



OCP

Software Defined Memory:
Workstream Update

**SDM Team
(Manoj Wadekar)**

Agenda

- FTI and SDM - introduction
- Team charter
- Operation methods
- Current activities and Status
- SDM Graduation!

FTI Focus Areas (Started in 2021)

- **AI HW-SW Design Collaboration**
- **Cloud Service model**
- **Software Defined Memory**



SDM Team Charter

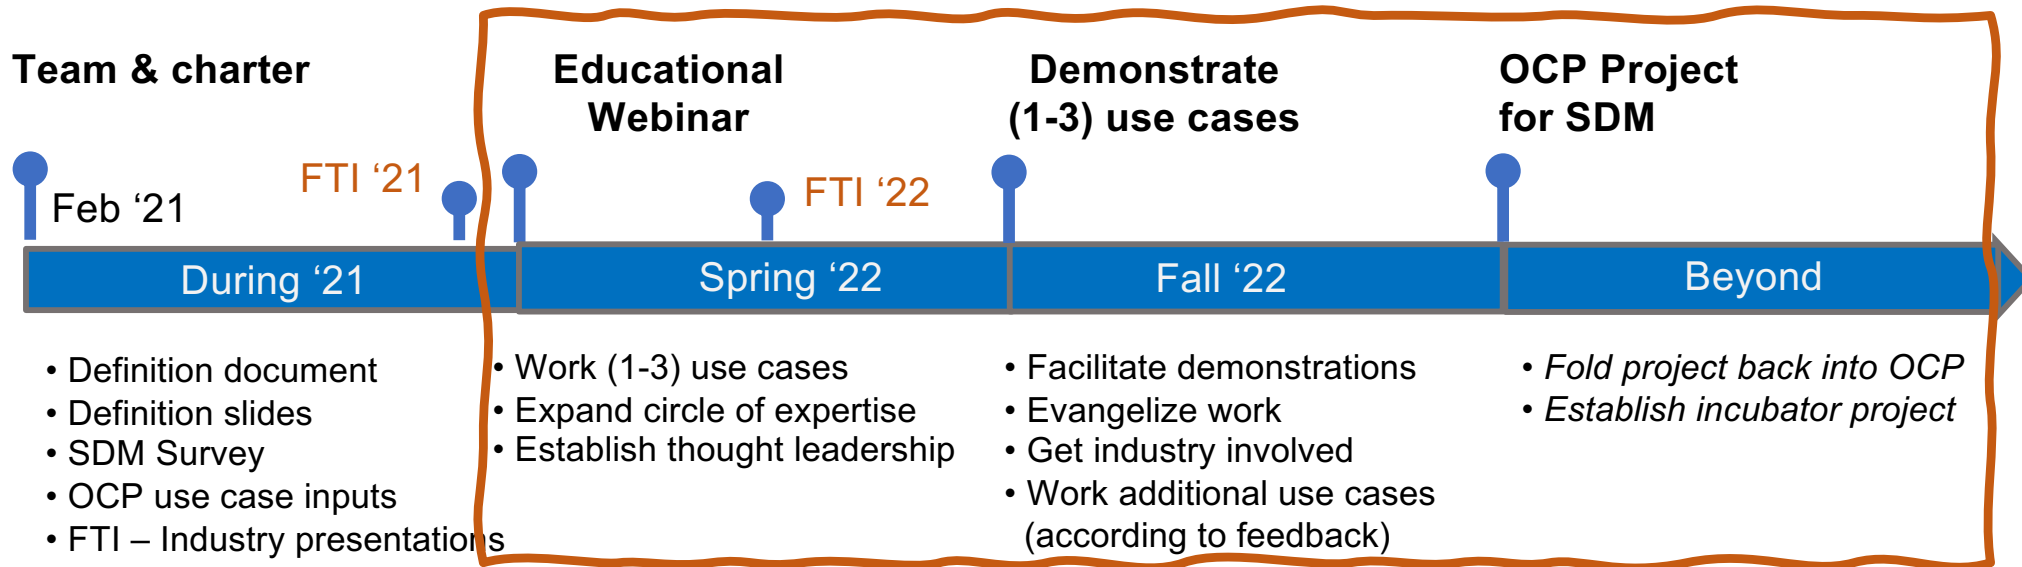
- **Identify** key applications driving adoption of Hierarchical/Hybrid memory use cases
- Establish architecture and **nomenclature** for such Systems
- Offer **benchmarks/metrics** that enable validation of novel ideas for HW/SW solutions for such systems

