SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Providing Capacity & TCO benefits to Applications using VMware Software Memory tiering

Sudhir Balasubramanian – Sr Staff Solution Architect, VMware

Arvind Jagannath - Sr Product Line Manager, Cloud Platform, VMware

Agenda

- Motivation
- Memory Tiering and CXL Use-cases
- Optimizing Oracle Workloads – Use Case
- Key Takeaways

VMware Competencies

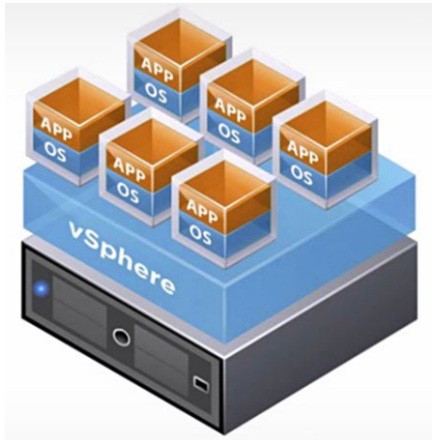

Virtualization ideal for transparent memory tiering

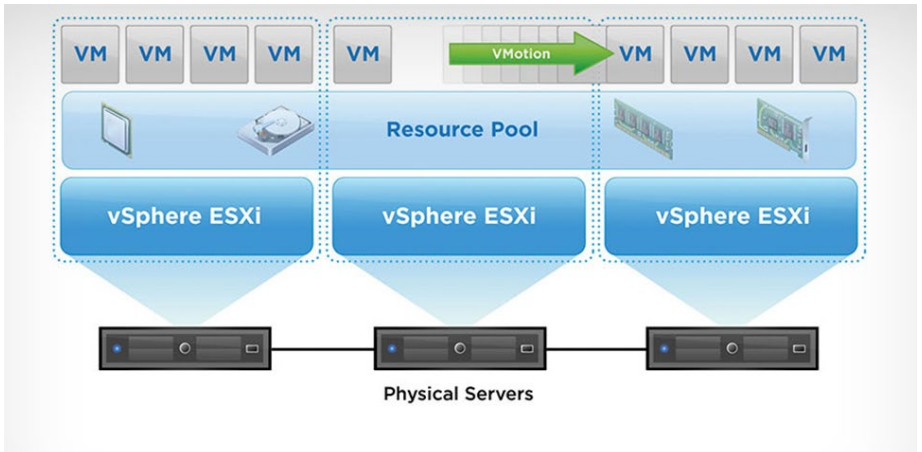

Resiliency - Cluster-wide DRS load balances/mitigate risks

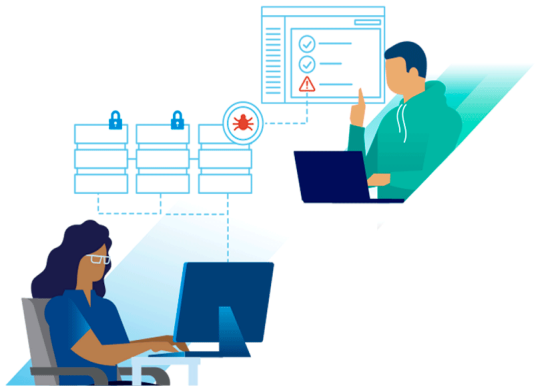

Strong Ecosystem of partners

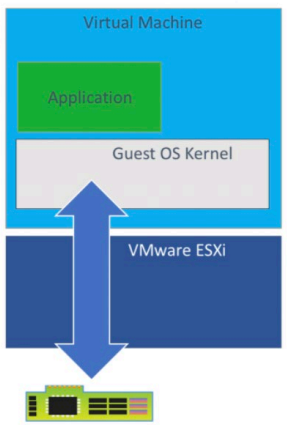

Passthrough devices

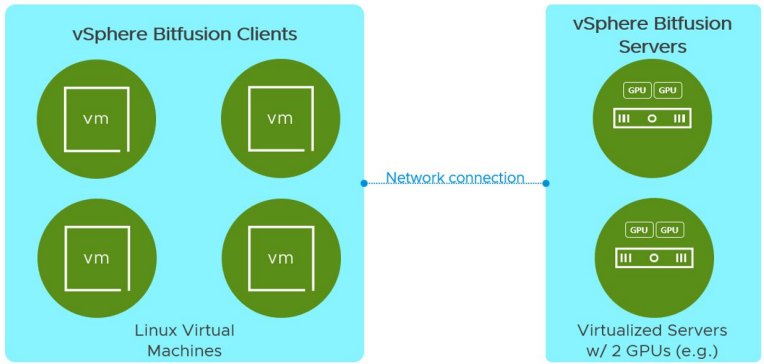

GPUs, sharing and Assignable hardware

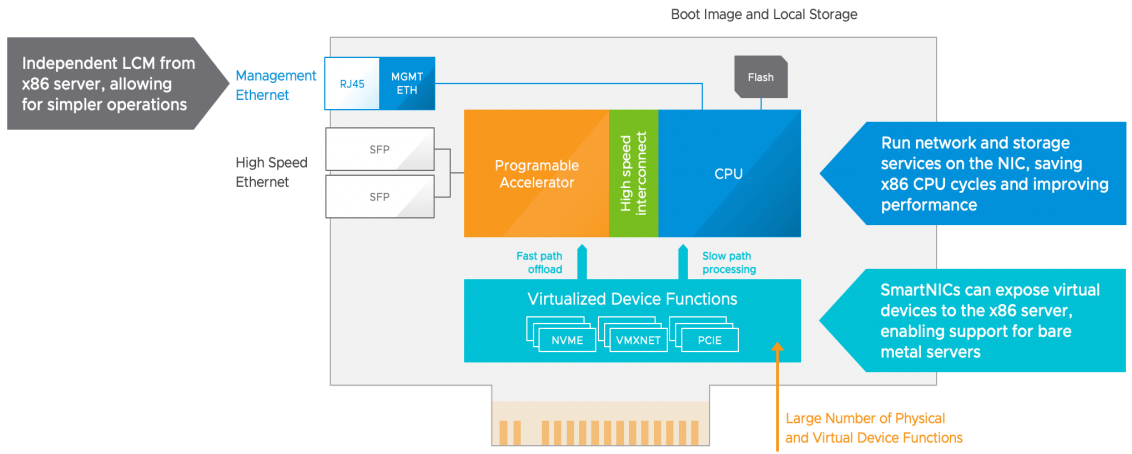

SmartNICs and Accelerators

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Memory Tiering and CXL Use-cases

