SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

CXL Ecosystem Innovation Leveraging QEMU-based Emulation

Anisa Su, Adam Manzanares
Samsung Semiconductor

What is QEMU?

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

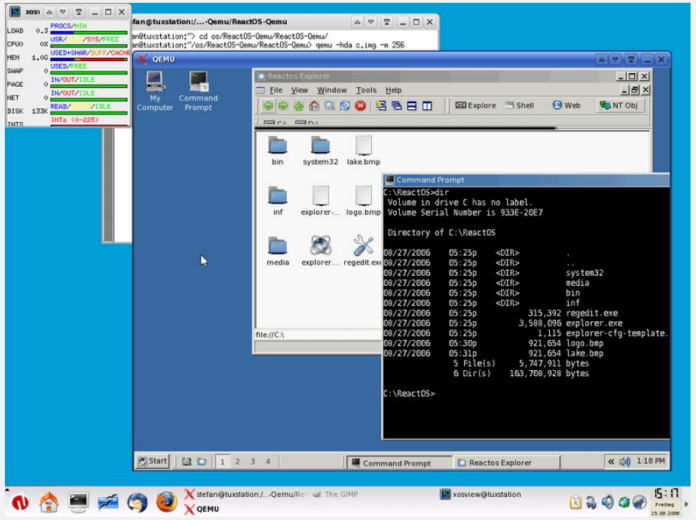

Full-system emulation

Run operating systems for any machine, on any supported architecture

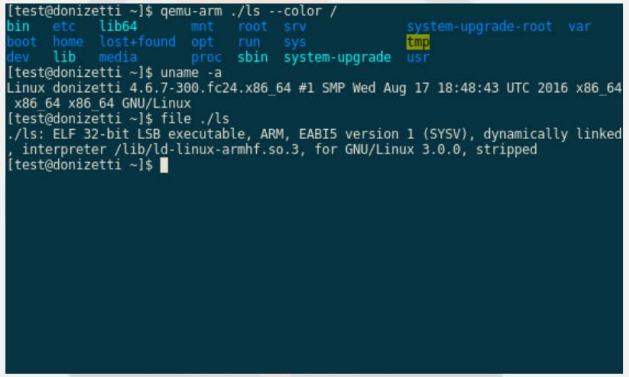

User-mode emulation

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

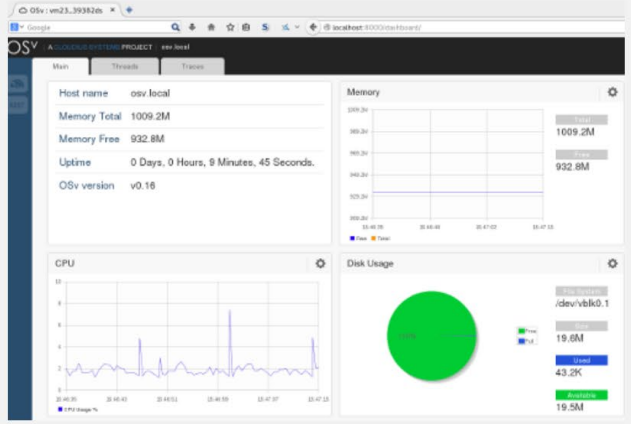

Virtualization

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?

- Next gen. HW access is limited
  - Build end-to-end SW without waiting for HW samples
- Reproduce success cases for NVMe
- Who benefits from emulation?
  - Operating Systems
    - Driver development, application prototyping
  - Management
    - BMC, Fabric Manager