Drive HW-SW Co-design for SDM Use Cases

(Ad-hoc) Industry Participation

- **Device/System Vendors:** Intel, Micron, Samsung, SMART Modular, Cisco
- **CPU Vendors:** AMD, ARM, Intel
- **CSP:** Meta, Microsoft, Uber, VMware
- **ISVs:** Microsoft, VMware, MemVerge
- **Analyst:** Informa

SDM 2022 plan



Investigation approach:

- **Evangelize**: SDM opportunities in industry with **SDM Solution Blueprint**
- **Enable**: adoption through “**SDM_spec**” **Benchmarks** for work-load specific use cases
- **Demonstrate**: OCP Community Lab POCs, White/Research Papers, Webinars, Spring workshop

Focus Area Teams

Databases/Caching

- **Chagam Anjeneya (Reddy) (Intel)**
- Meta (Manoj), Intel, Samsung (Rekha), Micron (Puja), Vmware (Renu), Microsoft (Jui-Hao), AMD (Alex)

AI/ML & HPC

- **Samir Mittal (Micron)**
- Microsoft (Jui-Hao Chang, Daniel, Samir Rajadnya)
- AMD (Alex Branover)
- Micron (Puja Zalavadia)
- Meta (Manoj W.)
- Samsung (Rekha P)
- Intel (Reddy)
- Arm (Eddie)

Virtualized Servers

- **Renu Raman (Vmware)**
- Meta, Uber (Vikrant), Vmware, ARM, Microsoft (Daniel Berger, Samir Rajadnya), AMD (Alex), Intel (Reddy)

University/Research

- **Rekha Pichumani (Samsung)**
- Samsung, Meta, Micron (Samir)

Industry/Analyst

- **Dennis Hahn (Informa)**
- Informa, Meta, Cliff (OCP)

Focus teams



Focus Areas

Databases, AI/ML &
HPC, Virtualized
Servers

University/Research
collab,
Industry/Analyst collab



Participation

SDM members can
participate in any/all
focus teams



Deliverables

Develop consistent
blueprint that enables
demonstration of SDM
value for the focus
application



Meetings

Focus teams will
meet/work offline to
develop the content

Bi-weekly Joint
meeting assess
progress

Blueprint for HW-SW Co-design: SDM

- Description of the use cases
 - What's the application and how SDM fits into
- HW Solutions
 - Different HW components necessary to demonstrate
- SW Solutions
 - SW components necessary to demonstrate
- How to test: Benchmark (open source): SDM_bench_<app>
 - Workload/Profiles
- How to measure: Monitoring/Metrics
- Identify co-travelers: Associated Standards/Research Bodies



Collaboration: details

- Bi-weekly Project progress meeting
 - Coordinating industry discussions, presentations, OCP activity within FTI
 - SDM Webinar (May'22)
 - FMS Round Table and presentations - (Aug '22)
 - SNIA presentation, Global OCP SDM group meeting (Sept '22)
- Bi-weekly Use case discussion meeting
 - AI/ML and HPC Use case:
 - Meta: “Memory Requirements of Meta AI Workloads”
 - Microsoft: “Azure DL/ML use cases”
 - AMD: “Composable Architecture for Data centers and ML”
 - Virtualized Server Use case:
 - Microsoft: “First-generation Memory Disaggregation for Cloud Platforms”
 - Uber: “CXL memory pooling for deploying containerized microservices”
 - VMware: “Introducing Project Capitola: Software defined memory for data centric workloads”
 - Cache and Database Use case:
 - Intel/Meta:” Memache use case and Cachebench ”
 - Meta: Transparent Page Placement