VMware's Big Memory and Disaggregation Vision starts with CXL and Tiering

VMware Memory Tiering

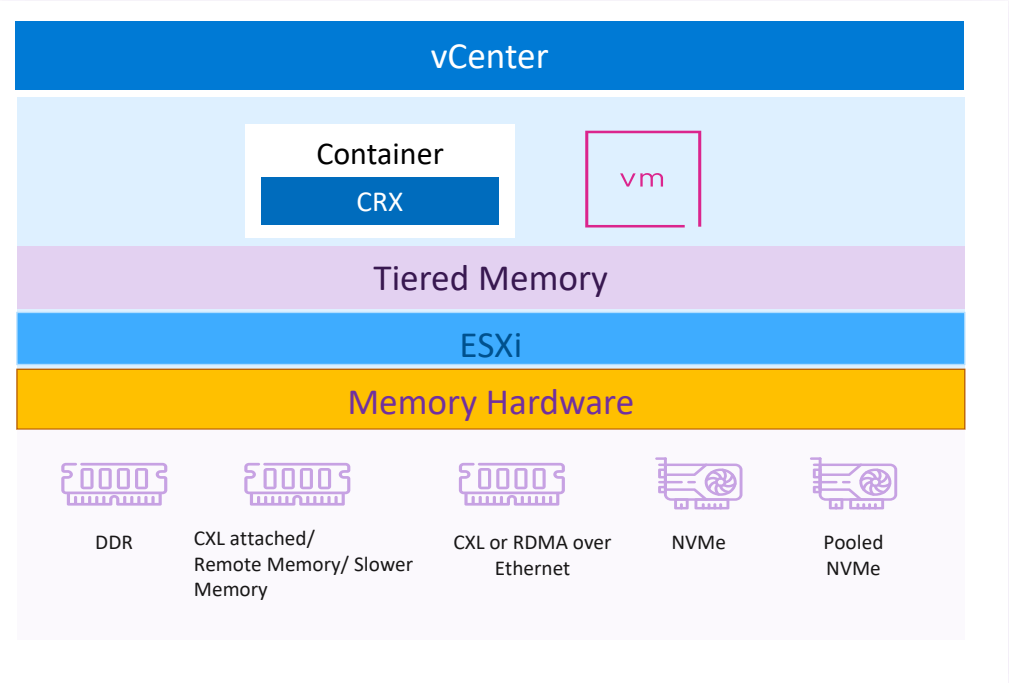

Benefits

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

Value over traditional approaches

- Virtualization
  - Independent underlying hardware changes
- 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/device heuristics fed to DRS
- Ensure Fairness across workloads
  - Consistent performance
- Min Configuration changes
  - No special tiering settings
- Minimum Performance Degradation
- Processor specific monitoring
  - vMMR monitors at both VM- and Host-levels

Key Use Cases emerging with CXL

Memory Expansion with NUMA-like latencies	Memory Tiering	Memory pooling across hosts on a cluster using memory appliances	Memory sharing	CXL switching and shared access (future)
-Increase capacity/scale -Flat (non-tiered) expansion -Consolidate server memory -Improve bandwidth -Improve core utilization	Lower TCO – combinations of lower cost memory with DRAM	Consolidate memory usage on a cluster	Utilize stranded memory on hosts	Disaggregation and Composability

