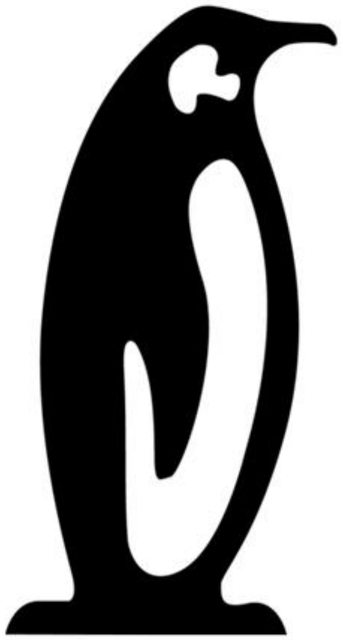

Linux Plumbers Conference

Vienna, Austria | September 18-20, 2024

CXL RAS Page offline for Corrected Errors

Srinivasulu Thanneeru(Micron)

LINUX
PLUMBERS
CONFERENCE Vienna, Austria / Sept. 18-20, 2024

Micron Confidential

Agenda

- Current Status
  - Type-3 Device Fault Domains
  - Error Coverage
  - Error Reporting
    - DRAM Event Records Logs
- CXL Page offline support
  - Predictive Failure Algorithm (PFA)

Micron Confidential

CXL Type 3 - Fault Domains

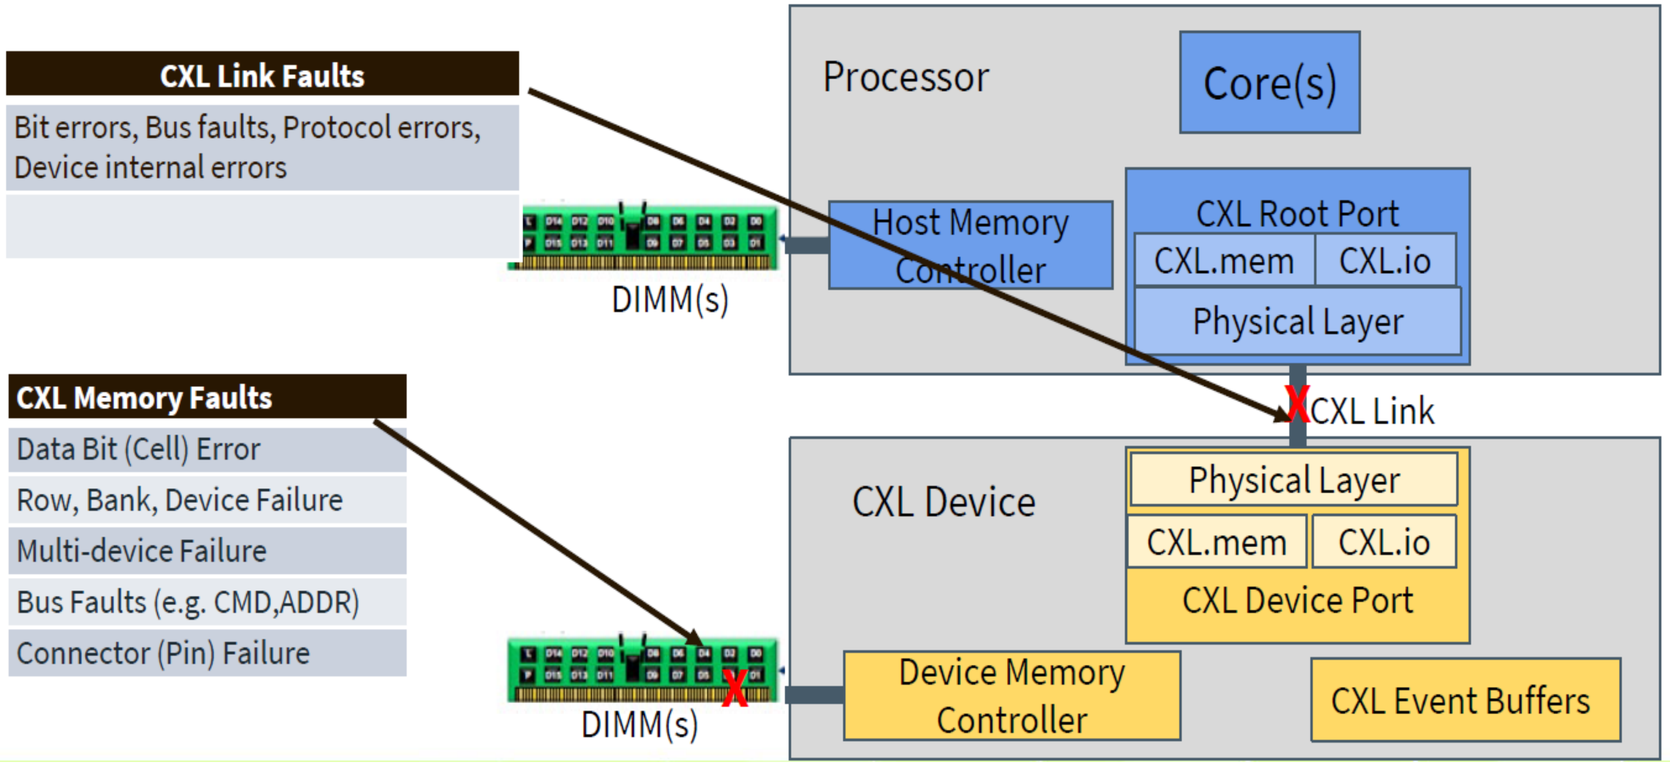

CXL Type 3 - Memory Faults Coverage

Possible Fault Type	Fault Causes (Examples)	Coverage (RAS Feature)
Data Bit (Cell) Error	High energy particle strike. Soft Error (SE). Transient error.	ECC, Demand Scrub, Patrol Scrub
Row Failure	**Marginality. Persistent error**	**OS soft page offline, PPR, Chipkill**
Bank Failure	**Marginality. Persistent fault**	**Chipkill, OS soft page offline**
Device Failure	**Marginality. Persistent fault**	**Chipkill, OS soft page offline**
Multi-device Failure	Persistent device failures	Contained by Poisoning (UCR) , MCA Recovery, OS hard page offline

ECC: Error Check and Correction
RAS: Reliability, Availability, Serviceability
PPR: Post Package Repair
UCR: Uncorrected Recoverable Error

CXL Type 3 - Memory CE Error Reporting Example

- Type-3 Device Logs the CE event.
- Type-3 Device Triggers interrupt (OS-First) on Corrected Error when PFA threshold triggers.
- OS gets the event records and populate traces.
- RAS-Daemon captures DRAM event records.

