

# CXL Linux

## Open Source Collaboration

Hosted by: Intel CXL Linux Development Team

10/25/2022

# CXL Linux Sync: Ground Rules

- Do not share confidential information
- Do not share confidential product details
- Do not disclose CXL consortium confidential information
- Do discuss any Linux questions about **released** CXL specifications:
  - <https://www.computeexpresslink.org/spec-landing>
- Do use IRC as a supplement for this sync meeting
  - #cxl on irc.oftc.net

# CXL Linux Sync: Agenda 10/25

- Opens:
  - FSDAX page reference counting rework (merged in mm-unstable)
  - FSDAX -> notify\_failure() regression work still pending
  - Code First ECR: ['SP' attribute in SRAT](#)
  - QEMU emulation status update
  - Others?
- Fixes pending for v6.1-rc
- Features in flight for v6.2
- Rough plans for post v6.2 work

# CXL Linux Sync: v6.1 Fixes

<https://git.kernel.org/pub/scm/linux/kernel/git/cxl/cxl.git/log/?h=fixes>

Queued:

- Mailbox input payload fix
- Decoder commit crash
- LSA payload handling fix
- CFMWS NUMA Node setup

Pending:

- Fix switch attached to single-port host-bridge
- BUG in create-region when no more intermediate port decoders available

# CXL Linux Sync: v6.2 Features

In rough priority order, feedback welcome:

- RCH Support (including DVSEC Range Register enumeration)
- Cache invalidation for region physical invalidation scenarios
- RAS Capability Tracing on RCH and VH AER events
- CXL Events to Linux Trace Events (including interrupts)
- EFI CPER record parsing for CXL error records
- Forward and reverse address translation (DPA  $\Leftrightarrow$  HPA)
- Volatile Region Discovery
- Volatile Region Provisioning
- Security commands (including background commands)
- CXL perf monitoring
- Miscellaneous cleanups and renames

# CXL Linux Sync: Post v6.2 Features

- Dynamic Capacity Device support
  - Sparse DAX Region infrastructure
  - DCD event plumbing
- Maintenance Feature Support (DRAM PPR)
- Switch mailbox CCI
  - Multi-head device mailbox tunneling
- Default "Soft Reserved" (EFI\_MEMORY\_SP) handling policy (cxl-cli + daxctl)

# CXL Linux

## Open Source Collaboration

Hosted by: Intel CXL Linux Development Team

8/30/2022

# CXL Linux Sync: Agenda 8/30

- Opens:
  - FSDAX -> notify\_failure() fixes
  - FSDAX page reference counting rework
- Linux v6.0-rc1 and ndctl (ndctl, daxctl, cxl-cli) v74 released
- Fix and Feature queue for v6.0-rc, v6.1 and ndctl-v75
- Rough plans for post v6.1 work for CXL 3.0 enabling

# CXL Linux Sync: Recently released

## Kernel:

- DPA Space Accounting
- PMEM Region Provisioning
- DOE Support in PCI core
- CDAT retrieval (for debug)

## User tooling:

- `cxl create-region`
- `cxl reserve/free-dpa`
- `cxl list -vvv`

# CXL Linux Sync: Next fixes and features

- 'arch\_flush\_memregion()'
- Fix validation of x1 switch topologies
- Volatile region provisioning
- Region labels
- Security commands support
- Trace events for CXL events (including interrupts)
- 'cxl monitor' command
- CXL AER handling
- Address translation

# CXL Linux Sync: Future work

- Performance monitoring
- Maintenance Feature Support (DRAM PPR)
- Dynamic Capacity Device support
- Default "Soft Reserved" (EFI\_MEMORY\_SP) handling policy

# CXL Linux

## Open Source Collaboration

Hosted by: Intel CXL Linux Development Team

7/26/2022

# CXL Linux Sync: Agenda 7/26

- Opens:
  - FSDAX page reference counting rework
- What is queued for v6.0 (and ndctl-v74)?
- Late v6.0 updates
- Post v6.0 work

# CXL Linux Sync: Queued for v6.0

- DOE Support in PCI core
- CDAT retrieval (for debug)
- DPA Space Accounting
- PMEM Region Provisioning

# CXL Linux Sync: In review for v6.0

- Interleave granularity fixes
  - Fix host-bridge x1 interleave constraint
  - [Fix region granularity > host-bridge granularity handling](#) (scale factors must match)

# CXL Linux Sync: Post v6.0 material

- Pre-existing region enumeration
- Volatile region provisioning
- XORMAP interleave support
- Trace Events for CXL Events
- List Poison
- Scan Media
- Address translation
- Region persistence in labels
- Region enumeration via labels

# CXL Linux

## Open Source Collaboration

Hosted by: Intel CXL Linux Development Team

6/28/2022

# CXL Linux Sync: Agenda: 6/28

- Opens:
  - CXL Device Tree Support
  - MEM\_HWINIT\_MODE=0
  - QEMU mainline CXL support is live
- What is in review for v5.20 (and ndctl-v74)
- What else might make v5.20?
- What is post v5.20 material?

# CXL Linux Sync: v5.20 in review

- [DOE + CDAT](#)
- [PMEM Region Provisioning](#)

# CXL Linux Sync: v5.20 on deck

- Pre-existing region enumeration
- Region persistence in labels
- Region enumeration via labels
- Address translation foundation

# CXL Linux Sync: Post v5.20 material

- List Poison
- Scan Media
- XORMAP interleave support
- Trace Events for CXL Events
- Address translation (in cxl-cli) for all kernel supported Events, List Poison, and Scan Media

# CXL Linux

## Open Source Collaboration

Hosted by: Intel CXL Linux Development Team

5/31/2022

# CXL Linux Sync: Agenda: 5/31

- What is in v5.19?
- What is on deck for v5.20?
- What is post v5.20 material?
- Opens

# CXL Linux Sync: v5.19 / ndctl-v73

## Kernel

- lockdep annotations
- CXL\_OSC (native CXL hotplug + error "handling")
- Disable suspend
- Mem\_enable fixes

# CXL Linux Sync: v5.20 / ndctl-v74

## Kernel

- Region Provisioning
- DOE Core
- CXL CDAT Retrieval
- Event record handling core
  - Scan Media records
  - Event Interrupts
  - Background command timesharing

## Userpace

- 'cxl create-region'
- Region listing support
- Scan media / Event records to json
- Address translation

# CXL Linux Sync: Post v5.20 / v6.0

## Kernel

- SPDM Attestation
- IDE
- Security commands

## Userspace

- Attestation helper process
- CXL Device-DAX Policy