

CXL Readiness with vSphere Virtualization



EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

Trends Vs Customer Needs



Trends

Explosive Growth of the Data



Desire to get more out of the data



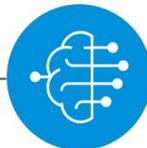
More data need to be processed in real-time



Software has led the innovation in the cloud, Hardware is catching up



Digital Transformation



Customer Needs



Need to scale infrastructure to address data growth



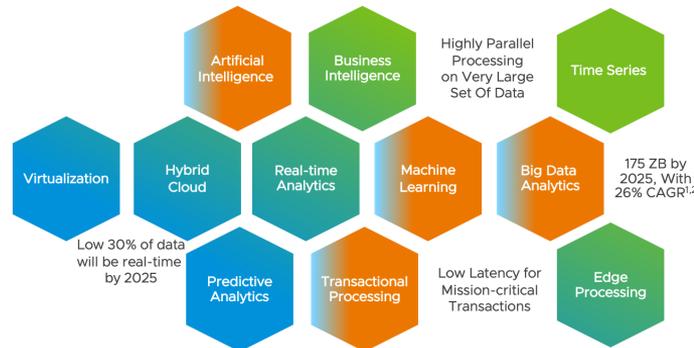
More in-memory computing to process faster



DRAM is expensive and lacks high densities



Need Enterprise class monitoring and remediation

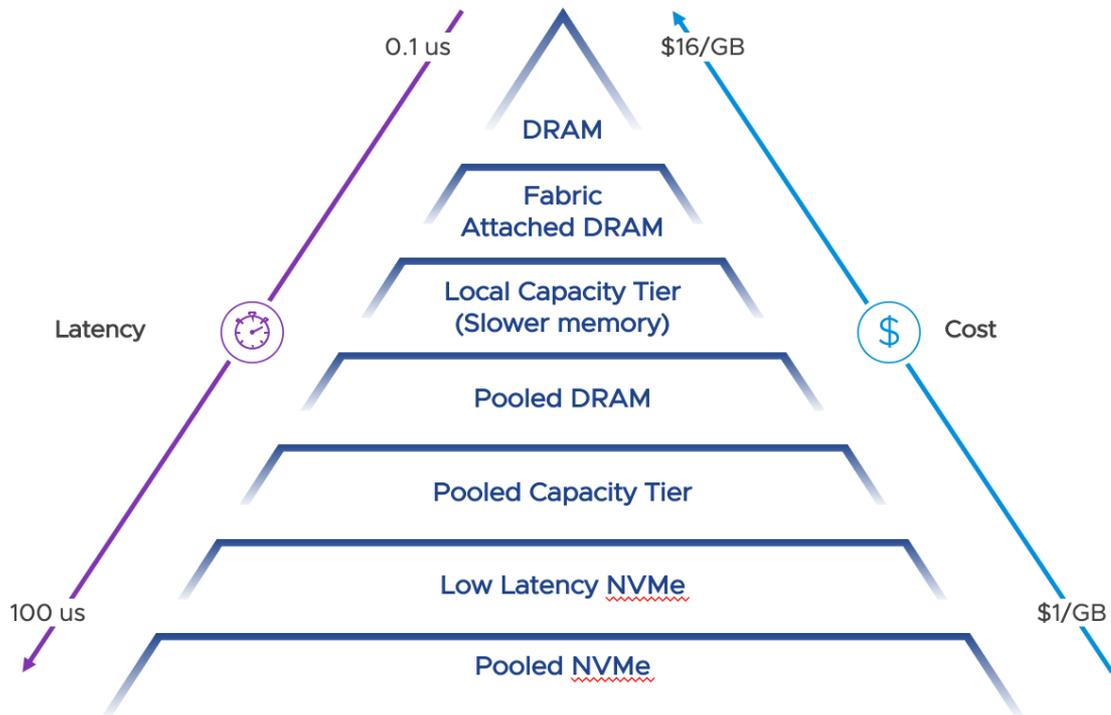


OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Problem Statement: Higher costs and lower utilization



OCP
GLOBAL
SUMMIT

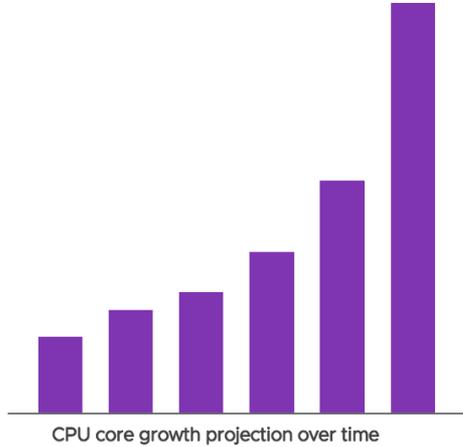
OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

