VMware Memory Vision for Real World Applications

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

Presented by

Sudhir Balasubramanian – Sr Staff Solution Architect

Arvind Jagannath – VCF Product Management

Customer Memory and VM trends

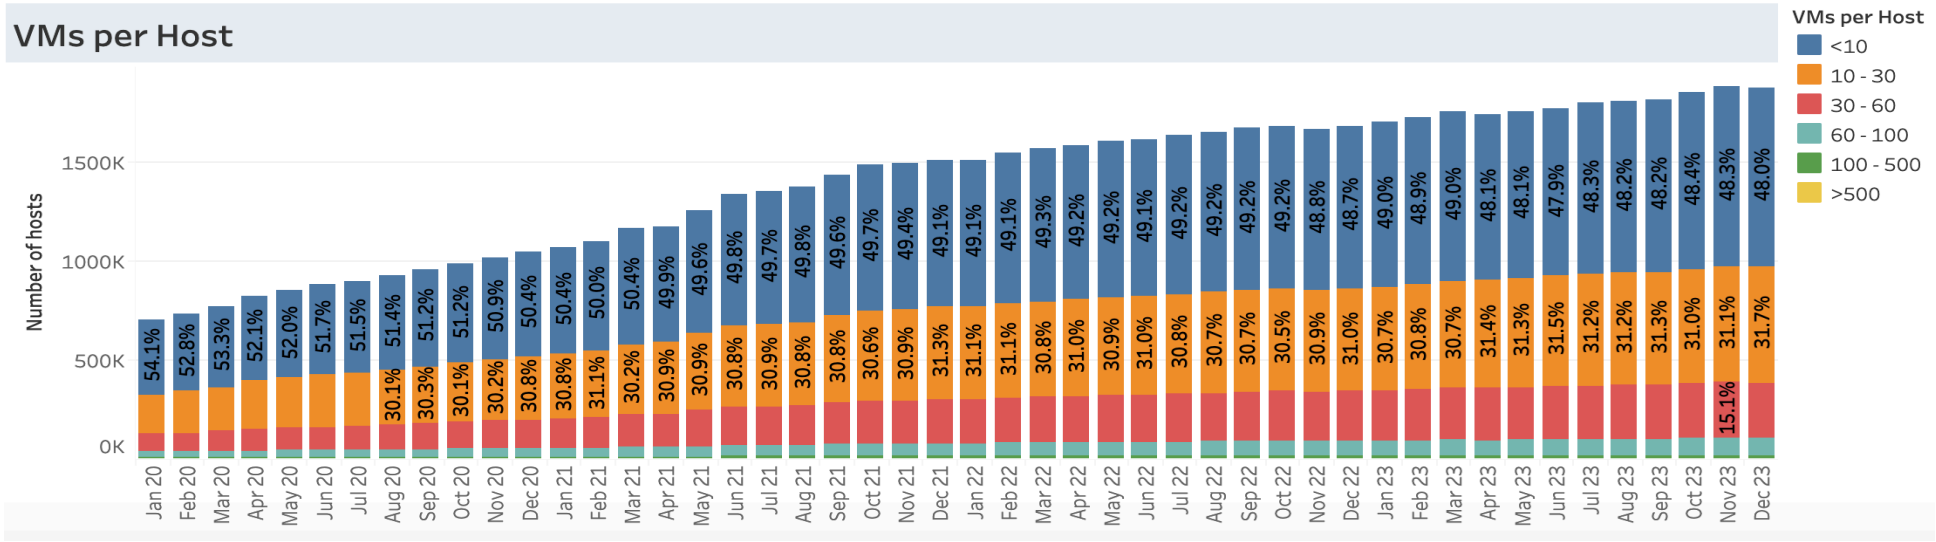

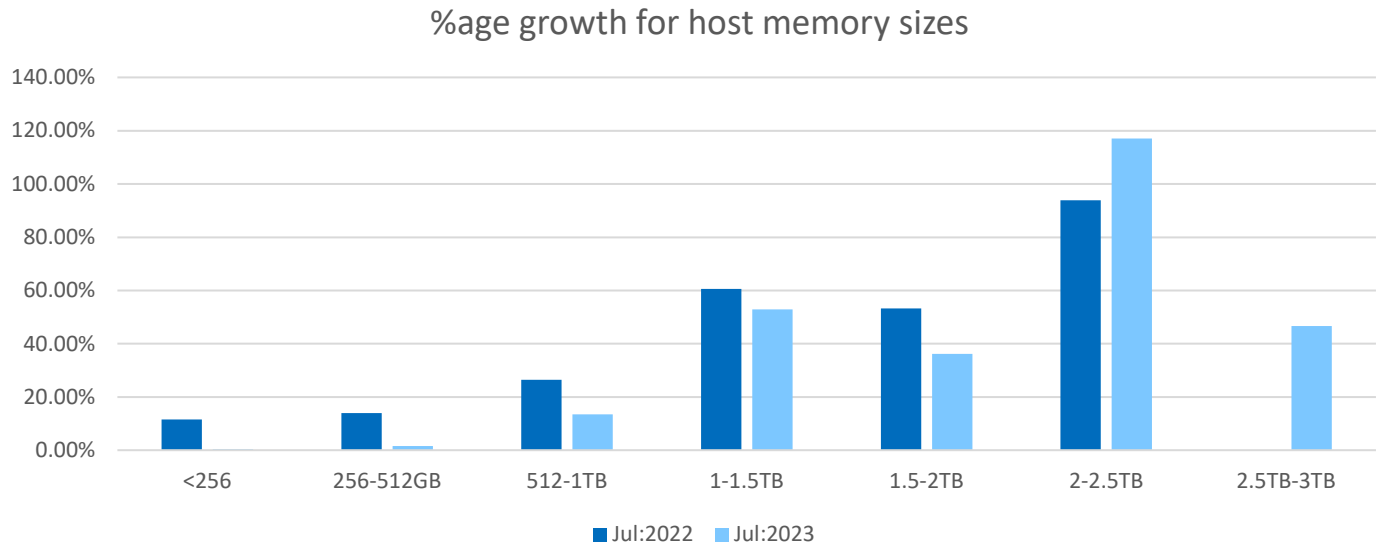

Host Memory in the larger buckets growing faster!!