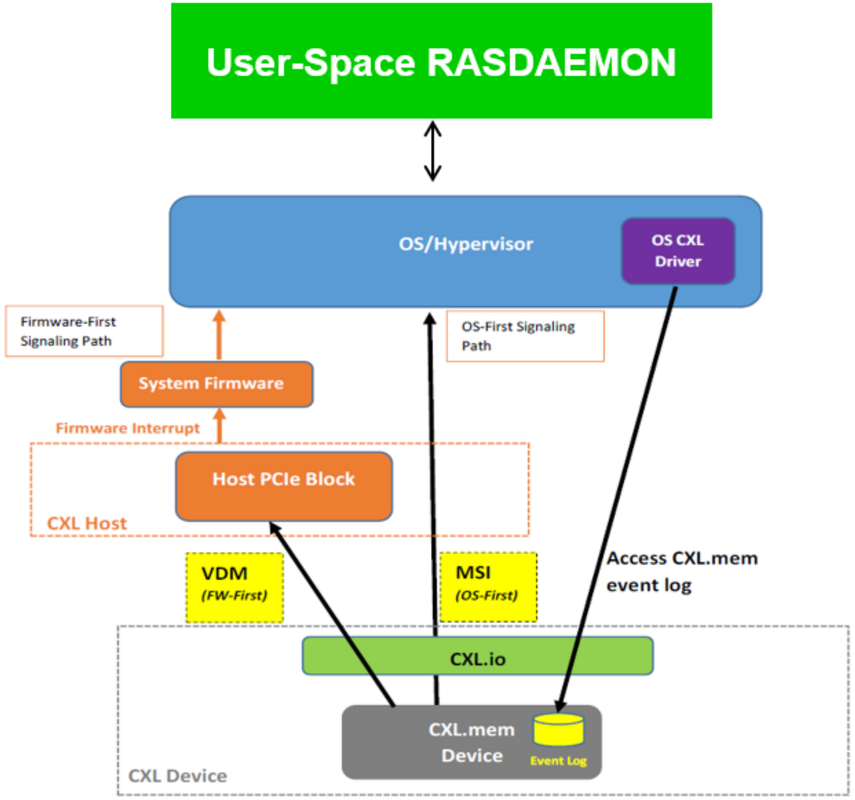

Micron Confidential

CXL Type 3 - Predictive Failure Algorithm (PFA)

- Devices implements an internal Correctable Error threshold or Predictive Failure Algorithm (PFA).
- Designed to assist the host in avoiding usage of memory locations that may degrade into an uncorrectable error.
- Device generates a DRAM Event Record with a Threshold Event descriptor and DPA pointing to the questionable memory location.

Micron Confidential

CXL Type 3 - Page offline support

Example Event Record:

Event Record Name: **DRAM Event Record**
Event Record Identifier: 601dcbb3-9c06-4eab-b8af-4e9bfb5c9624
Event Record Length: 128
Event Record Flags: 0
Event Record Handle: 0x8001
Related Event Record Handle: 0x0
Event Record Timestamp: 1717485377838306151 2024-06-04 07:16:17.838306
Maintenance Operation Class: 0
Event Record Reserved_2: 0
Event Record Data:
Physical Address: 0x41
Memory Event Descriptor: 0x0
Uncorrectable Event: 0
Threshold Event: 1
Poison List Overflow Event: 0
Memory Event Type: Media ECC Error
Transaction Type: Internal Media Management
Validity Flags: 0xff
Memory Event Location:
Channel: 2 (Valid: 1)
Rank: 0 (Valid: 1)
Bank Group: 0 (Valid: 1)
Bank: 0 (Valid: 1)
Row: 0x0 (Valid: 1)
Column: 0x0 (Valid: 1)
Nibble Mask: 0x4000 (Valid: 1)
Correction Mask (Valid: 1)