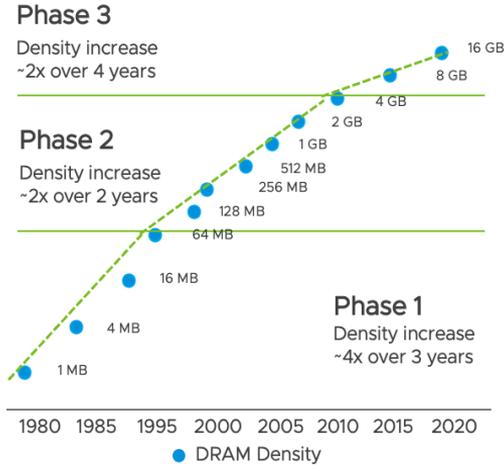
Problem Statement: Scale Challenges



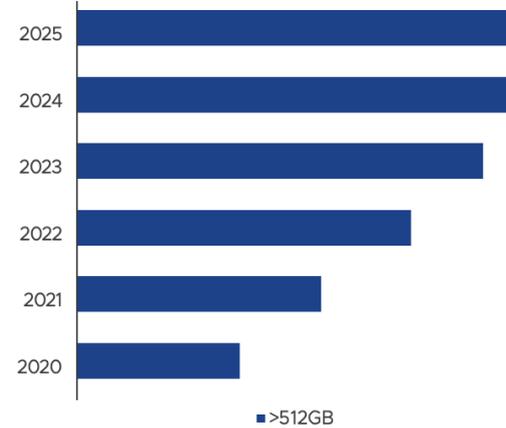
Compute demand accelerating



DRAM not scaling
Compound annual growth rate in data²



Large-memory systems growing in prevalence



Data-intensive workloads need hot data in memory, but DRAM is expensive and has limited capacity
Enterprise data is vulnerable, but security and encryption can compromise performance

² IBM, "3D NAND Technology – Implications to Enterprise Storage Applications," 2015 Flash Memory Summit, 2014. flashmemorysummit.com/English/Collaterals/Proceedings/2015/20150811_FM12_Yoon.pdf.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.



VMware's Big Memory Vision starts with Software Tiering

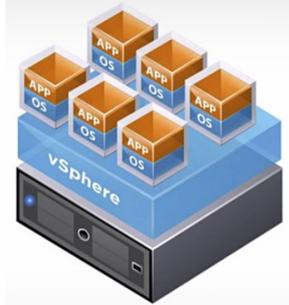


OCP
GLOBAL
SUMMIT

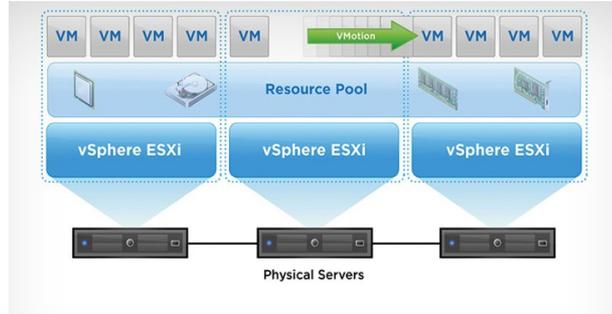
OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

VMware Competencies



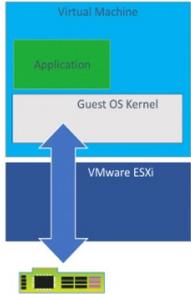
Virtualization - Transparent tiering



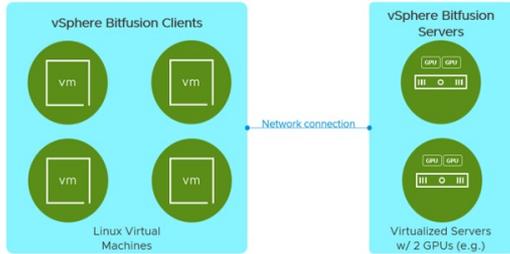
Resiliency - Cluster-wide DRS load balances/mitigate risks



