CXL Memory in Windows

Scott Lee – Principal Software Engineering Lead

Microsoft

Agenda

- Windows CXL Memory Architecture
- OS-First CXL Memory Reliability Availability and Serviceability (RAS)
- CXL Memory in Windows
- CXL Memory Virtualization
- Current Status
- Futures

Windows CXL Memory Architecture

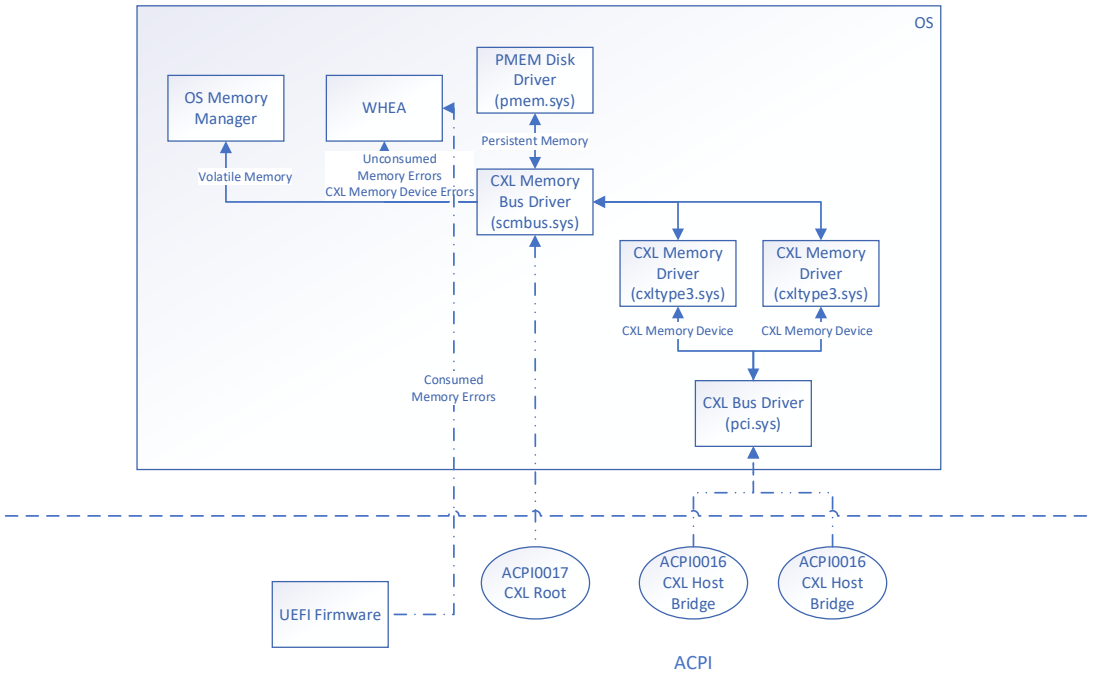

- CXL Bus Driver – CXL bus and root port level functionalities
- CXL Memory Driver – CXL Type3 device functionalities
- CXL Memory Bus Driver – system level CXL memory functionalities (e.g. interleave, Device to System Physical Address mapping, Specific Purpose memory)
- WHEA – Windows kernel subsystem for hardware event handling