Features You Can Emulate

- 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

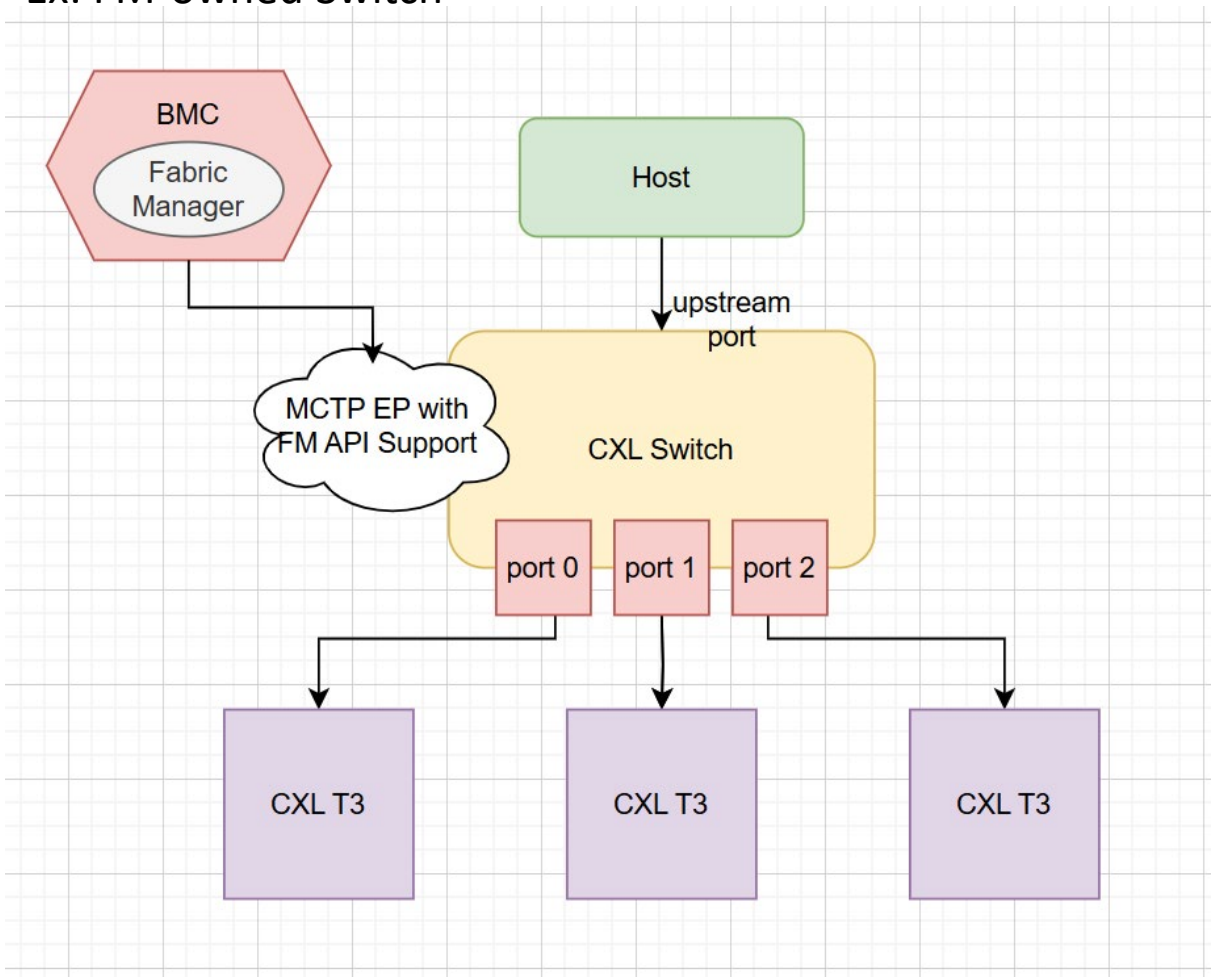

How is Memory Dynamically Added?