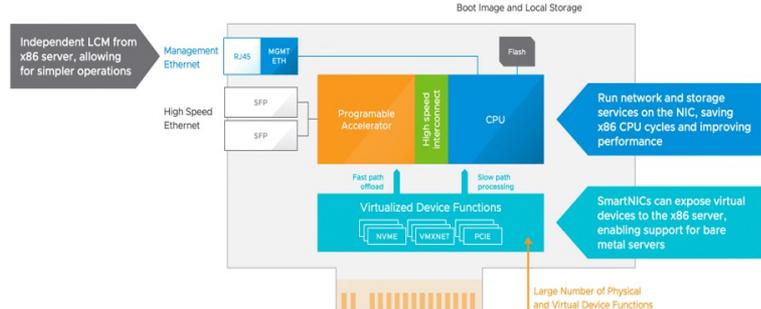
Integration - Strong Partner Ecosystem



Passthrough devices



GPUs, sharing and Assignable hardware



SmartNICs and Accelerators



OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

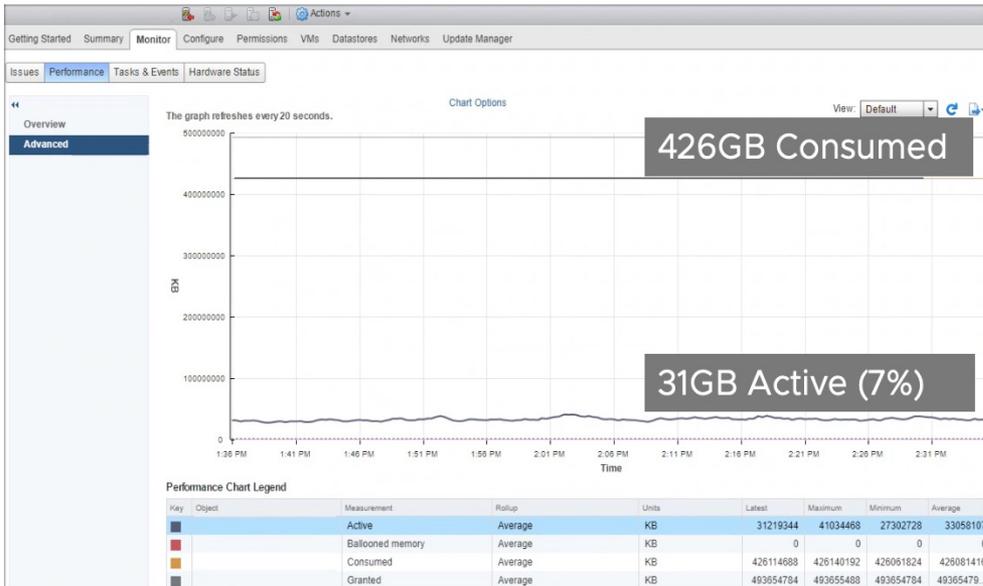
Active Memory in vSphere



Data from VMware customer production environments tends to show a lot of stranded memory



Active memory < 25-40% of consumed/physical memory is good fit for memory tiering



Results may vary



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.



Example – How Software Tiering can reduce Total Cost

2TB DRAM

2.5TB Memory (DRAM + Tier 2)

Host configuration DRAM only		List Price ¹
CPU	2 x 6348 (28C @ 2.6GHz)	\$6,144*
DRAM	2048GB (32 x 64GB)	\$58,368
Slower Memory	-	-
Storage capacity	15 X 3.84TB vSAS (57TB)	\$27,000
Cache drives	3 X P5800X 400GB	\$6,780*
Software	vSAN + vSphere ENT+	\$28,563
Total		\$126,855*

Host configuration DRAM + PMem		List Price ¹
CPU	2 x 6348 (28C @ 2.6GHz)	\$6,144*
DRAM	512GB (16 x 32GB)	\$14,592
Slower Memory	2048GB (16 X 128GB)	\$17,121
Storage capacity	15 X 3.84TB vSAS (57TB)	\$27,000
Cache drives	3 X P5800X 400GB	\$6,780*
Software	vSAN + vSphere ENT+	\$28,563
Total		\$100,200*

20% Platform level Cost Reduction

*Estimated
 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
 Your costs and results may vary



OCTOBER 18-20, 2022
 SAN JOSE, CA

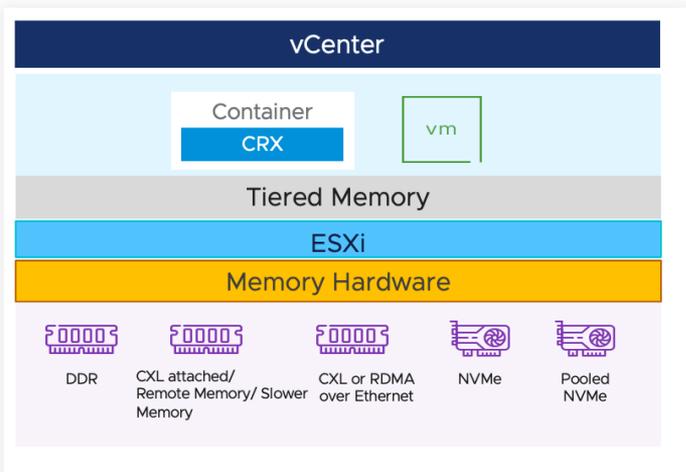
EMPOWERING OPEN.

VMware Software Tiering



Value over traditional tiering approaches

- **Transparent – Single volatile memory address**
 - No Guest or Application changes
 - Run any Operating System
 - ESX internally handles page placement
- **DRS and vMotion to mitigate risks**
 - Tiering heuristics fed to DRS
- **Ensure Fairness across workloads**
 - Consistent performance
- **Zero Configuration changes**
 - No special tiering settings
- **Minimum Performance Degradation**
- **Processor specific monitoring**
 - vMMR monitors at both VM- and Host-levels