OS-First CXL Memory RAS

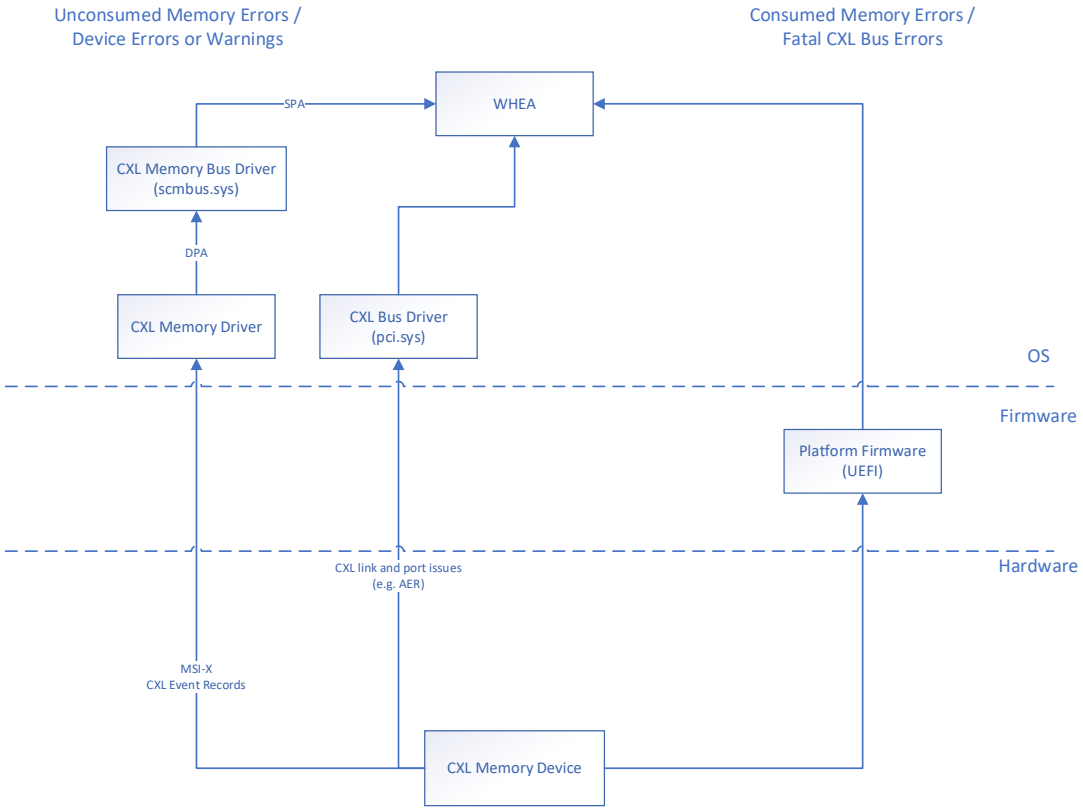

- CXL Memory Driver detected through interrupt and CXL event records
  - Unconsumed memory errors
  - CXL memory module level events
- CXL Bus Driver detected through interrupt
  - CXL bus and root port related events (e.g. Advance Error Reporting - AER)
- Platform Firmware detected through CPU/System on Chip (SoC)
  - Consumed uncorrectable memory error
  - Viral
- WHEA handles the hardware event
  - Events convey through Common Platform Error Record (CPER)
  - Offline memory at 4K page size or greater size depending on the error
  - Crash the system if can't offline affected memory

CXL Memory in Windows

- Two approaches for CXL memory usage
  - General Purpose (GP) Memory
    - Memory-only NUMA node for memory that have different characteristics than memory on NUMA (Non-Uniform Memory Architecture) node that has CPU
    - Memory available to any requester so can create unexpected issues if performance is slower
    - Default OS memory management prioritizes memory allocation from nearest NUMA node
  - Specific Purpose (SP) Memory
    - Each memory tier in memory-only NUMA node
    - Memory dedicated for specific usages and allocated through new APIs (Application Programming Interface)
    - Software stores data in different memory tiers based on its characteristics
    - Requires software code changes
    - Called Dedicated Memory in Windows
    - OS Memory Manager can use SP memory for various usages (e.g. pagefile)

CXL Memory Virtualization

- Do not see any current needs to virtualize CXL memory devices to VM (Virtual Machine)
- Can use CXL memory in VMs using generic reporting of memory as either General Purpose or Specific Purpose Memory

Current Status

- Working on support for CXL 2.0 or greater implementations with OS-First RAS with interested applications
- Initial focus is on memory expansion (i.e. direct attached CXL memory devices) scenarios with no CXL switch
- Have preview Windows builds available to try
  - NDA (Non-Disclosure Agreement) partners have access to more collaterals
  - Non-NDA partners can gain access through Windows Insider program
  - Support available in both Windows server and client OS
- Official support is TBD (To Be Determine)

Futures

- Firmware-First RAS
- Dynamic Capacity Device (DCD)
- Non-volatile/Persistent Memory
- Trusted Execution Environment Security Protocol (TSP)
- Hot-add and remove