Deployment Options

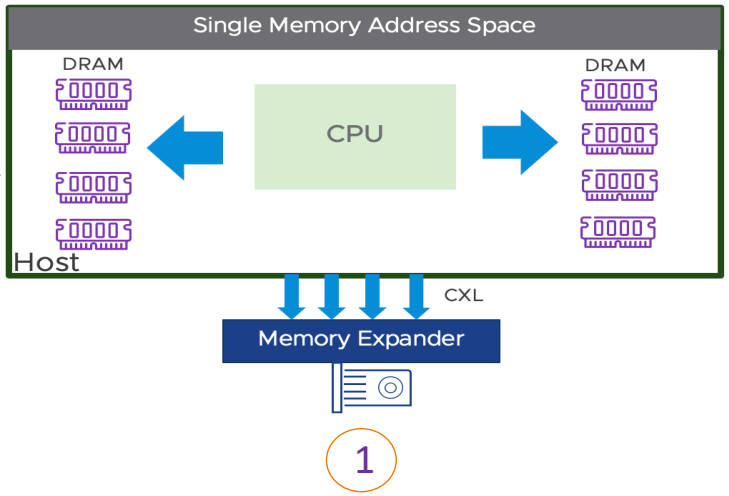

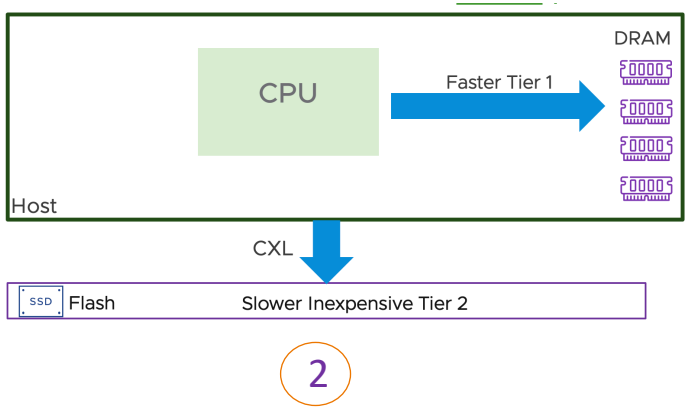

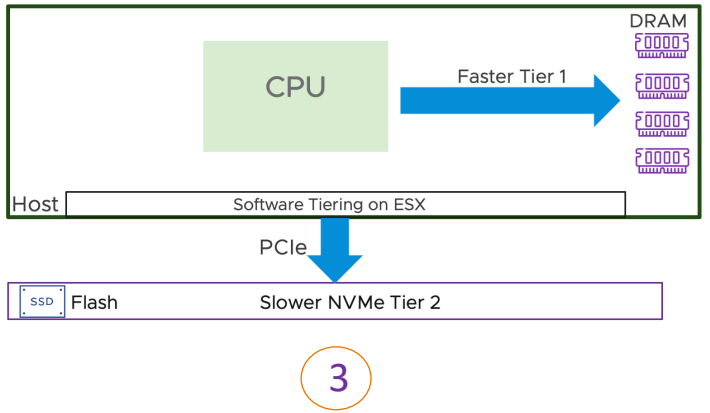

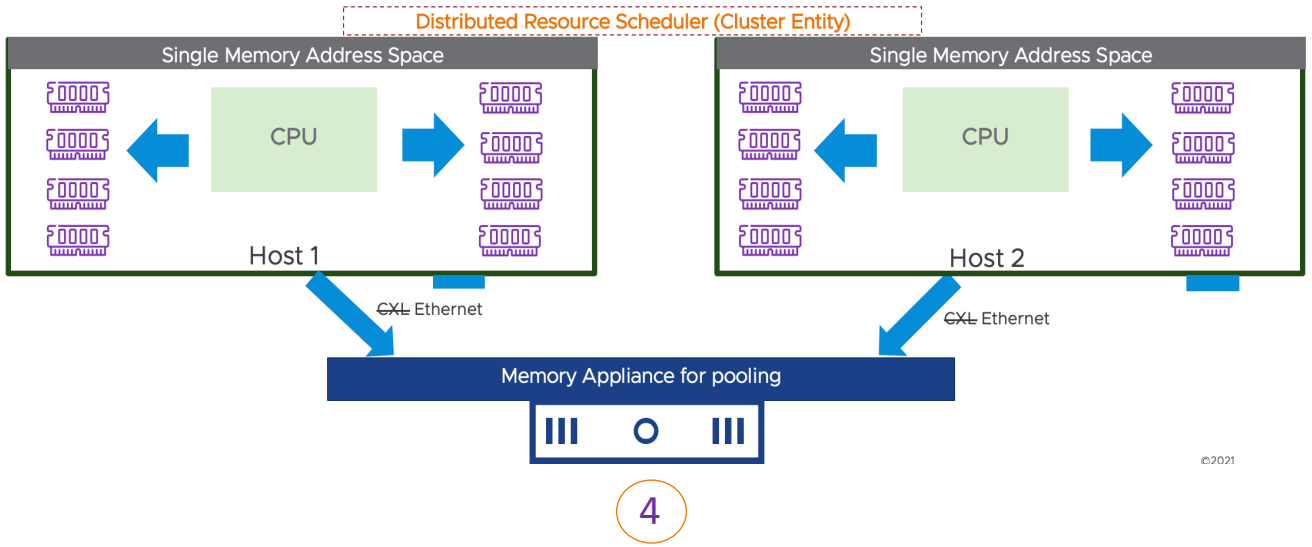

Example – Reducing Total Server Costs

2TB DRAM

Host configuration DRAM only		List Price[1]
CPU	2 x 6348 (28C @ 2.6GHz)	$6,144*
DRAM	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

1TB DRAM + 1TB on Tiered Memory

Host configuration DRAM + Low-cost Mem		List Price[1]
CPU	2 x 6348 (28C @ 2.6GHz)	$6,144*
DRAM	1024GB (32 x 32GB)	$29,184
Lower-cost Memory	1024GB	<3000
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		$46,682

