Driving HW-SW Co-Design for Software Defined Memory Systems

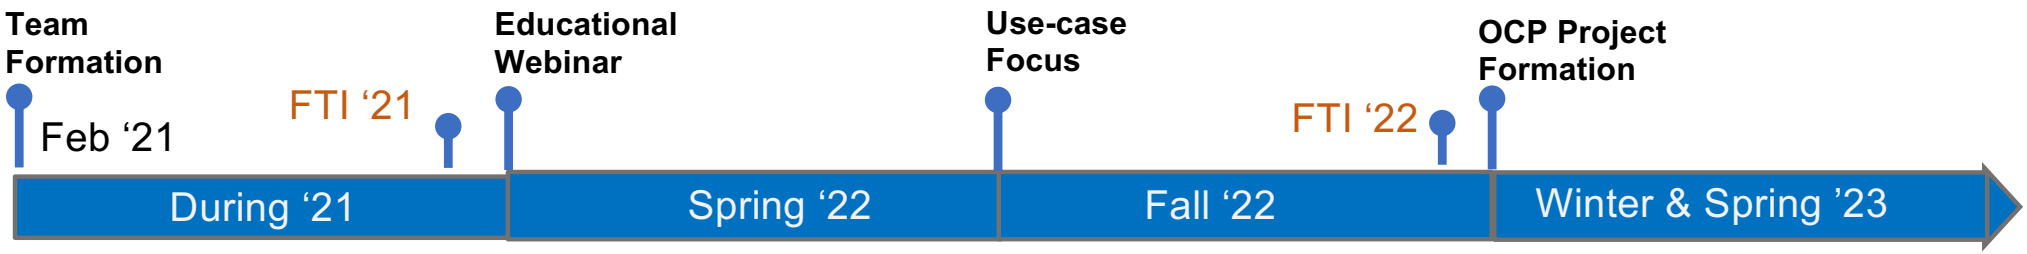

- Definition documents
- SDM Survey
- Industry presentations

- Expand group membership
- Establish thought leadership
- Define charter & deliverables

- Databases & caching
- Virtualized servers
- AI/ML & HPC

- Fold project back into OCP
- Establish incubator project

Charter

Focus on key applications driving CMS adoption

Establish CMS architecture & nomenclature

Identify gaps in specifications across full-stack

Offer benchmarks enabling innovations in new and emerging use cases

Industry Participation

Device vendors
Intel, Micron, Samsung, Smart Modular

CPU vendors
AMD, ARM, Intel

CSP
Meta, Microsoft, Uber, Vmware, Cisco

ISV and others
Microsoft, Vmware, Informa, MemVerge

Deliverables

Evangelize with CMS **solution blue-prints**

Enable **OCP_CMS specifications** (where needed) for consistent design

Demonstrate with OCP labs, research papers, webinars and workshops

SDM System Workstreams & Scope of Deliverables

	Frontier A Local Memory Expansion
Solution blue-prints [primary use cases]	Databases / Caching / Virtualization AI/ML, HPC, others
HW configurations [new capabilities to improve solution ROI]	CPU / GPU / Mem. expanders / Accelerators
HW specifications [for emerging devices]	Form factor, thermal, device mgmt., security
SW ecosystem	Use case specific SW platform readiness
SW & HW integration [specification gaps closure]	Caching controls Page migrations for memory tiering (for emerging memory, low-cost memory) Hot / cold page mapping
Open-source benchmarking [ecosystem consistency]	Cachebench, Deepspeed, others
Compute near memory	Basic Semantics

SDM System Workstreams & Scope of Deliverables

	Frontier A Local Memory Expansion	Frontier B Pooled Memory Expansion	Frontier C Switched Memory Fabrics
Solution blue-prints [primary use cases]	Databases / Caching / Virtualization AI/ML, HPC, others	Virtualization	To be developed
HW configurations [new capabilities to improve solution ROI]	CPU / GPU / Mem. expanders / Accelerators	Multi-port and/or Multi-host configuration & specifications for memory controllers	Switch based fabrics (config. differences for within & across racks)
HW specifications [for emerging devices]	Form factor, thermal, device mgmt., security	Device failure handling RAS definitions Memory fencing and other such issues	+ Switch definitions & specs Multi-protocol (network, CXL, others) Variable payload efficiency (byte to block) Electrical & Optical interconnects
SW ecosystem	Use case specific SW platform readiness	+ Dynamic memory allocation capability Memory overcommit capability	Needs focused work to sharpen scope
SW & HW integration [specification gaps closure]	Caching controls Page migrations for memory tiering (for emerging memory, low-cost memory) Hot / cold page mapping	+ Pooled memory system management Including interconnect config & mgmt.	
Open-source benchmarking [ecosystem consistency]	Cachebench, Deepspeed, others	mlc, others	
Compute near memory	Basic semantics →		Rich semantics