Bridging the CPU-Memory imbalance

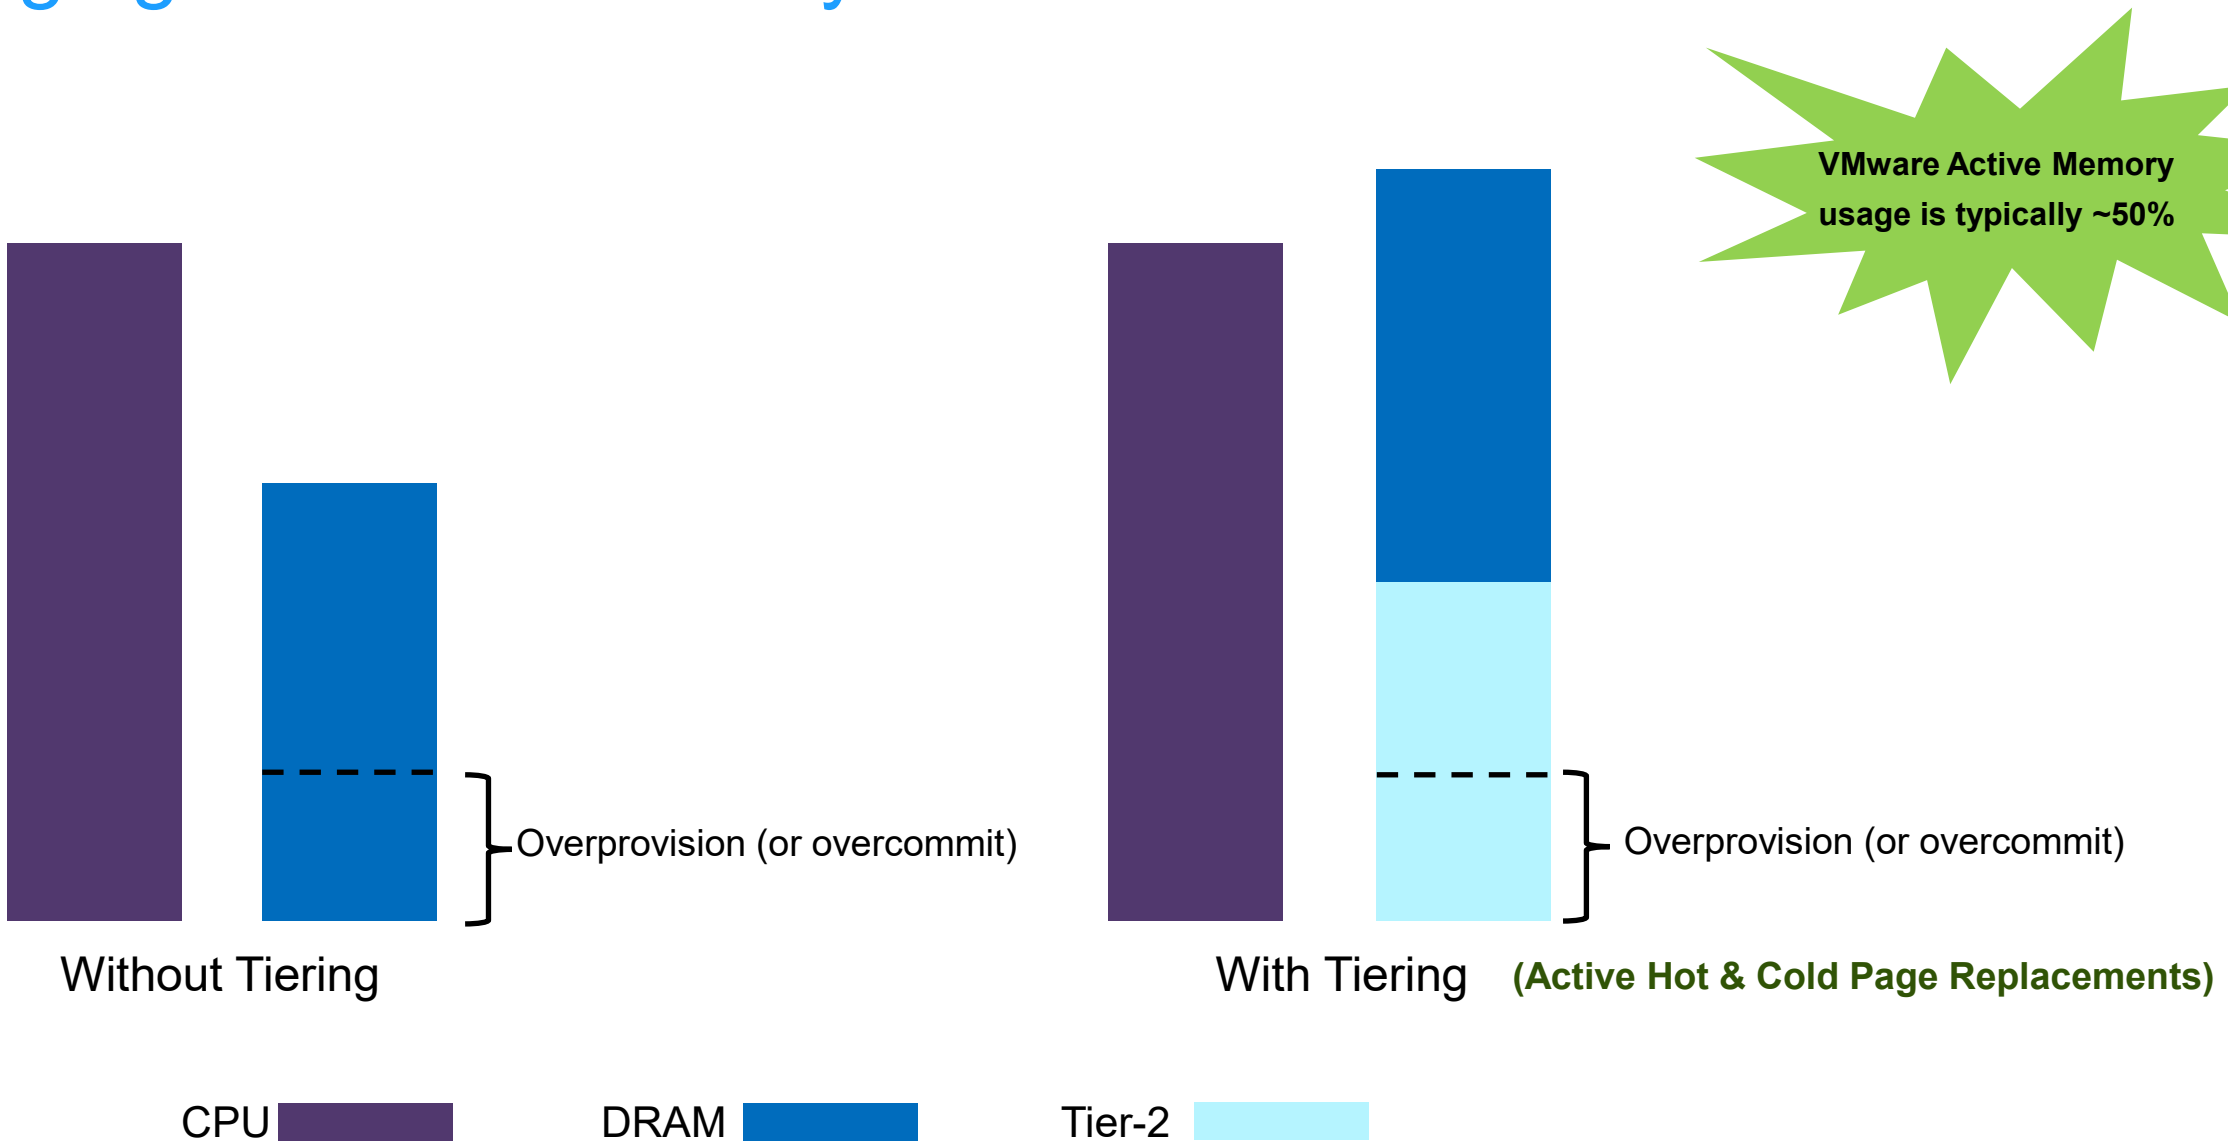

Memory Tiering Overview

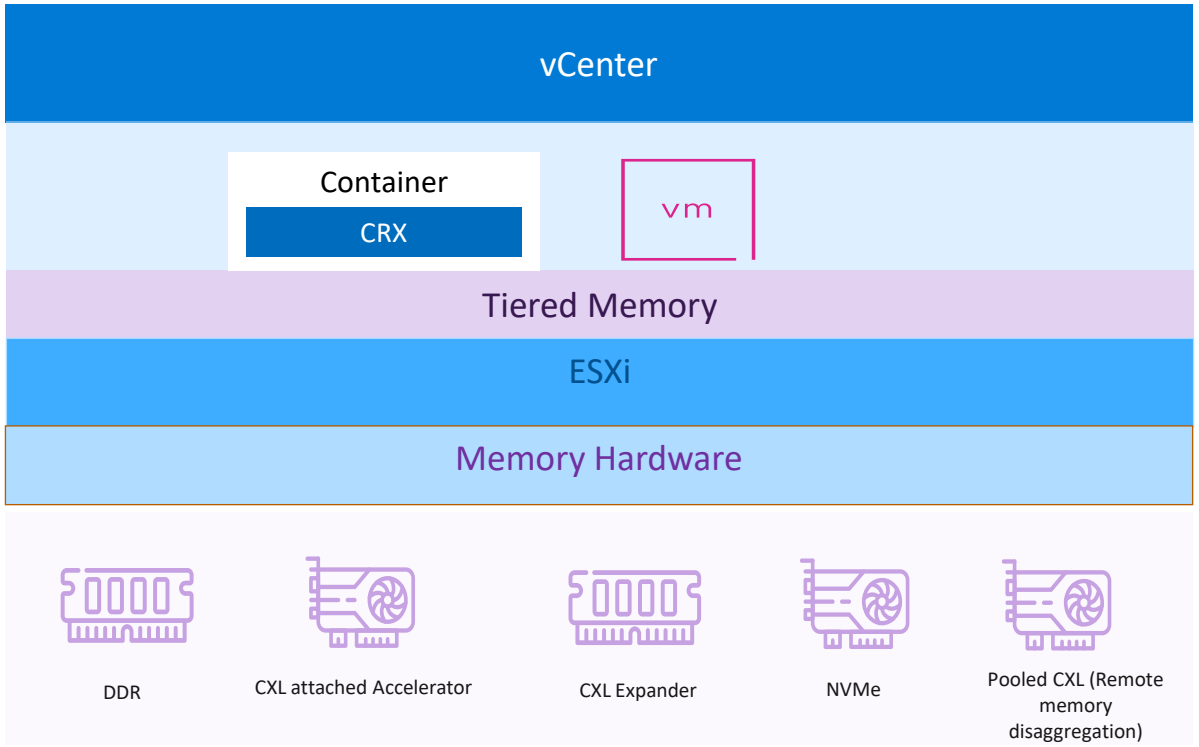

Reduce CapEx

- Provide lower TCO by (50%) DRAM substitution
- Improve capacity and bandwidth
- Improve host CPU core utilization
- Consolidate host resources

Reduce OpEx

- Expand memory on-demand
- Improve vMotion speeds
- Reduce failure probability
- Reduce host evacuation time for maintenance

Other future use-cases

- Pooling/sharing
- Remote mem tiering (stranded mem)
- Memory tracking/scrubbing & proactive maint.
- App-specific acceleration (AI/ML, Analytics)
- Encryption/dedupe/compression)