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

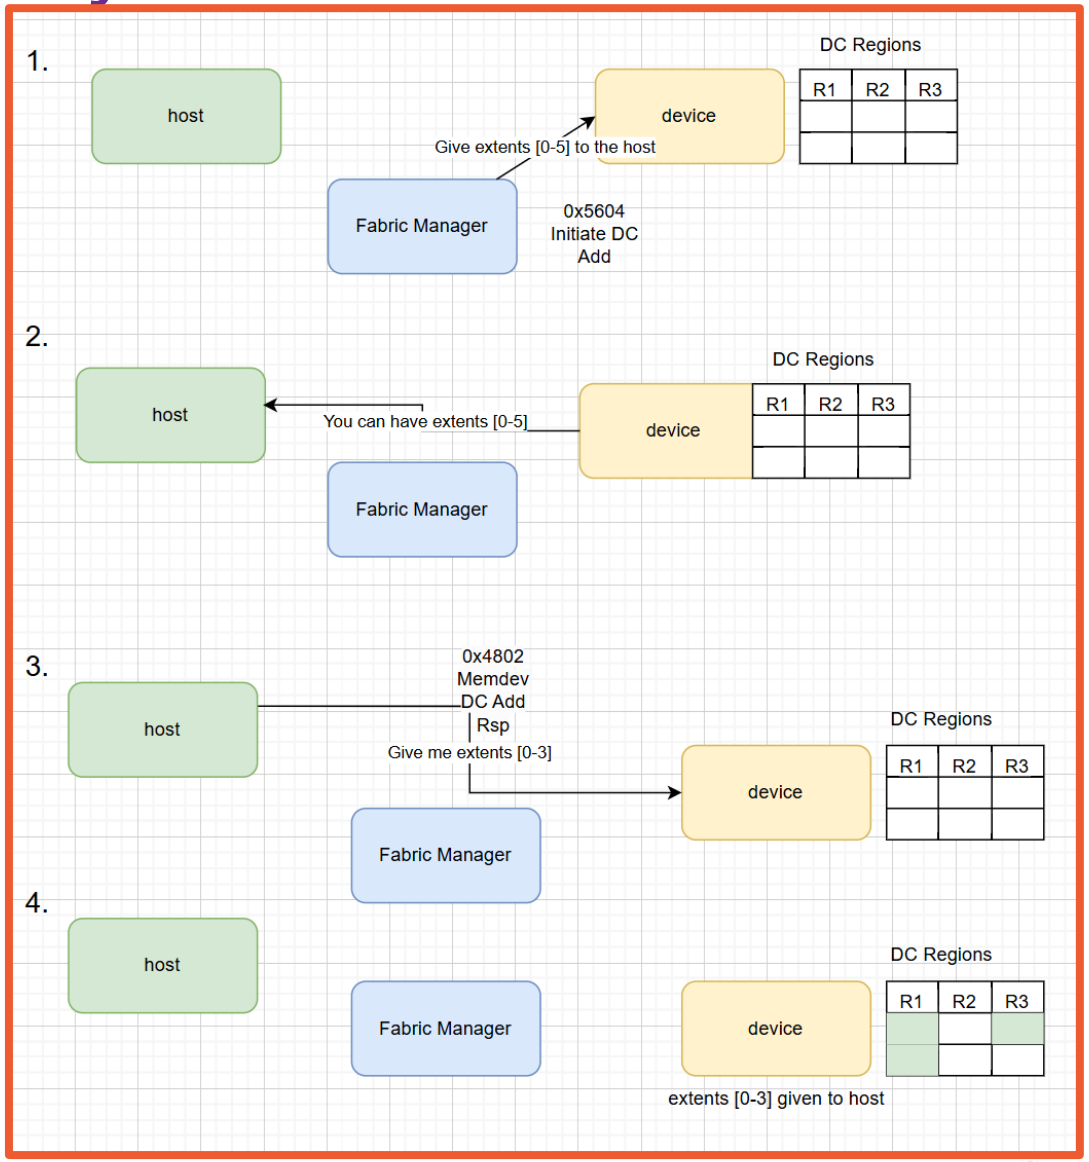

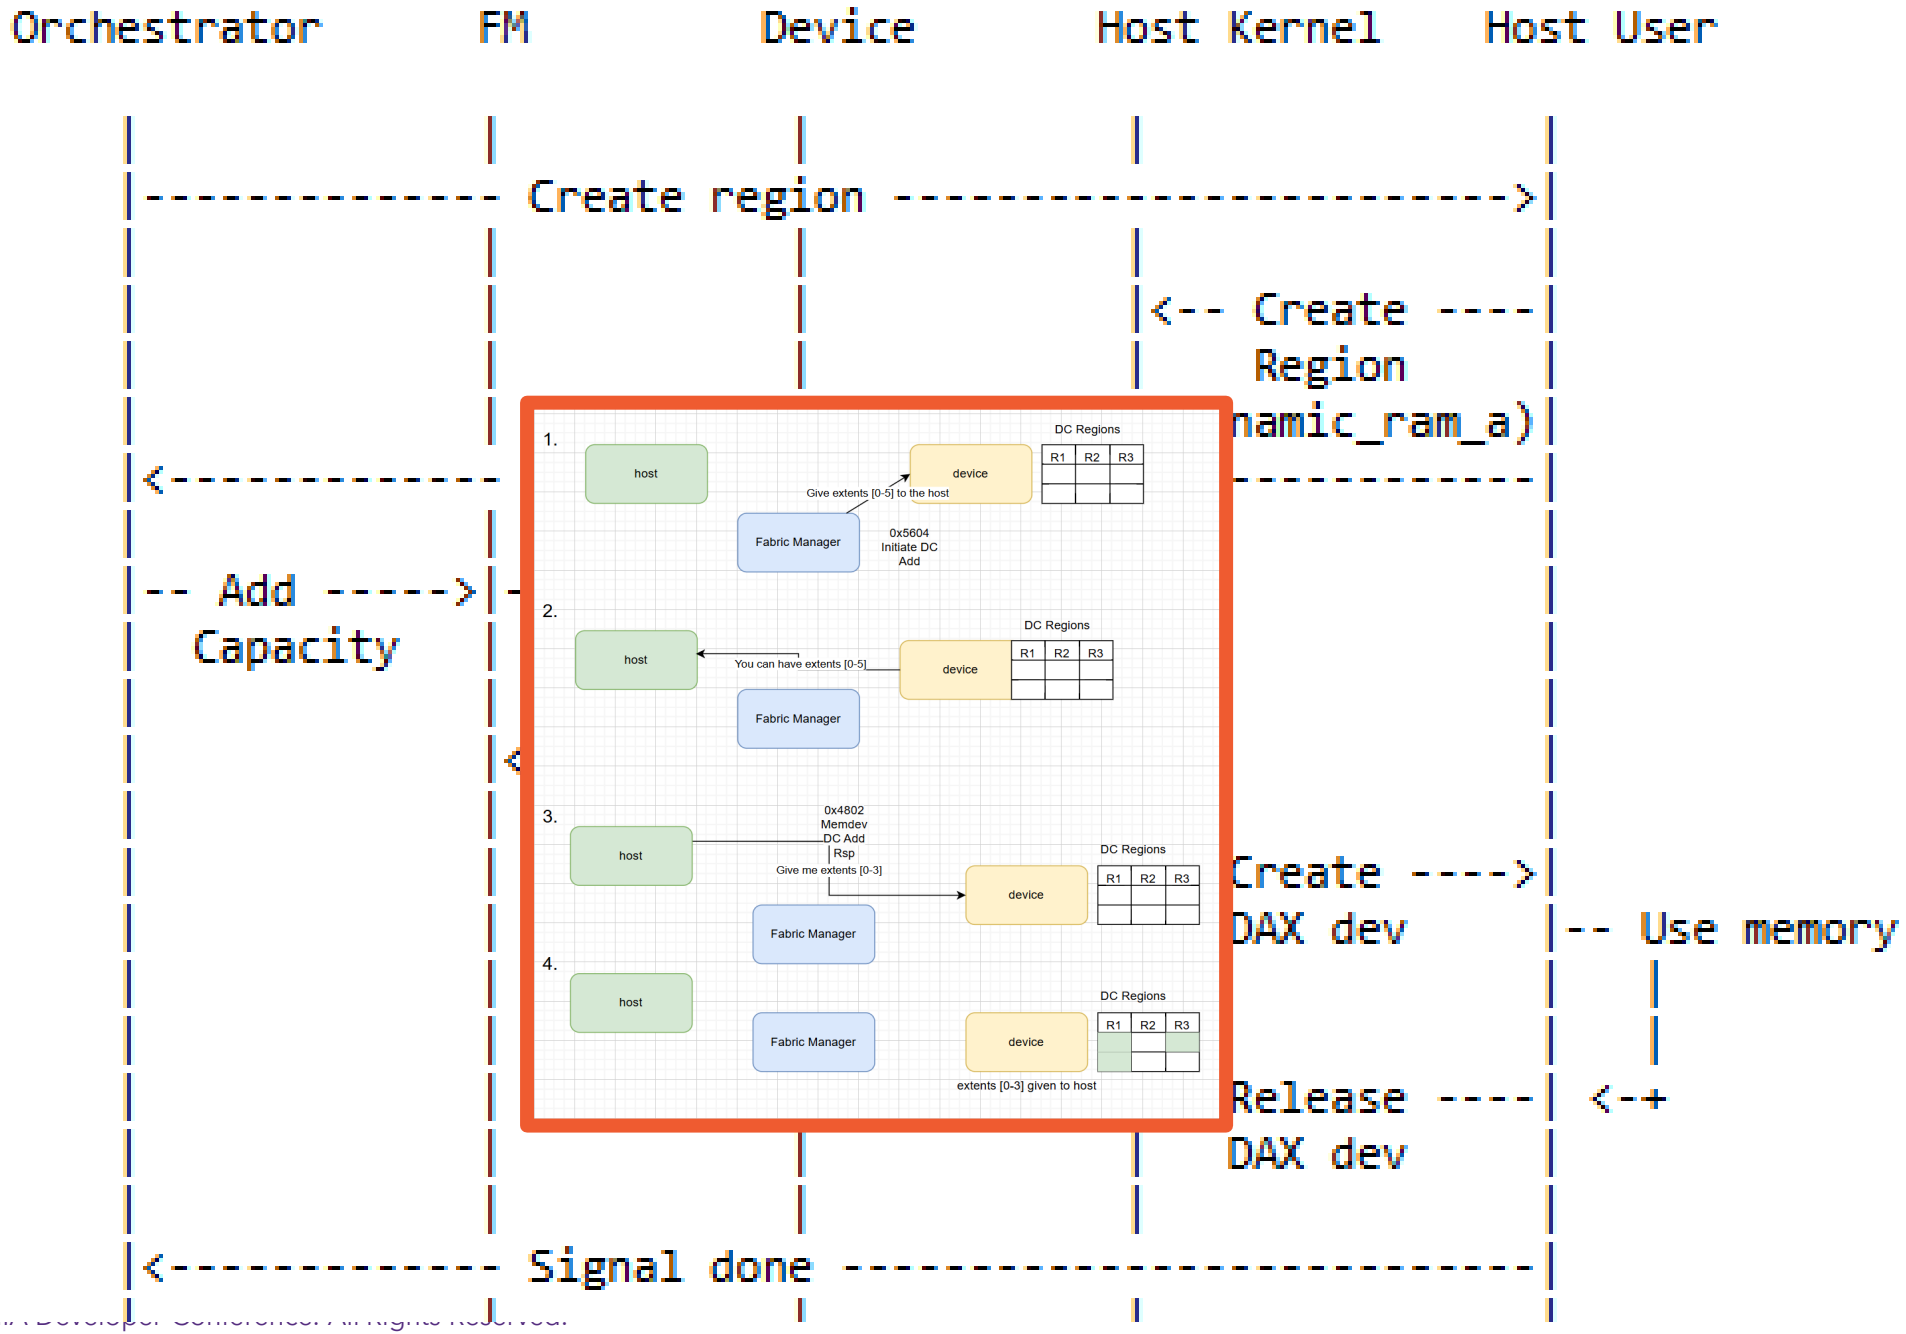

How is Memory Dynamically Added – Roles

- Orchestrator: initiator, drives when to add/release capacity
- Fabric Manager: "wrapper" for the FMAPI
- Device
- Host Kernel
- Host User