Benefits

- Higher Density - core utilization
- Lower TCO
- Larger bandwidth



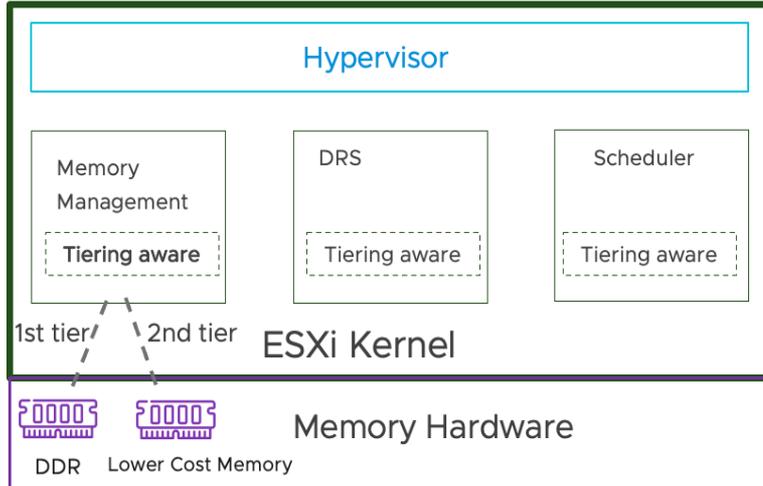
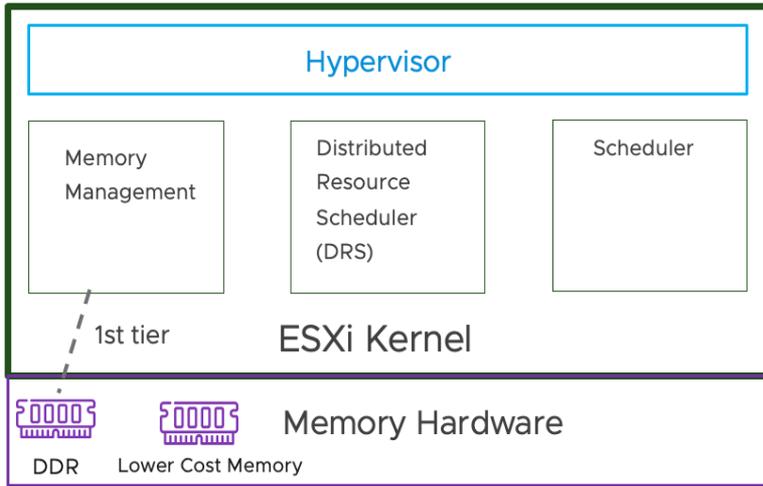
OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

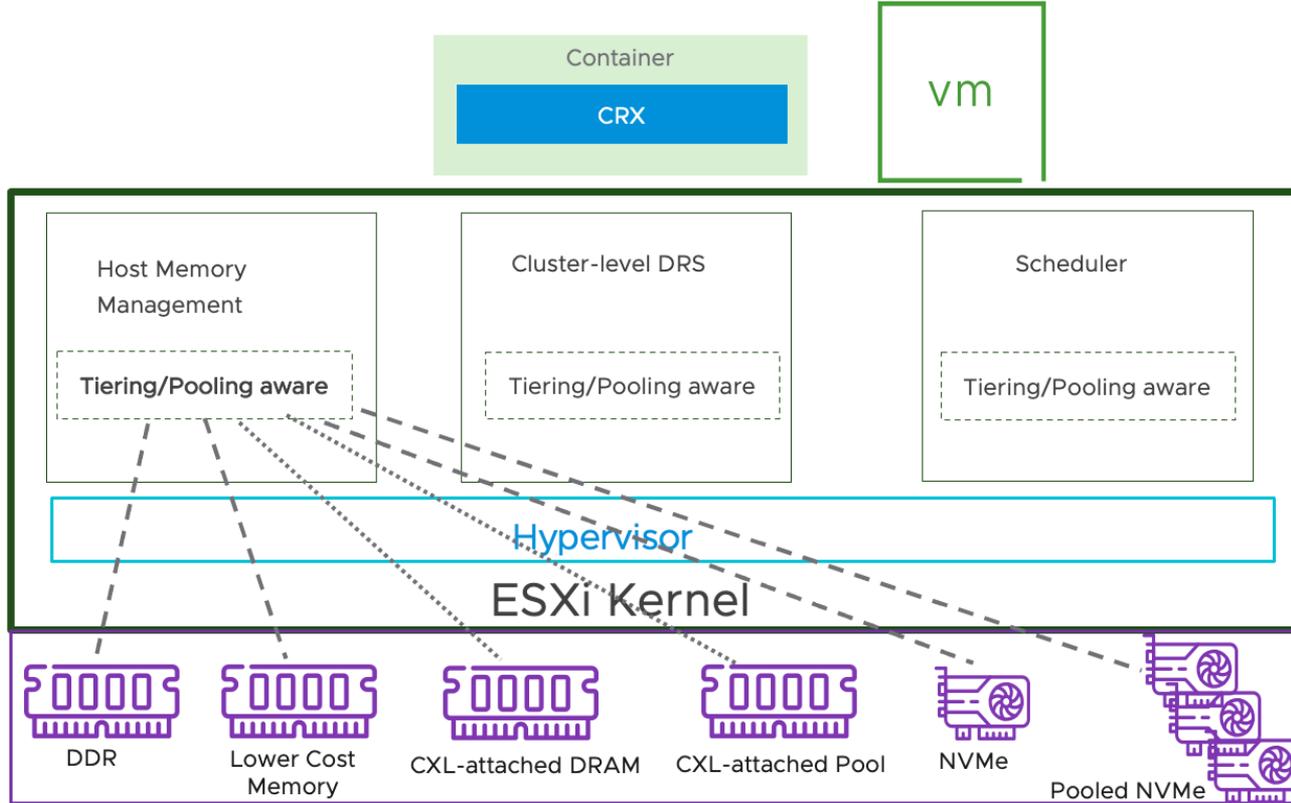
Software Tiering: How Does it Work?

Phase-1: Host Local memory tiering



Software Tiering: How Does it Work?