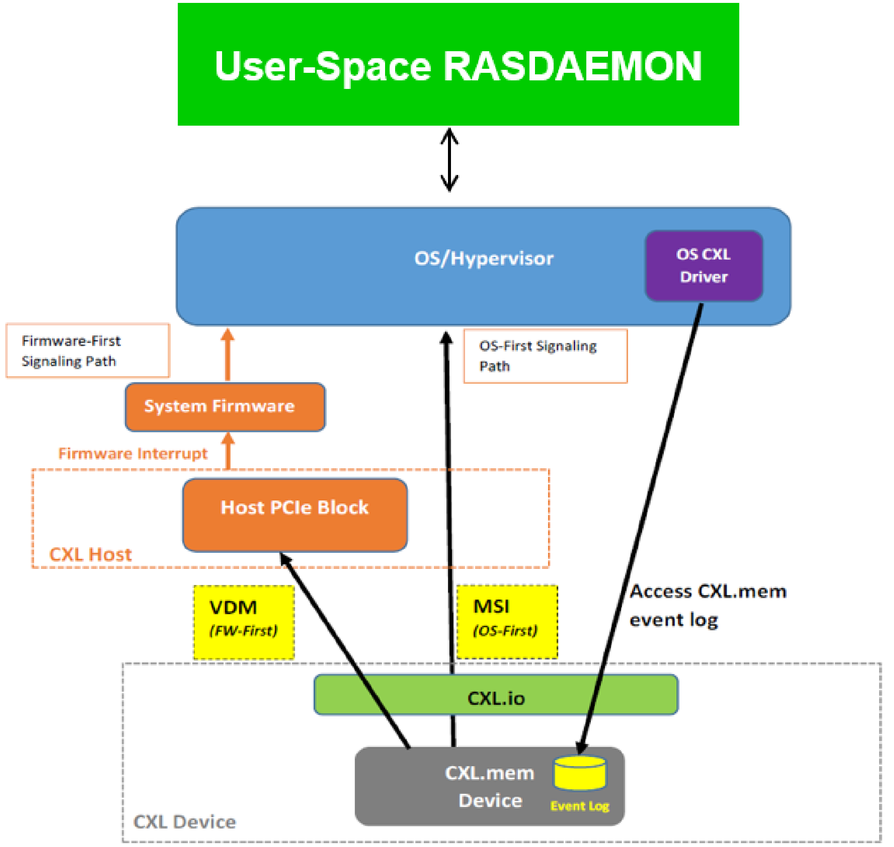

CXL Type 3 - Page offline support

- Type-3 Device logs the CE event
- Type-3 Device Triggers interrupt (OS-First) on Corrected Error when PFA threshold triggers.
- OS gets the event records and populate traces
- RAS-Daemon captures DRAM event records
- CE with threshold set, then do page-offline for DPA/HPA.