SDM Workstreams

- **Local Memory Expansion**
 - Devices: CXL Memory Buffer, CXL Memory, Device Form Factors
 - System Configuration: GP Servers with CXL, GPU/Accelerator systems with CXL
 - SW Subsystem: Kernel support for memory tiering, Applications with memory tiering support,
 - Use Cases: Cache/DB, Inference servers (Meta, Uber?)
 - Benchmarks: Cachebench, db_bench, mlc
- **Pooled Memory Expansion**
 - Devices: CXL Memory Buffer, CXL Memory : with multi port support
 - System Configuration: Pooled memory systems
 - SW Subsystem:
 - Use Cases: Virtualized Servers (Microsoft, VMware, Uber)
 - Benchmarks: mlc
- **Memory Switches/Systems**
 - Devices: CXL Switches, Memory Systems with multi-port
 - System Configuration: CXL Switches, Memory Systems
 - SW Subsystem:
 - Use Cases: Virtualized Servers
 - Benchmarks: TBD
- **Near Memory Compute**

Workstream Contributions (planned)

- **HW Specs:**
 - OCP Memory Module Spec :
 - Memory Expansion Board,
 - Memory Module (Form Factors, connectors – pointers/additions)
 - **Benchmarks:** Use Case based Benchmarks and Metrics
- **SW Specs:**
 - Kernel SW:
 - E.g. Kernel/User Space page placement (TPP)
 - E.g. Instrumentation Spec (TBD)
 - BIOS/FW
 - Tools
 - Orchestration, manageability, etc.

SDM is graduating from FTI

- SDM graduates as an OCP subproject under Server Track
- **Composable Memory Systems**

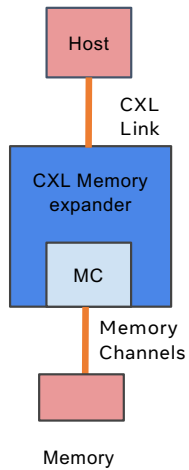


Call to Action

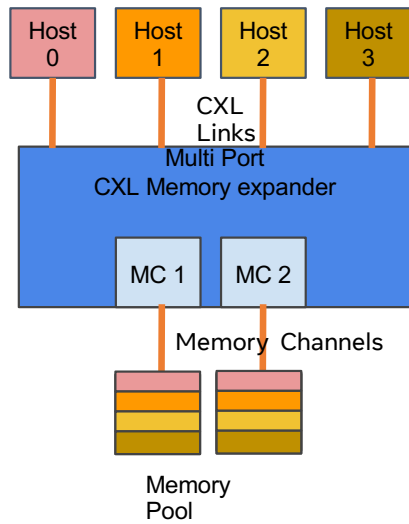
- Help us tackle Memory Challenges through Hardware and Software innovation!
- Develop TCO-optimized CXL memory
- Please join CMS community as it is formalized as a sub-project under Server track
 - CMS Group meets to plan for the details, Wednesday 10/19/2022
 - 2.30-5.00pm, SJCC - Lower Level - LL20BC
 - CMS Workgroup: [Link](#)

Backup

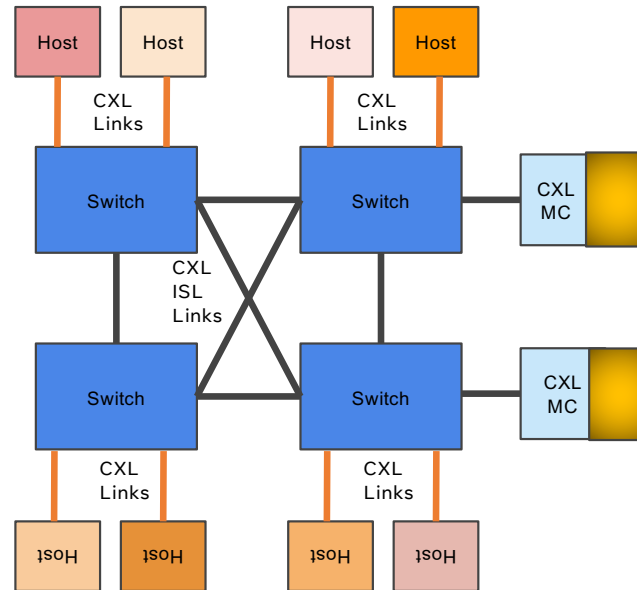
CMS Workstreams



Local Memory



Pooled Memory



Memory Switches/Systems