Future



OCP
GLOBAL
SUMMIT

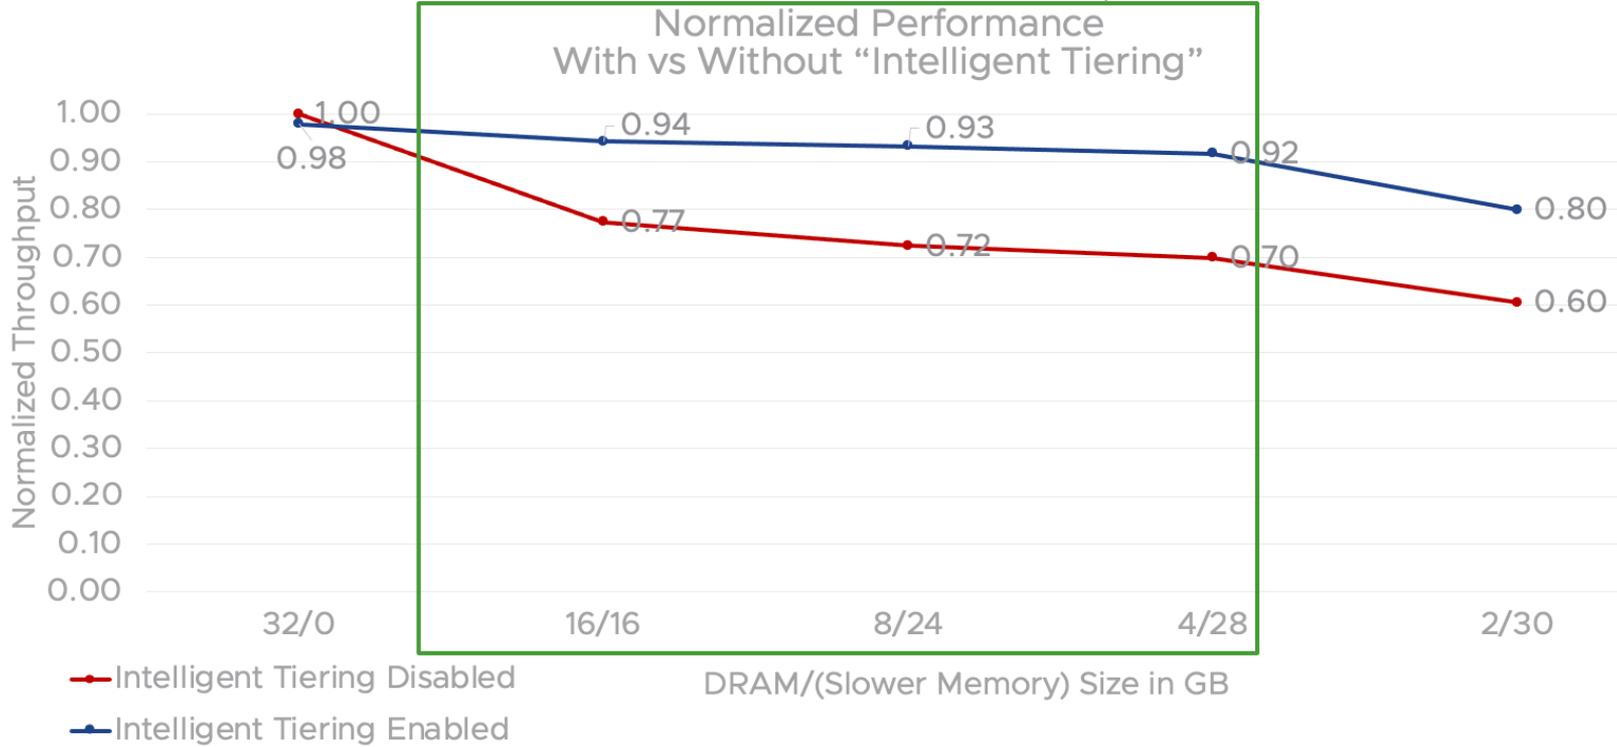
OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

HammerDB using SQL Server

Performance Results (Higher is better)