Customer Value Prop – How Memory Tiering reduces Total Server Costs

2TB DRAM (Socket attached)

Host configuration	(DRAM only)	List Price[1]
CPU	2 x 6348 (28C @ 2.6GHz)	$6,144*
DRAM/Pooling	2048GB (32 x 64GB)	$58,368
Lower-cost Memory	-	-
Emulex LPE36002 Dual port HBA	64Gb FC	$1,230
Storage capacity	2 X 3.84TB vSAS (7.68TB)	$3,600
Software	vSphere ENT+	$3,500
Total		$72,842

1.5TB DRAM + 0.5TB NVMe

Host configuration	(DRAM + NVMe)	List Price[1]
CPU	2 x 6348 (28C @ 2.6GHz)	$6,144*
DRAM	1536GB	$43,776
Lower-cost Memory (3:1)	512GB	$512
Emulex LPE36002 Dual port HBA	64Gb FC	$1,230
Storage capacity	2 X 3.84TB vSAS (7.68TB)	$3,600
Software	vSphere ENT+	$3,500
Total		$58,432

1TB DRAM + 1TB on CXL Accelerator

Host configuration	(DRAM + CXL Accelerator)	List Price[1]
CPU	2 x 6348 (28C @ 2.6GHz)	$6,144*
DRAM	1024GB (32 x 32GB)	$29,184
CXL Accelerator Board	1024GB	<$4048
Emulex LPE36002 Dual port HBA	64Gb FC	$1,230
Storage capacity	2 X 3.84TB vSAS (7.68TB)	$3,600
Software	vSphere ENT+	$3,500
Total		$47,706

Up to 36% Cost Reduction with 1:1 DRAM:CXL Accelerator
(19% with NVMe tiering with 4:1)

*Estimated (costs and results may vary depending on server vendor)

1 https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r750/pe_r750_14794_vi_vp?configurationid=b605e5ac-c8b9-4578-b0e2-7d9b15772b04