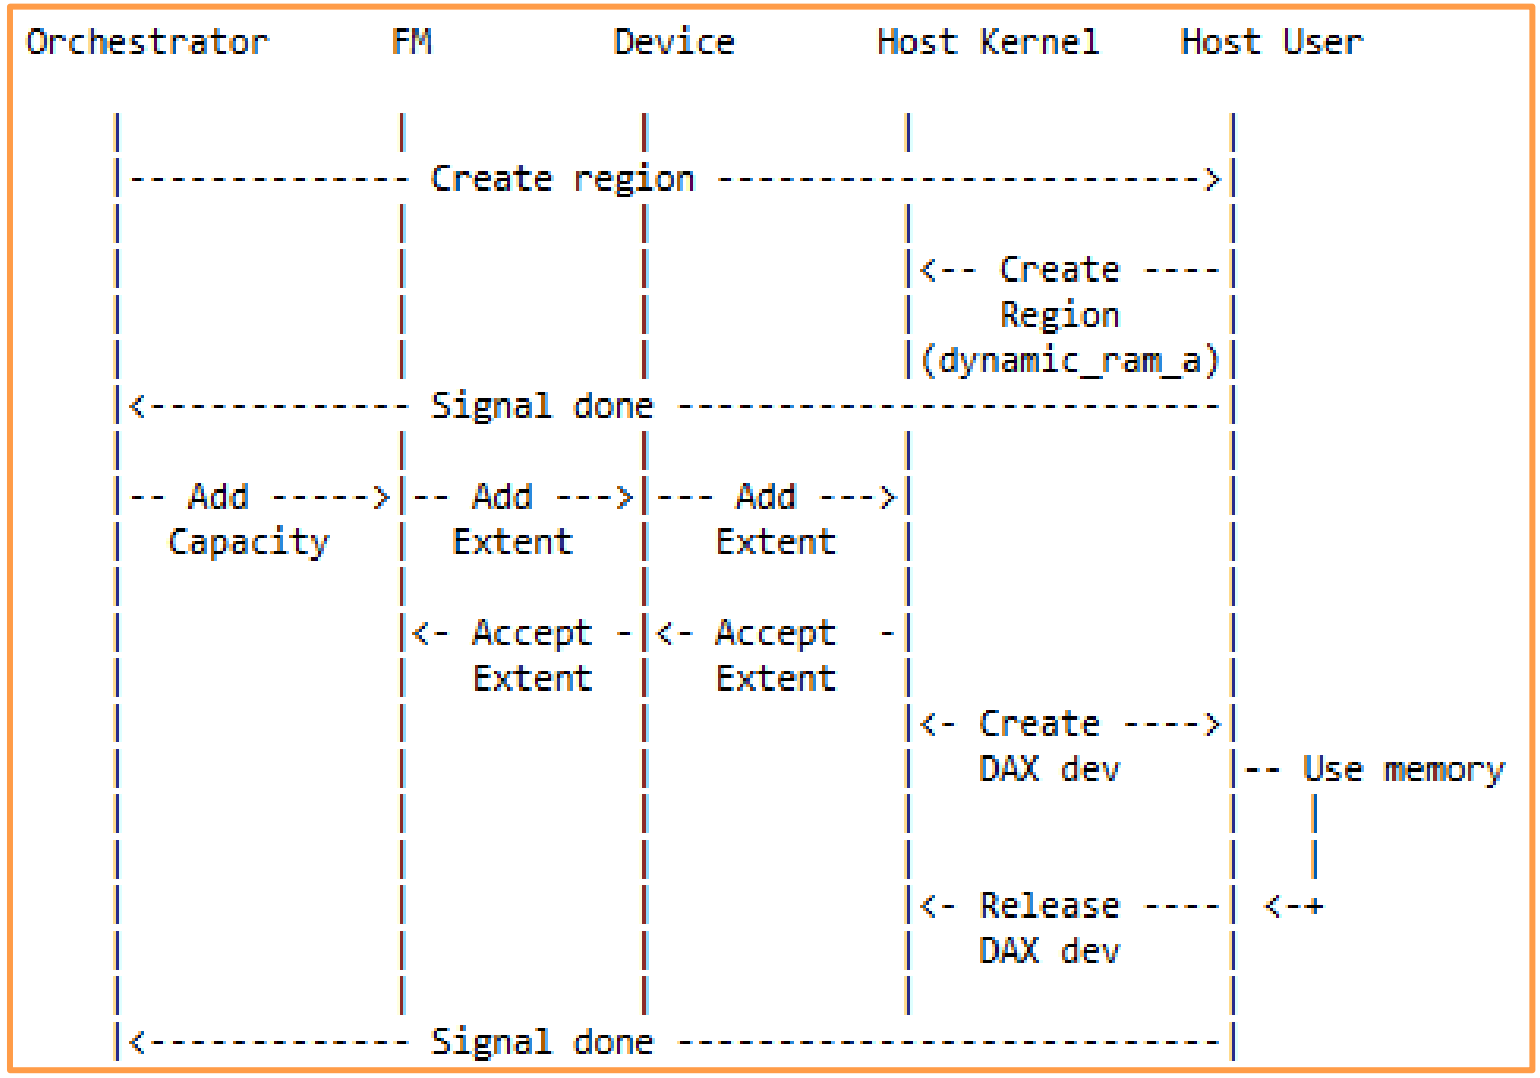

[2] CXL Driver DCD Support Patchset

How is Memory Dynamically Added – Full Workflow

- Create Region:
  - Kernel CXL Memory Region != Device Region
  - Logical construct to expose CXL device memory through DAX device(s)
- Add Capacity
- Direct Access (DAX) Device: exposes memory through a fd
  - Use mmap() to map to proc's virtual address space
  - Reconfigure to system-ram mode and use as RAM