Significantly reduce memory cost without significant drop in performance

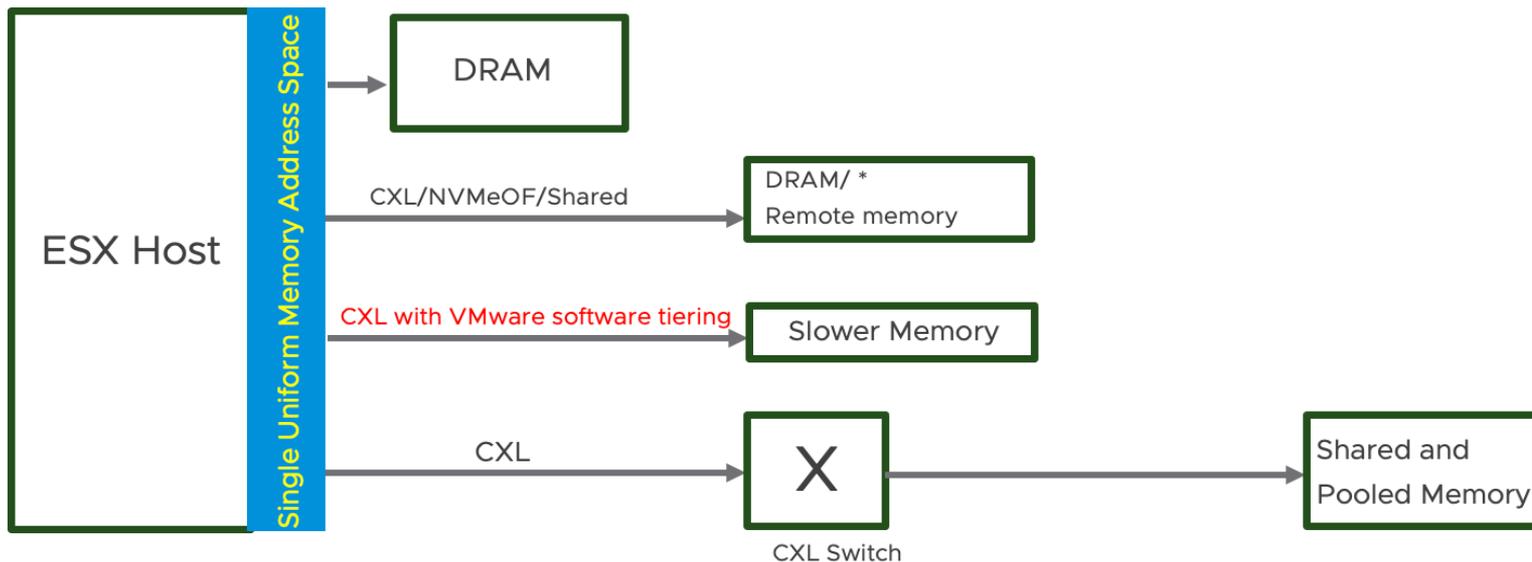


OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Software Tiering and CXL

Future



**Prototyping with CXL1.1*



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Key Use Cases emerging with CXL



- VMware working with CPU and Memory/NVMe vendors
- Lower TCO – Use combinations of lower cost slower memory with DRAM
 - More capacity
 - **Needs Software Tiering!!**
- Memory Expansion with NUMA-like latencies
 - More capacity
 - Flat (non-tiered) expansion
- Memory Expansion for higher bandwidth
- Memory sharing and pooling across multiple hosts in a cluster
 - Memory appliances providing both bandwidth and utilization
- CXL switching and shared access (future)



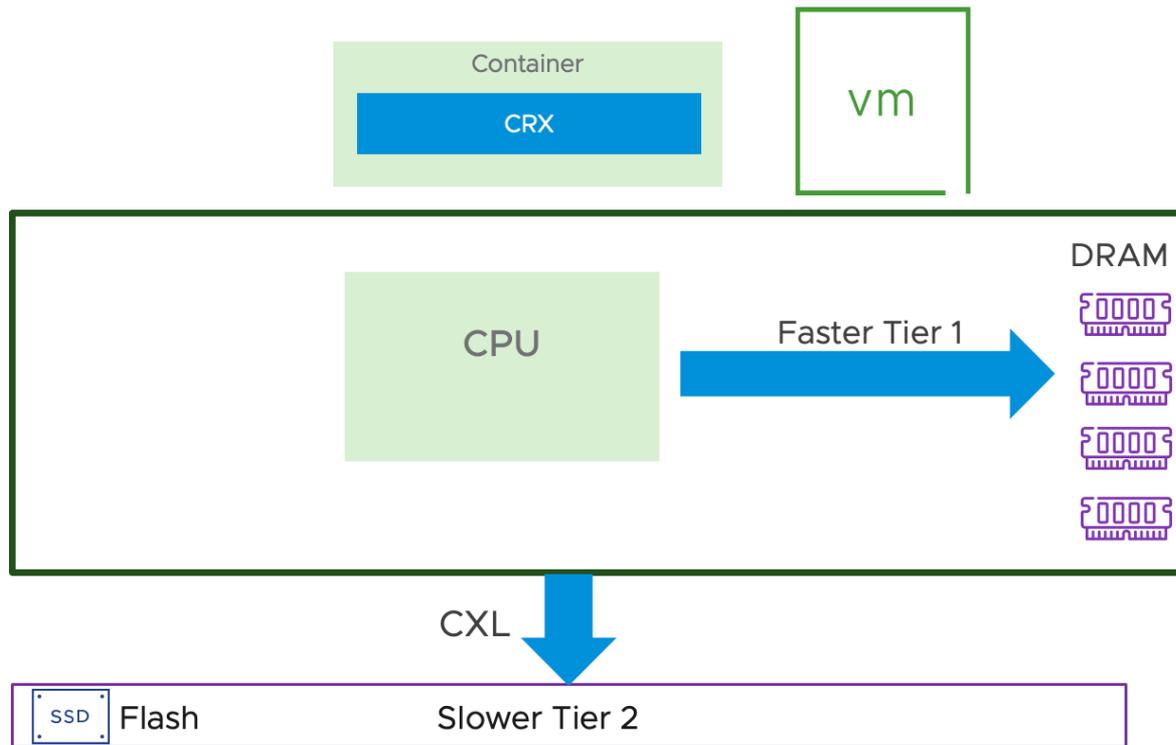
OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

CXL Use case – VMware Software Tiering

“Memory-SSD”