Solution 1: NVMe-based tiering

Solution 1: NVMe-based tiering

Tech Preview: 8.0U3

GA: 9.0

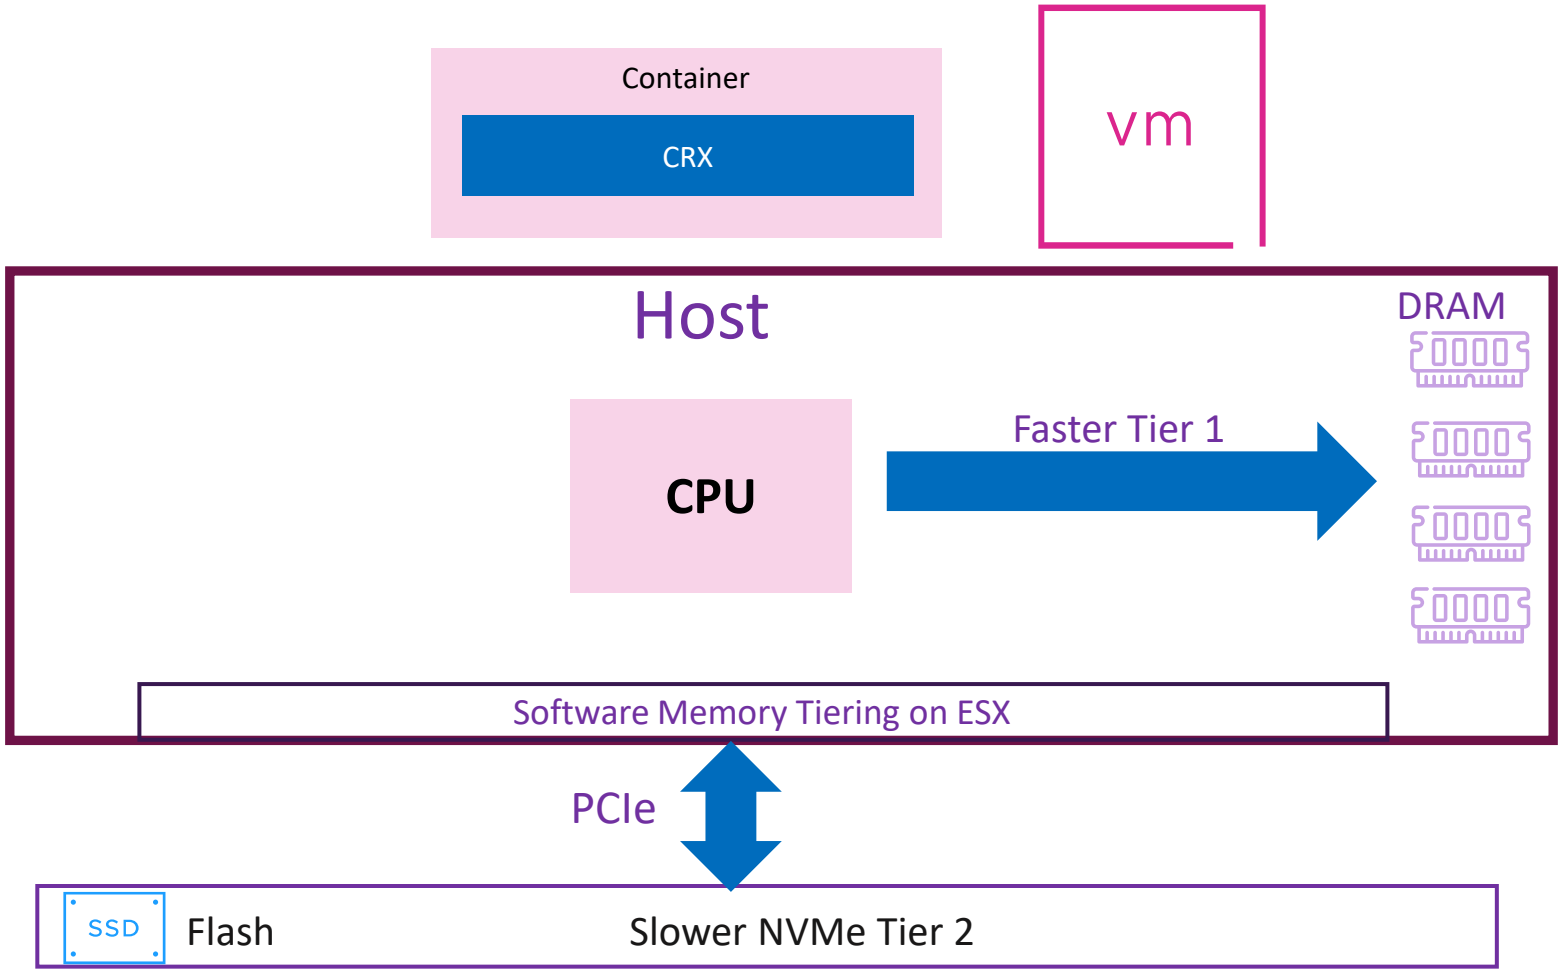

Solution 2: CXL Accelerator

- VMware Hardware-Software co-design
  - CXL Type-2 Accelerator (Intel GNR)
- 3/4 length single slot/width PCIe form-factor
  - Intel Agilex FPGA based design
- Suited for a wide variety of workloads
  - VDI, Databases (Traditional & In-memory), Web Services etc.
- Seamless deployment and maintenance
  - Intel vanilla BIOS
- Up to 4TB of additional memory
  - 1:1 DRAM:CXL Accelerator configuration

