli
  - nvme-cli/plugins at master · linux-nvme/nvme-cli
- libcxlmi provides similar flexibility for CXL devices:
  - userspace library for command framing
  - send any command to any CXL device
  - not tied to any specific interface to the device

Demo: Configuration

- libcxlmi acts as the FM, runs on the host

- *Need MCTP support in kernel and QEMU

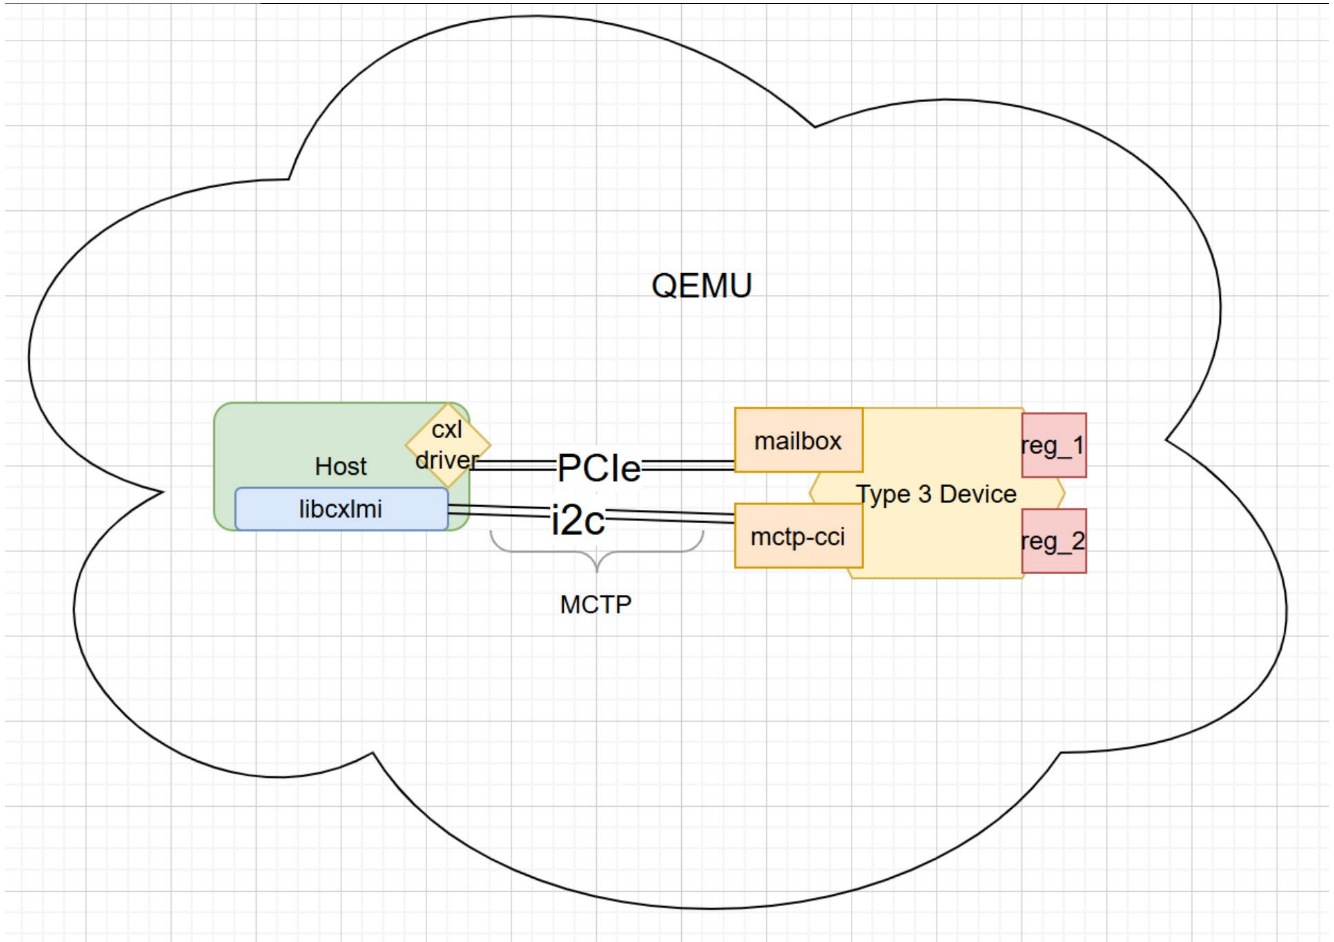

DCD Demo

Tools to Help Get Started

- Emulating CXL topology can be tricky: need some background knowledge of hardware components
    - CFMW
    - Host Bridge
    - Root Port
    - Device
- cxl-test-tool:
    - Beginner-friendly
    - Flexible for experts

Acknowledgements

- Jonathan Cameron (Huawei)
- Ira Weiny (Intel)
- Gregory Price (Meta)
- Ben Widawsky (Google)
- Davidlohr Bueso (Samsung)
- Fan Ni (Samsung)
- Many others

Thank You