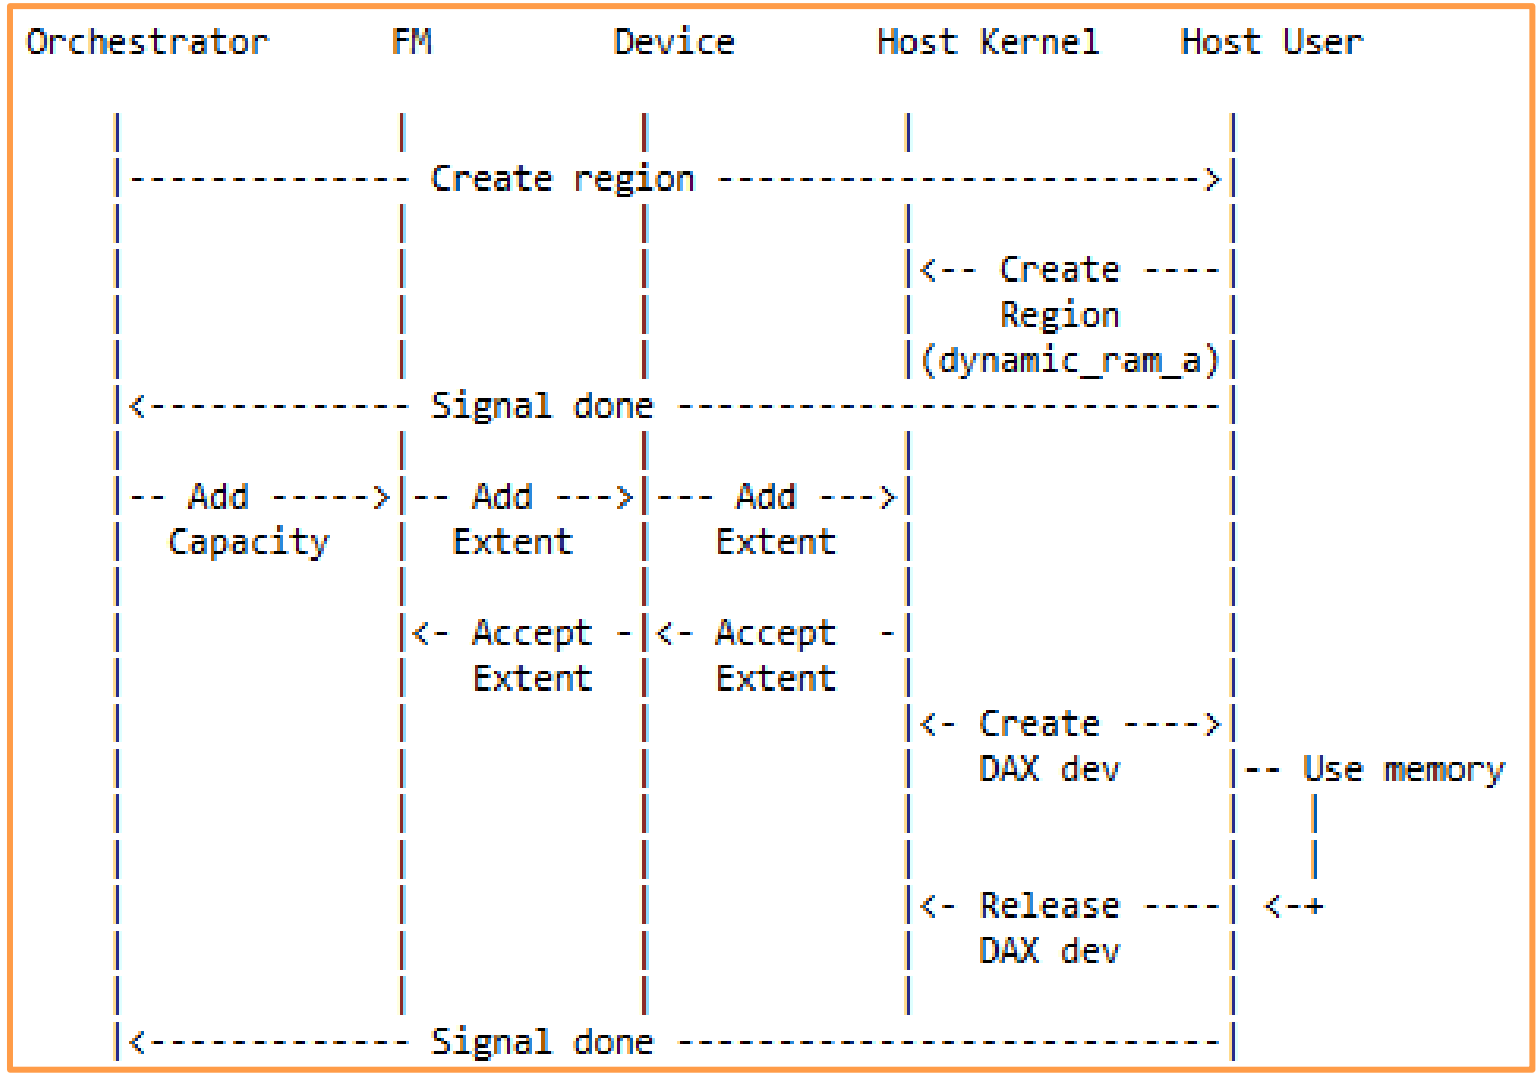

How is Memory Dynamically Added – Roles

- Orchestrator: initiator, drives when to add/release capacity
- Fabric Manager: "wrapper" for the FMAPI
- Device
- Host Kernel
- Host User