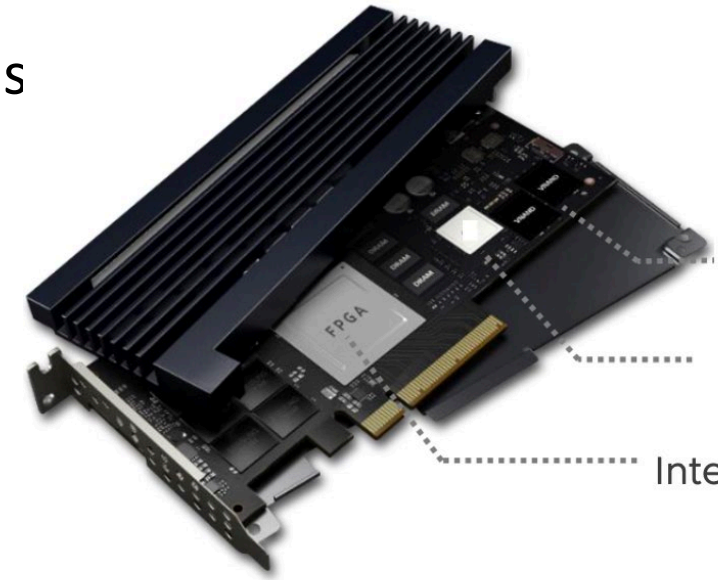

Vision for Memory Disaggregation – Building for the future

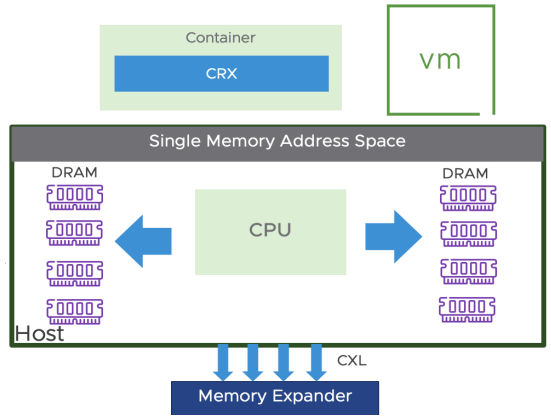

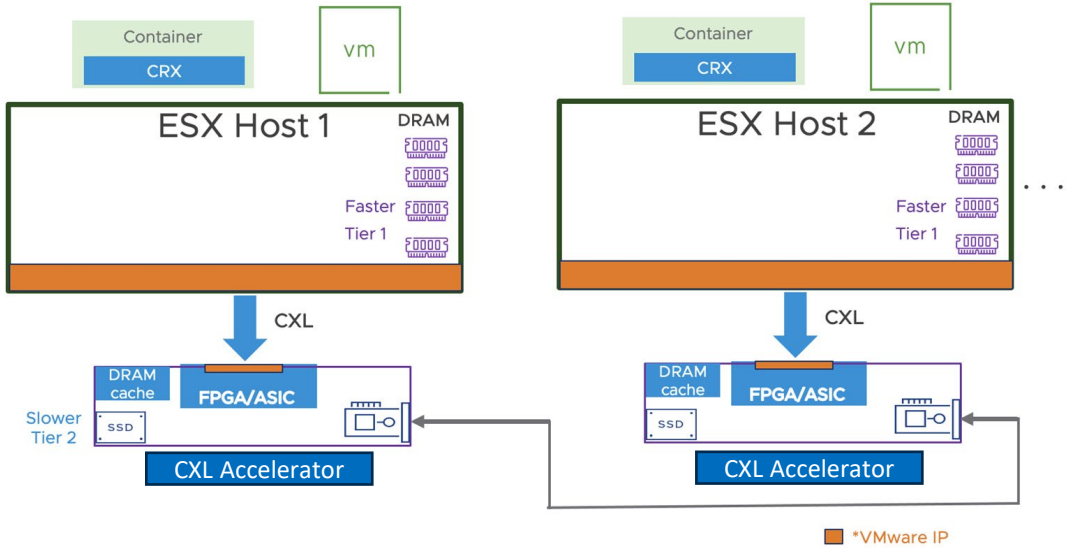

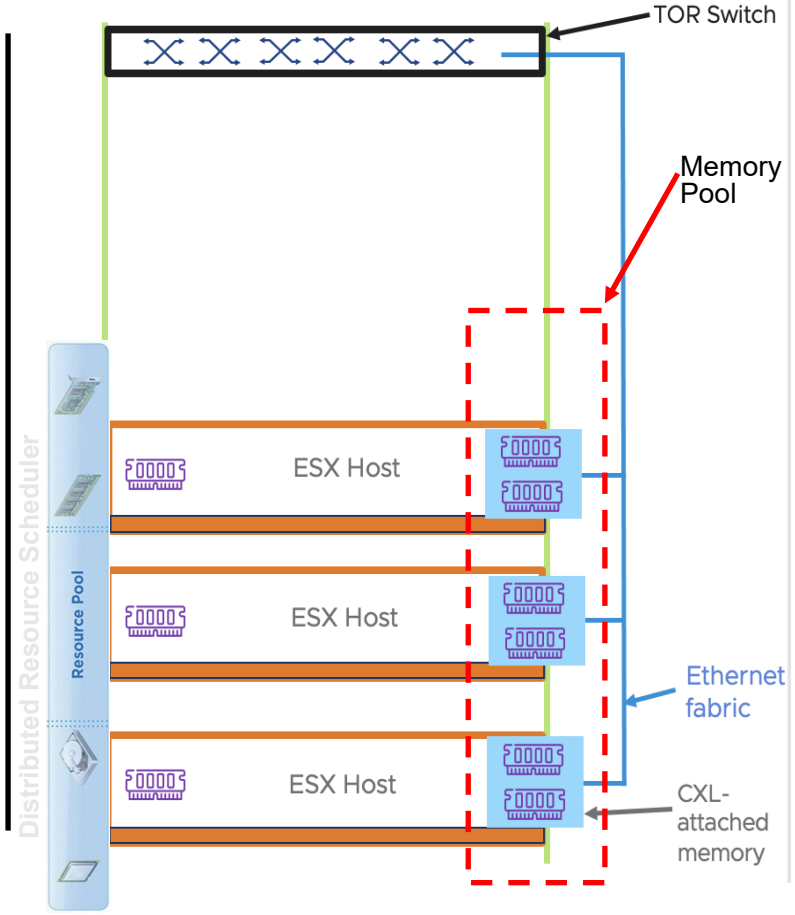

- Memory Tiering
- vMotion Perf

CXL Memory Expansion/Tiering

- Memory Pooling/Disaggregation
- Cloud and Edge Computing
- Application Memory (AI models, analytics..)
- Offloads - storage, networking use-cases