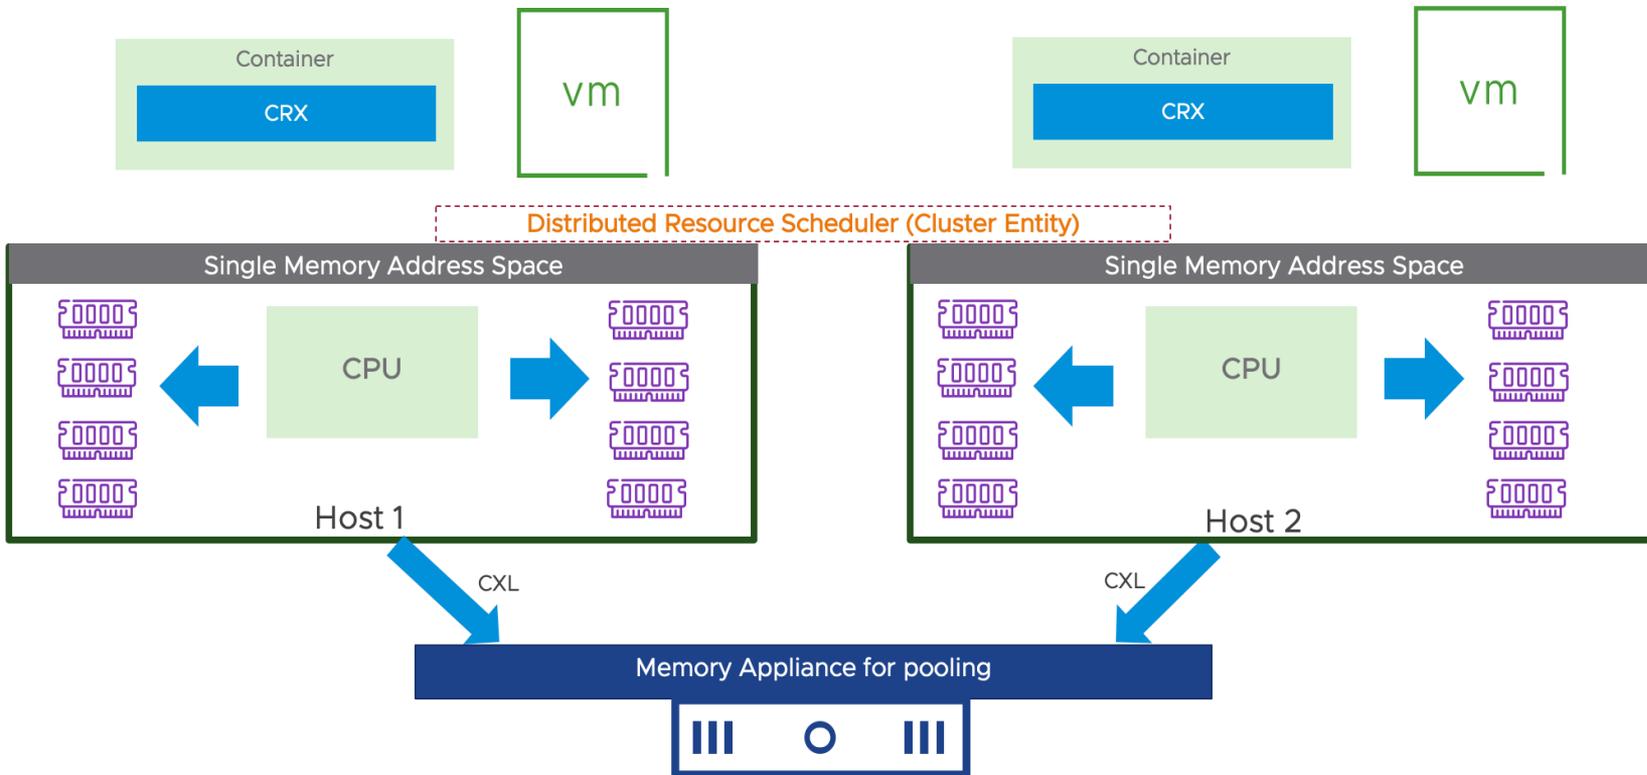
OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