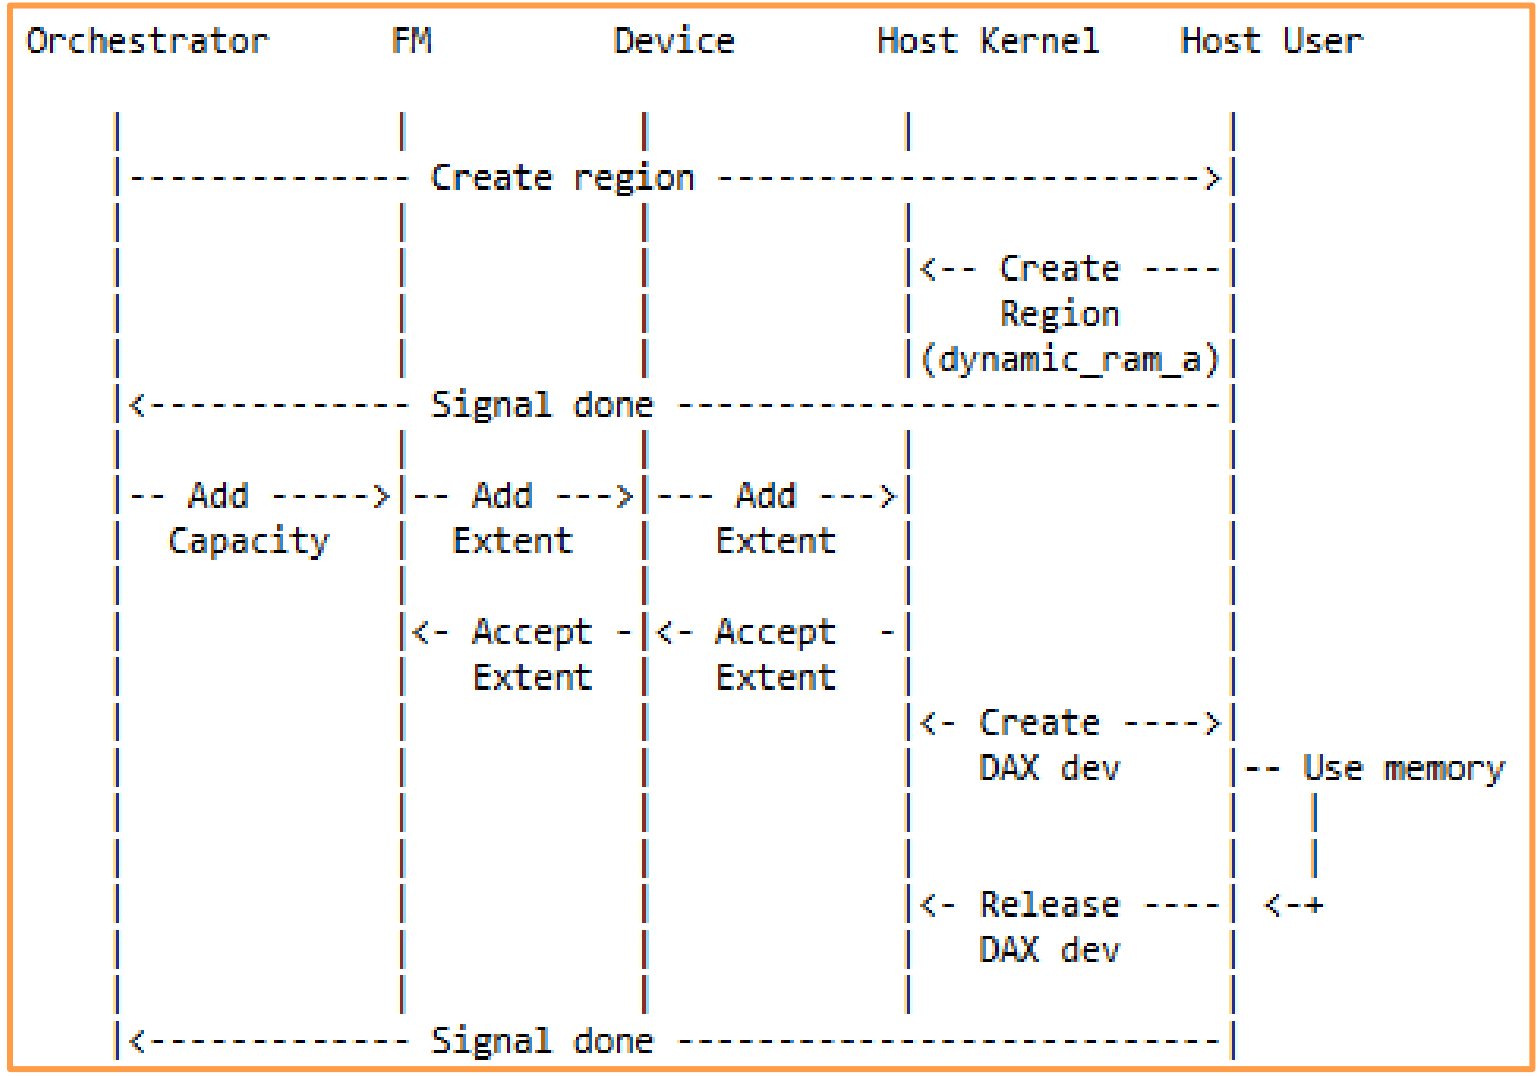

[2] CXL Driver DCD Support Patchset

Need for libcxlmi

- Main motivation: to 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

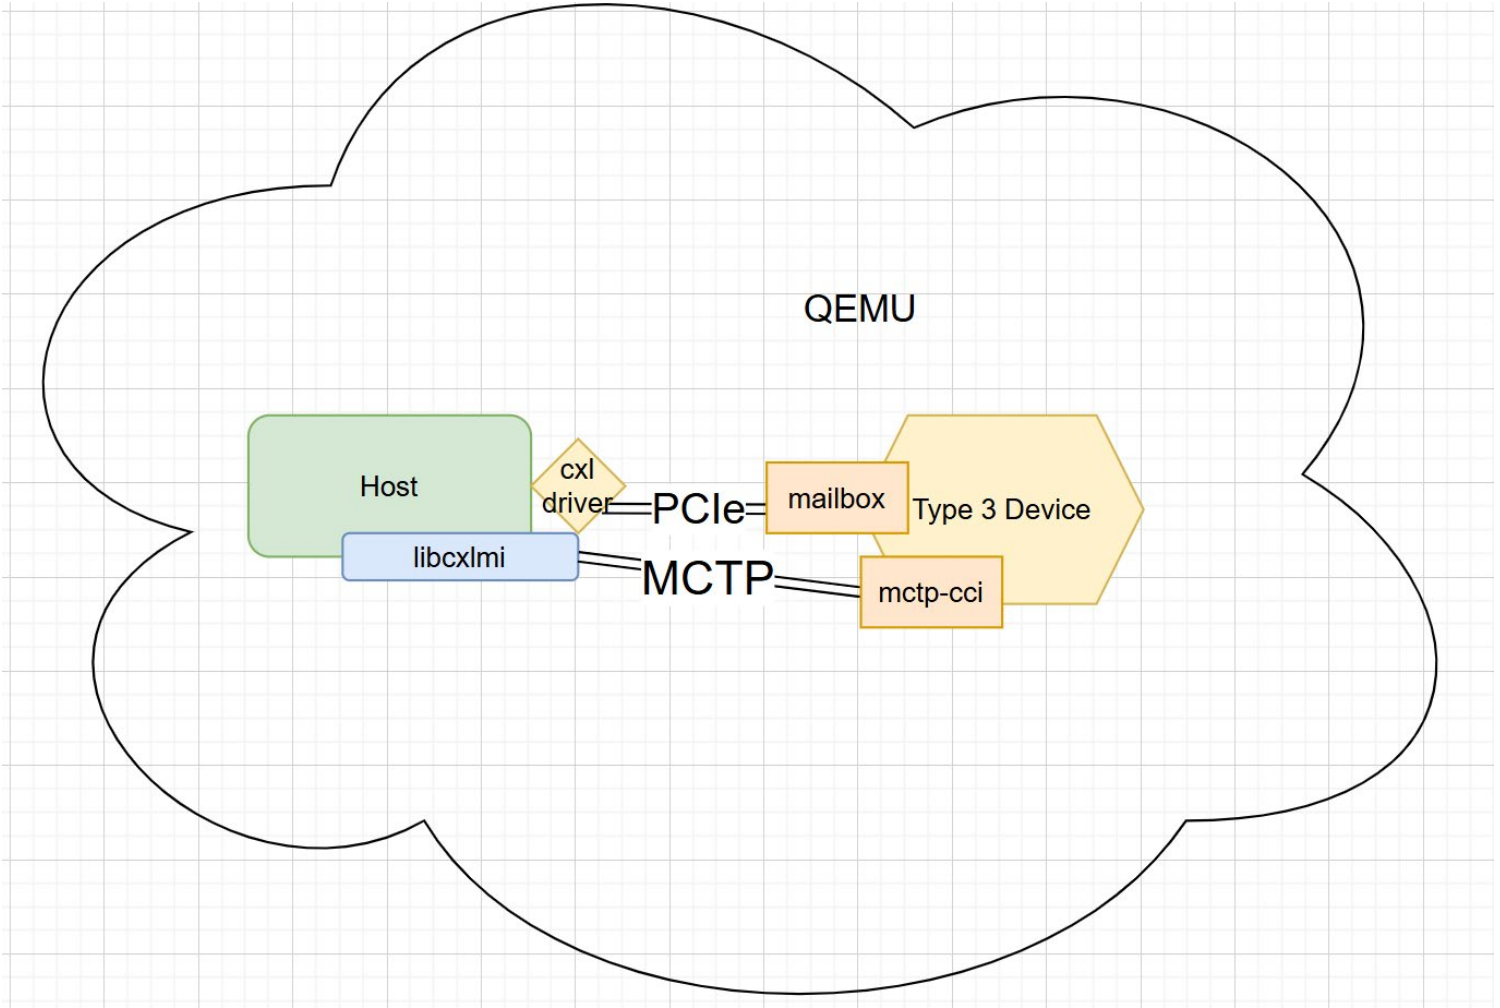

Demo

Resources

- Reproduce the demo: SNIA SDC 2025 CXL Emulation w/QEMU Demo
- cxl-test-tool
  - Easily automate VM startup with different topologies
- CXL Driver Documentation:
  - CXL Driver Operation — The Linux Kernel documentation
  - https://origin.kernel.org/doc/html/latest/driver-api/cxl/
- libcxlmi