CXL Type-2 Memory Accelerator

- Tiering as a platform feature
- Fully integrated hardware virtualization

Disaggregated/Composable Memory

Oracle on Memory Tiering - Metrics

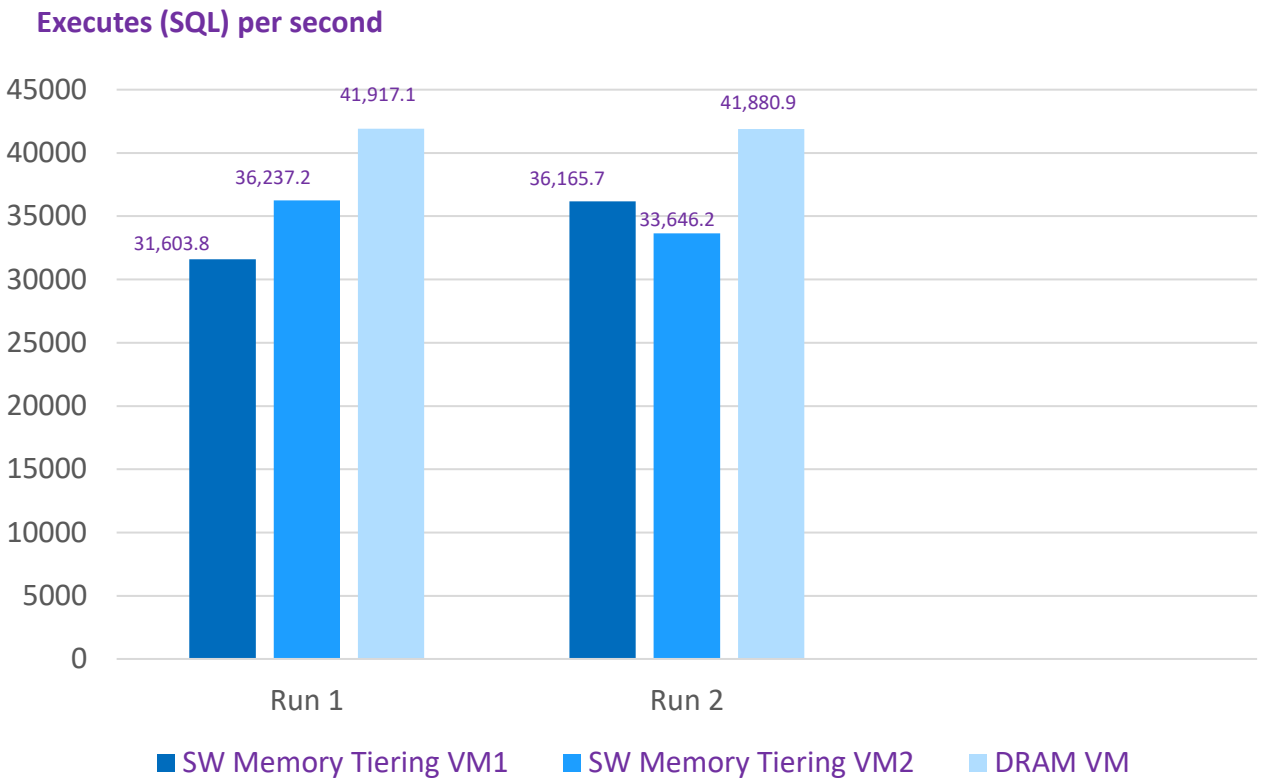

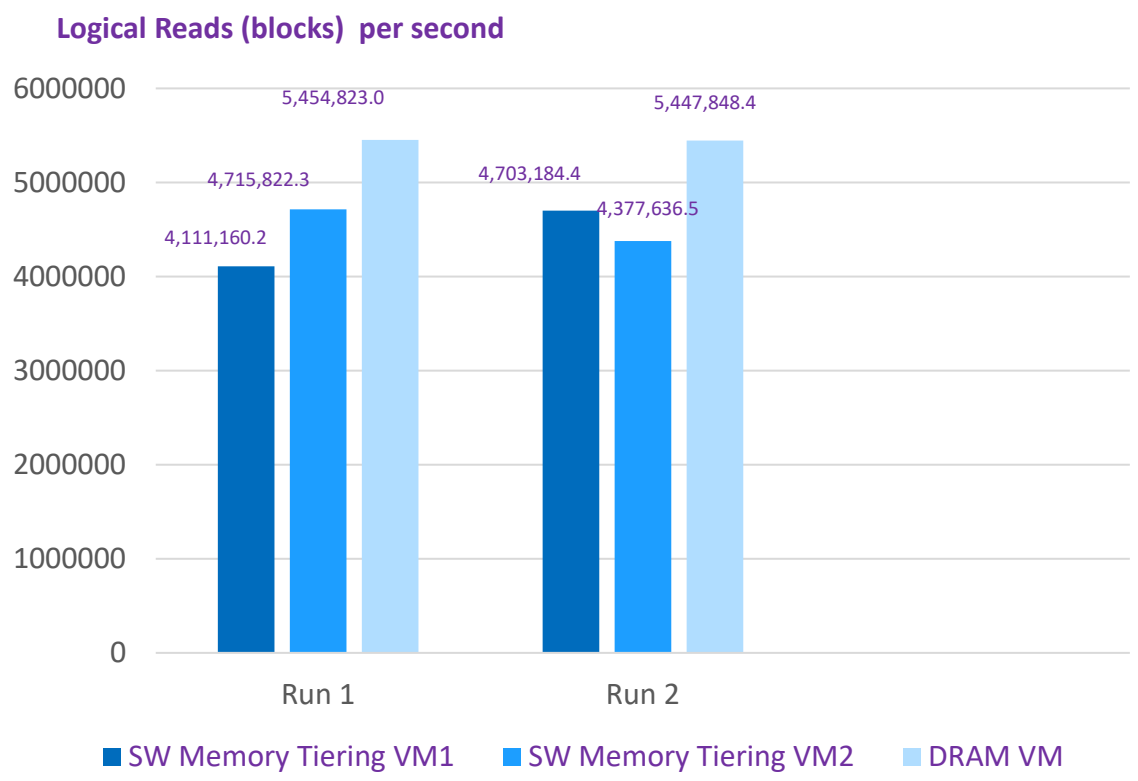

- Load Generator chosen as SLOB 2.5.4.0
  - UPDATE_PCT=0 - READ only test
  - RUN_TIME=1200 secs (20mins)
- Test Results
  - **Executes(SQL) / second**
  - Run 1
    - **Aggregate SW Tier VM1+VM2 = 69,841/sec V/S DRAM Mode VM** - 41,917.1/sec
  - Run 2
    - **Aggregate SW Tier VM1+VM2 = 69,811.9/sec V/S DRAM Mode VM** - 41,880.9/sec