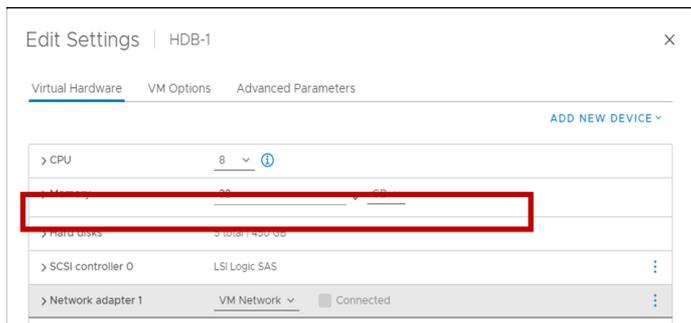
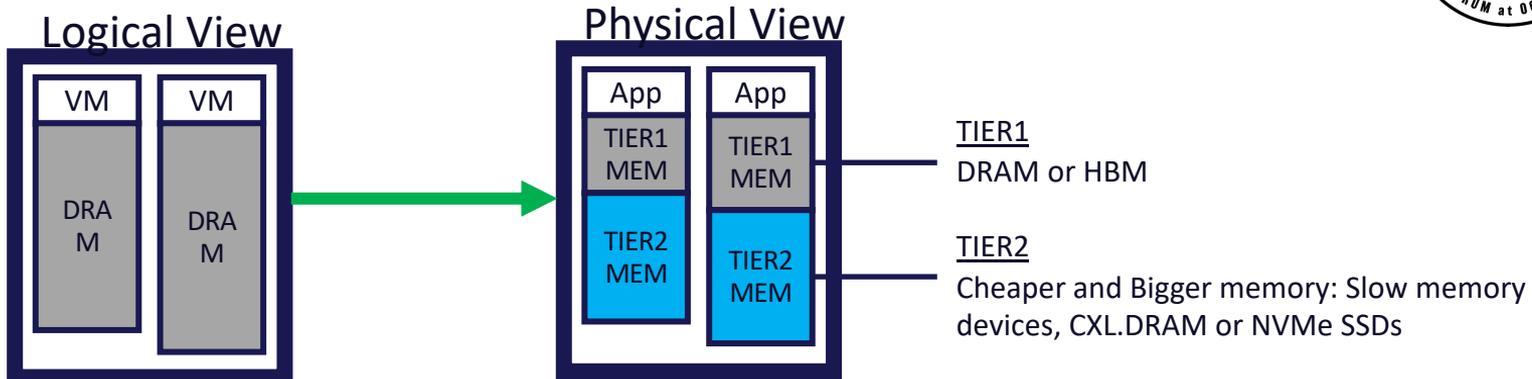
CXL Use case

Memory Appliance – Improve Cluster Utilization



Software Memory Tiering in vSphere

Built in vSphere, requires no modifications to applications or Guest Operating Systems



- ✓ vSphere decides what tiers to use and when
- ✓ Memory available to the host is sum of all memory tiers

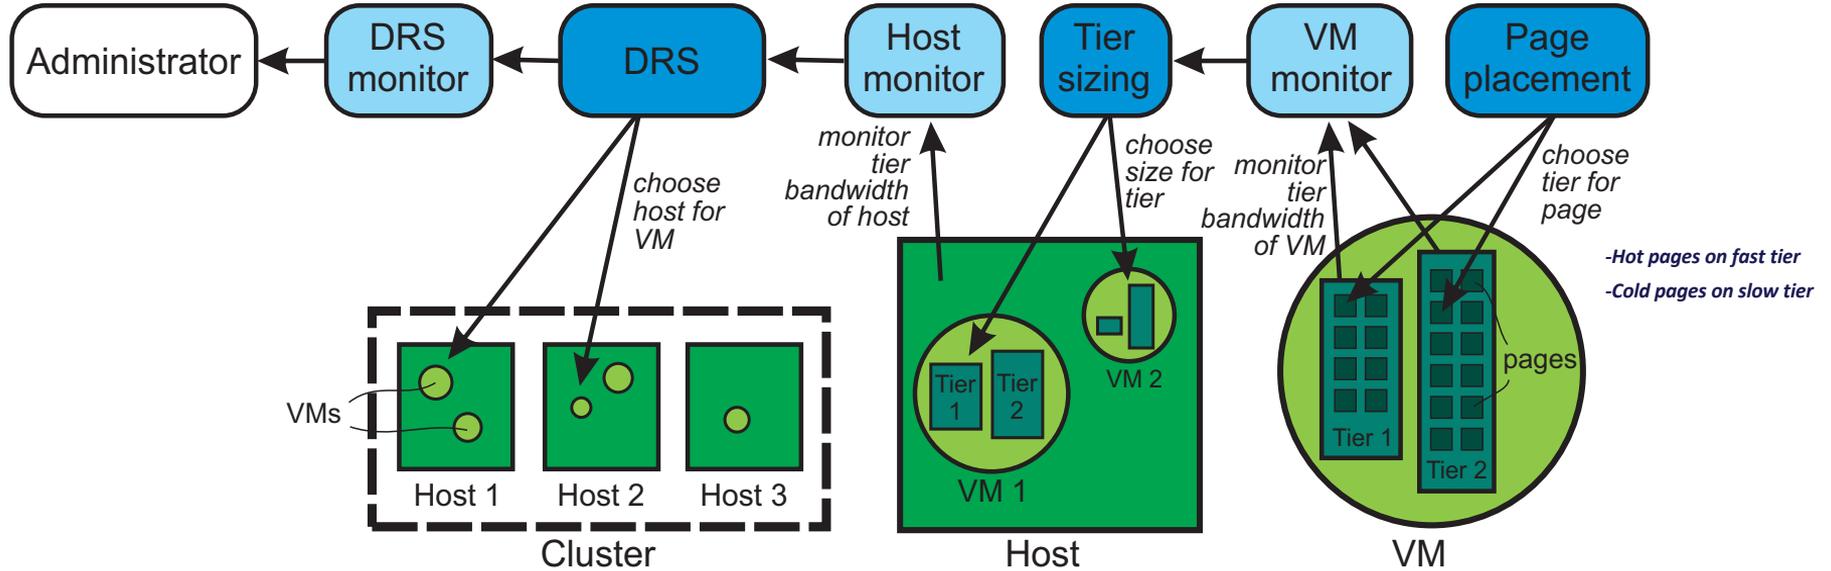


OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

How it all fits together in a vSphere Cluster?



Tier Target Size = f (BW of Tiers, Num_VMs, Active)



OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Thank you!



EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