Acknowledgement

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

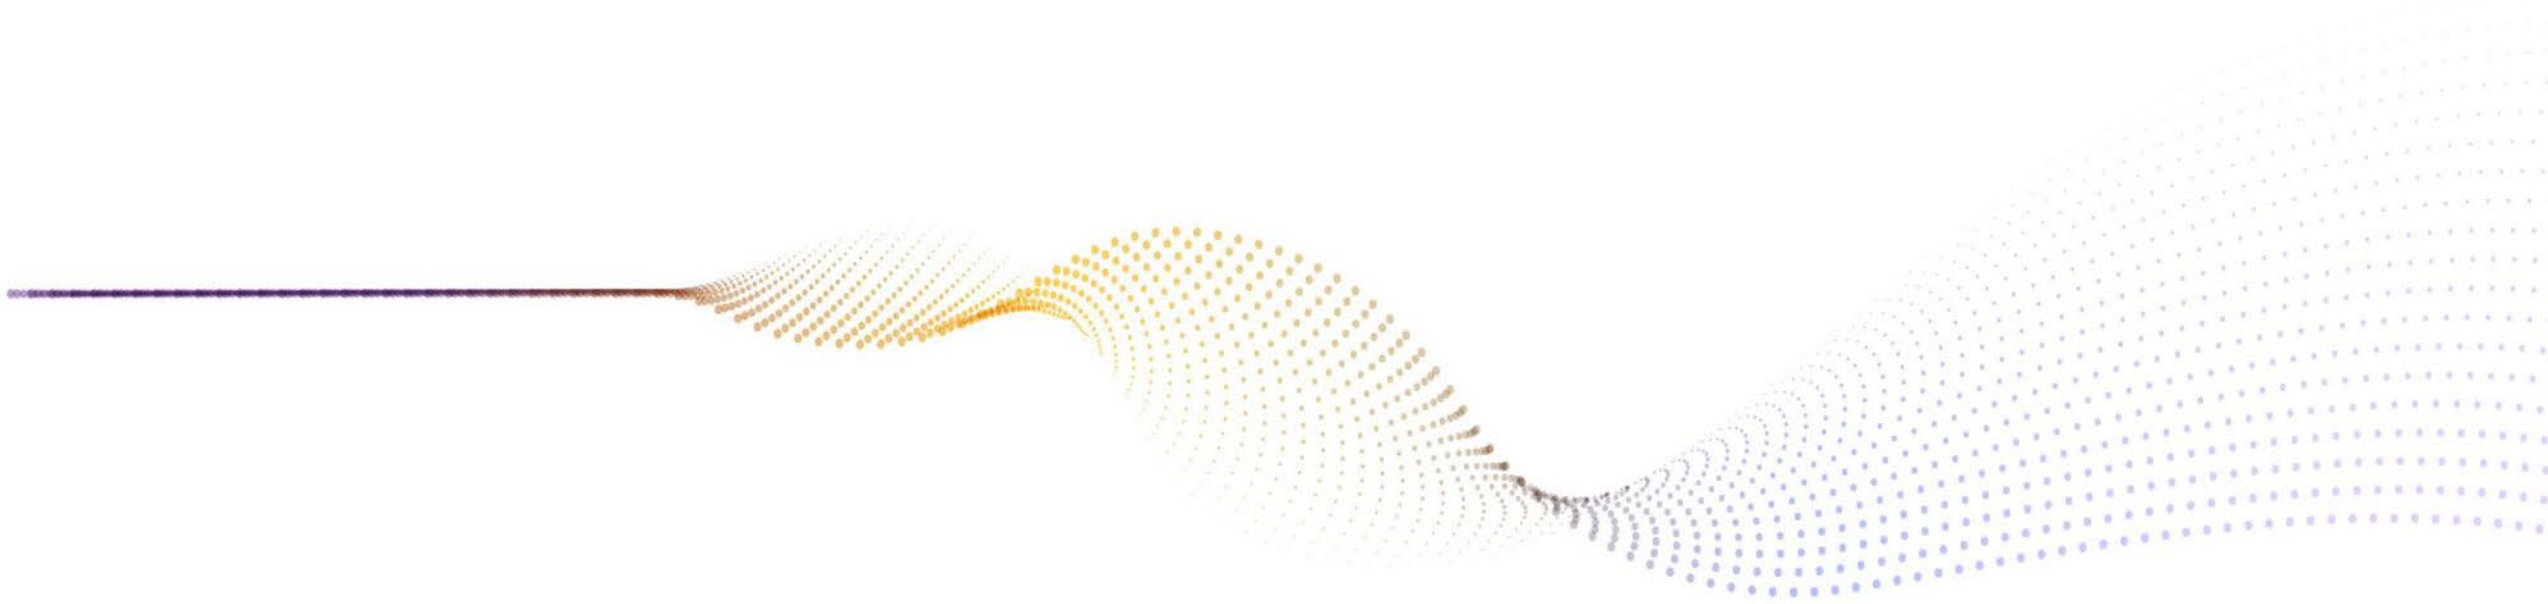

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