- Test Results
  - **Logical Reads (blocks) per second**
  - Run 1
    - **Aggregate SW Mem Tier VM1 + VM2 = 8,826,982.5/sec**
    - **DRAM Mode VM - 5,454,823.0/sec**
  - Run 2
    - **Aggregate SW Mem Tier VM1 + VM2 = 9,080,820.9/sec**
    - **DRAM Mode VM - 5,447,848.4/sec**

Summary for preliminary workload performance with each solution

NVMe Tiering

1. LoginVSI (VDI Benchmark):
   - 1:1 ratio: Compared to DRAM baseline –
     – 100% Increase in VM density:128->256 VMs with 1-2% performance drop, 5-6 additional cores usage
2. VMMark (Mixed DB, In-mem DB, MySQL, Web services, varied VM sizes (2-256GB), vCPUs (2-8)):
   - 4:1 ratio: (4 DRAM & 1 NVMe) Compared to DRAM baseline –
     – 33% increase in tiles (Tile=Mixed bundle) 3 tiles -> 4 tiles. Score (Measure of performance) Not affected.
     – More workloads are being planned

CXL Accelerator

1. HammerDB with Oracle TPC-H profile (End-to-End solution with NVMe usage)
   - Running on a 48GB VM – (1:1 - DRAM:CXL Accelerator)
     - 85-90% performance of DRAM
2. More workloads are being planned & performance optimizations in pipeline

Summary

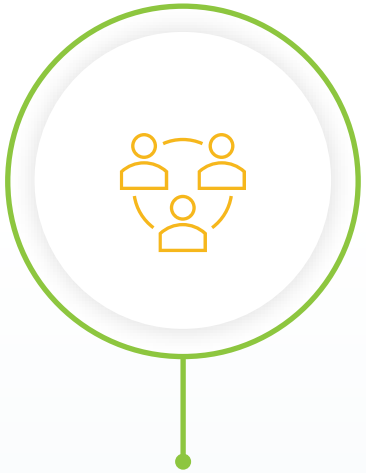

VMware has been collaborating with CXL partners to bring tiering to address cost, scale, and density challenges

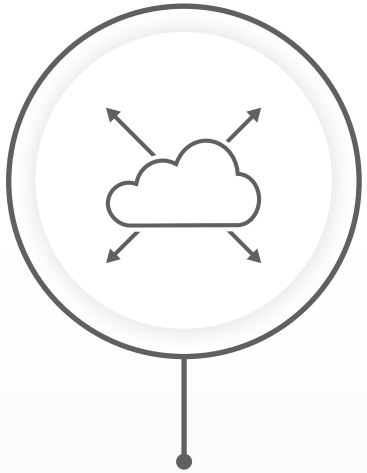

Business transformation is leading to larger datasets and real-time analytics that requires larger memory capacity and better CPU-memory balance leading to better server consolidation

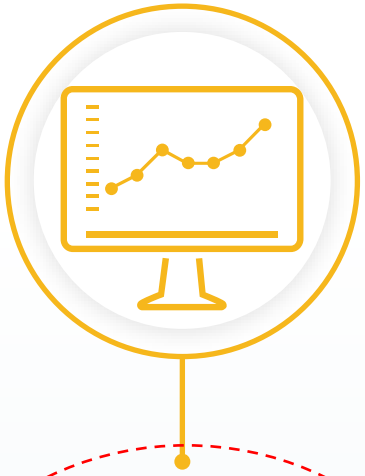

VMware is bringing multiple memory solutions – CXL expansion, NVMe tiering, and CXL-based accelerator solutions that provide large TCO benefits without impacting performance

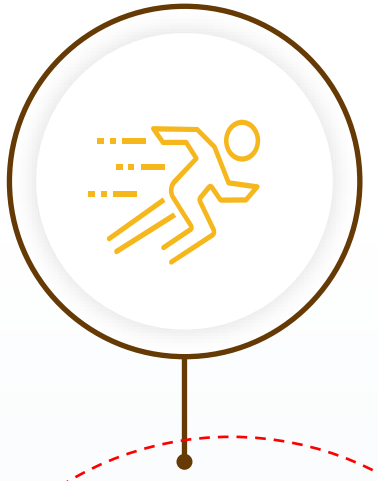

Memory Tiering will bring scale without adding any operational complexity. Tiering is also ready for future technologies that can bring pooling and disaggregation

Call to action

- Arvind Jagannath (Arvind.jagannath@Broadcom.com)
  - Please contact Arvind to collaborate/partner on these solutions
  - Also, for inquiries on POCs (VMware is ready for POCs on these solutions)

- Sudhir Balasubramanian (sudhir.balasubramanian@broadcom.com)
  - Please contact Sudhir for any questions on workloads/performance

Please take a moment to rate this session.

Your feedback is important to us.