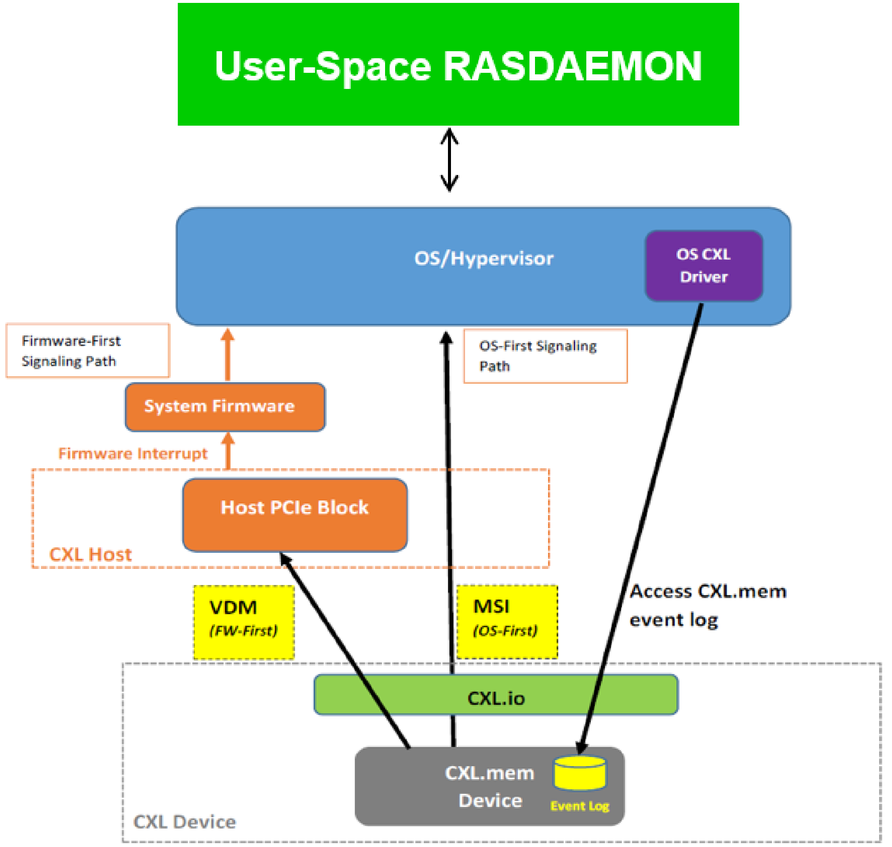

CXL Type 3 - Page offline support

```
diff --git a/ras-cxl-handler.c b/ras-cxl-handler.c
index 037c19c..04be770 100644
--- a/ras-cxl-handler.c
+++ b/ras-cxl-handler.c

+       if (tep__get_field__val(s, event, "hpa", record, &val, 1) < 0)
+               return -1;
+       ev.hpa = val;
+       if (trace__seq__printf(s, "hpa:0x%llx ", (unsigned long long)ev.hpa) <= 0)
+               return -1;
+
        if (tep__get__field__val(s,  event, "dpa__flags", record, &val, 1) < 0)
                return -1;
        ev.dpa__flags = val;
@@ -1005,6 +1017,12 @@ int ras__cxl__dram__event__handler(struct trace__seq *s,
                }
        }

+#ifdef HAVE_MEMORY_CE_PFA
+       /* Account page corrected errors */
+       if (!ev.descriptor.uncorreted && ev.descriptor.threshold )
+               ras__record__page__error(ev.hpa, PAGE_CE_THRESHOLD + 1, now);
+#endif
+
```

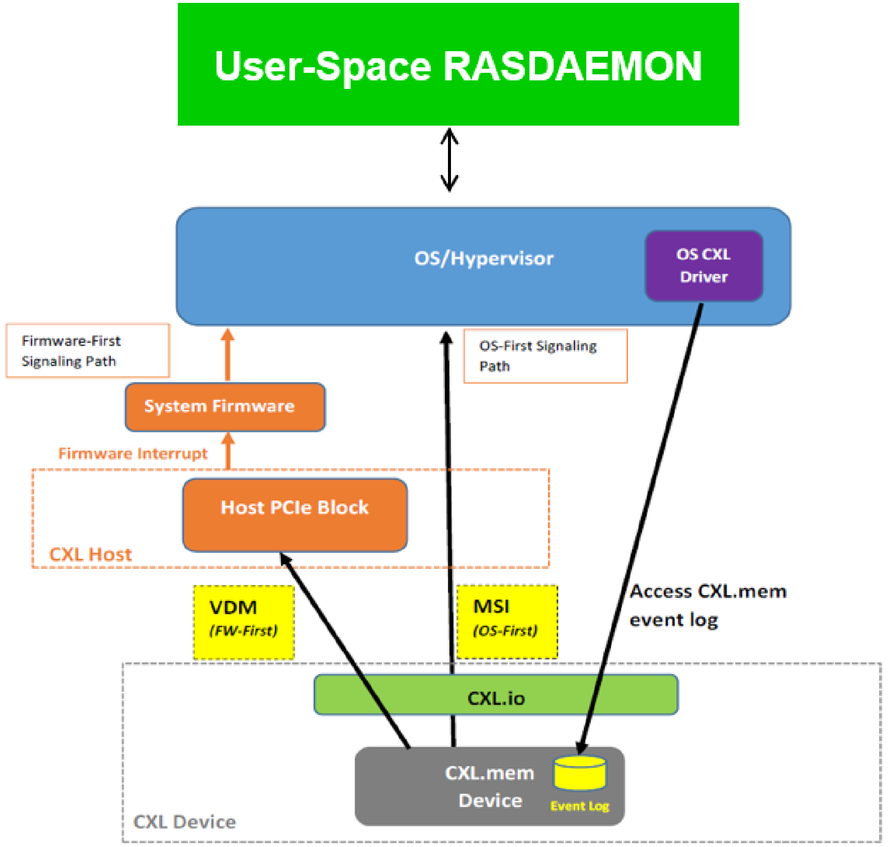

CXL RAS Page offline for Corrected Errors

LINUX
PLUMBERS
CONFERENCE Vienna, Austria / Sept. 18-20, 2024

Micron Confidential