Up to 36% Cost Reduction compared to DRAM-Only

*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

CapEx and OpEx improvements with CXL

Reduce CapEx

- Provide lower TCO by DRAM substitution
- Improve capacity and bandwidth
- Improve host CPU core utilization
- CPU Savings from page tracking/transfer offloads

Reduce OpEx

- Improve OpEx by reducing migration times
- Reduce failure probability
- Reduce host evacuation time for maintenance

Other future use-cases

- Memory tracking/scrubbing and proactive maintenance

Optimizing Oracle Workloads – Use Case

Oracle Workloads Using VMware Software Memory Tiering - Concept

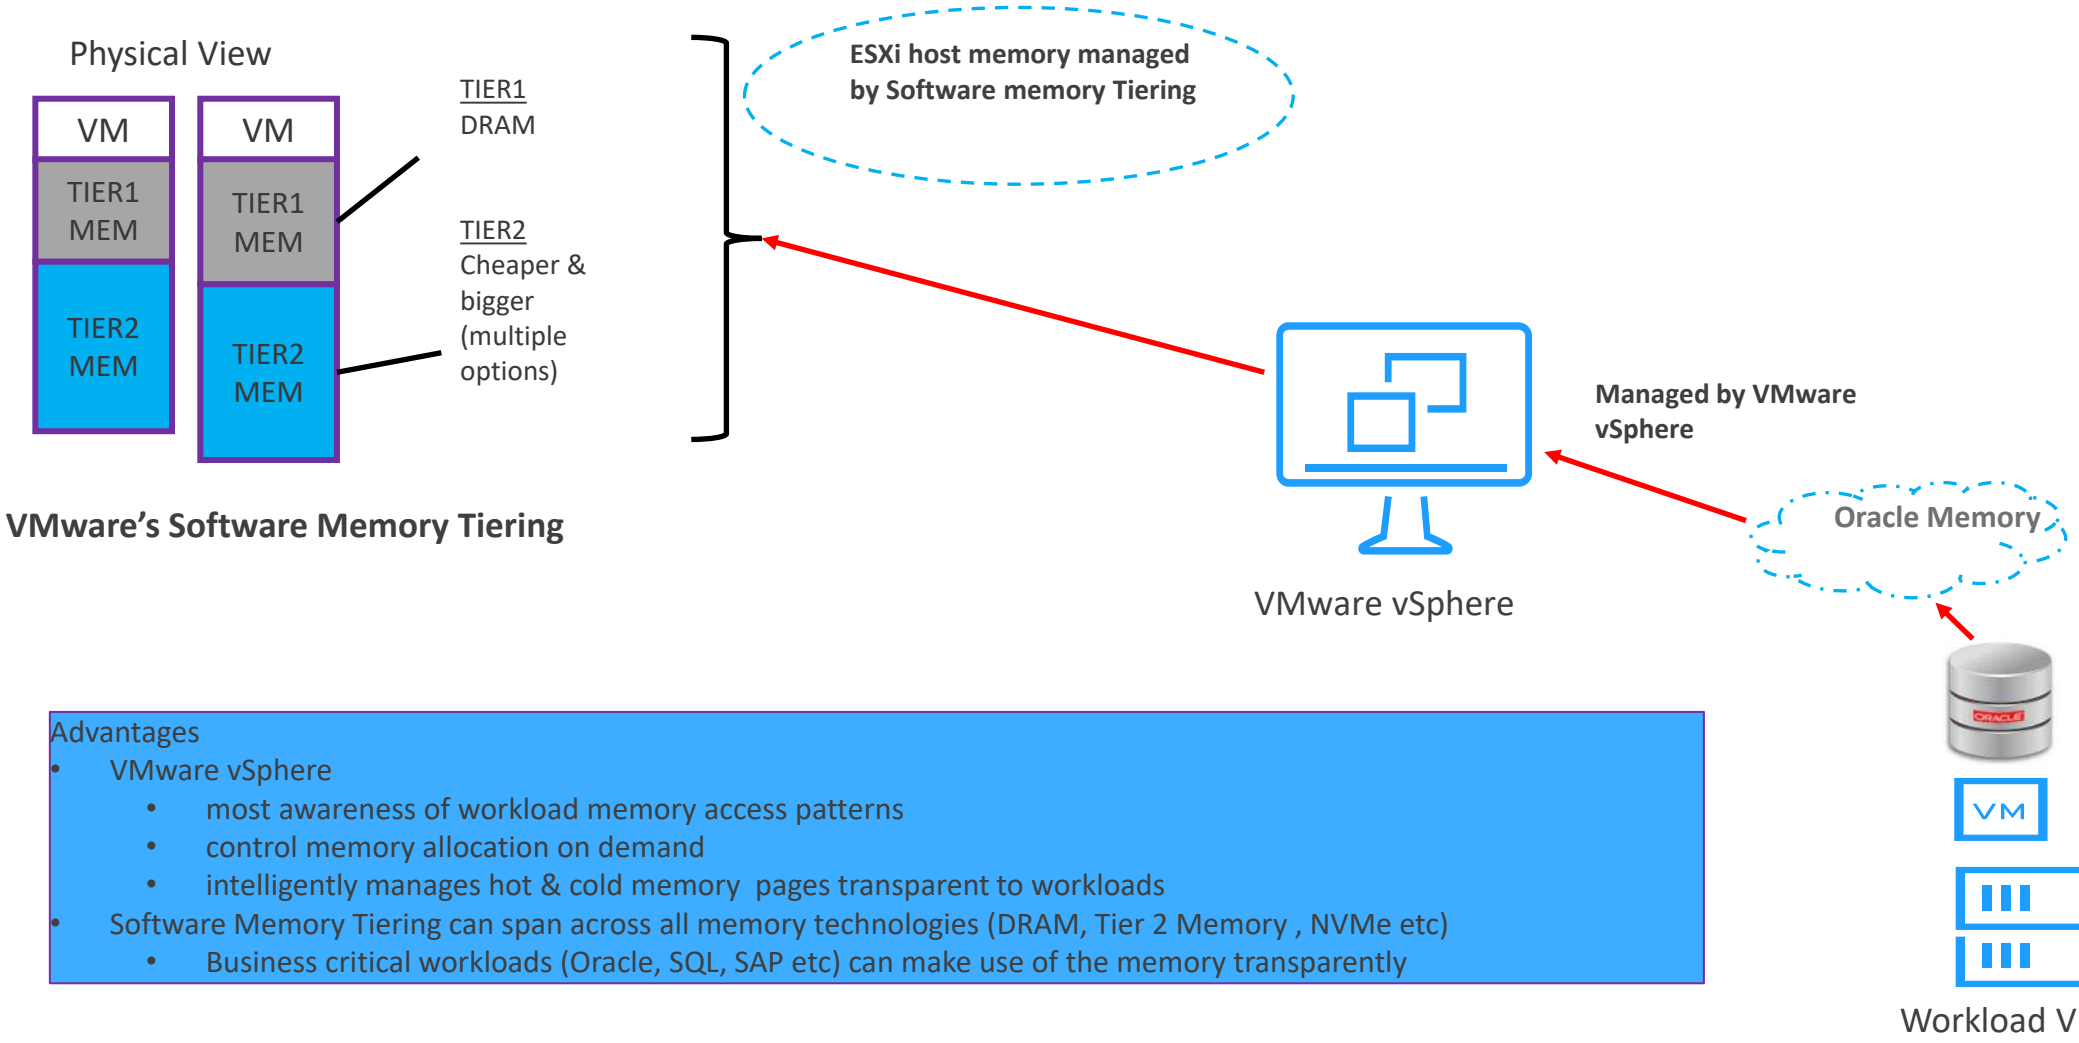

Advantages

- VMware vSphere
  - most awareness of workload memory access patterns
  - control memory allocation on demand
  - intelligently manages hot & cold memory pages transparent to workloads
- Software Memory Tiering can span across all memory technologies (DRAM, Tier 2 Memory , NVMe etc)
  - Business critical workloads (Oracle, SQL, SAP etc) can make use of the memory transparently

VMware Software Memory Tiering - ESXi Details

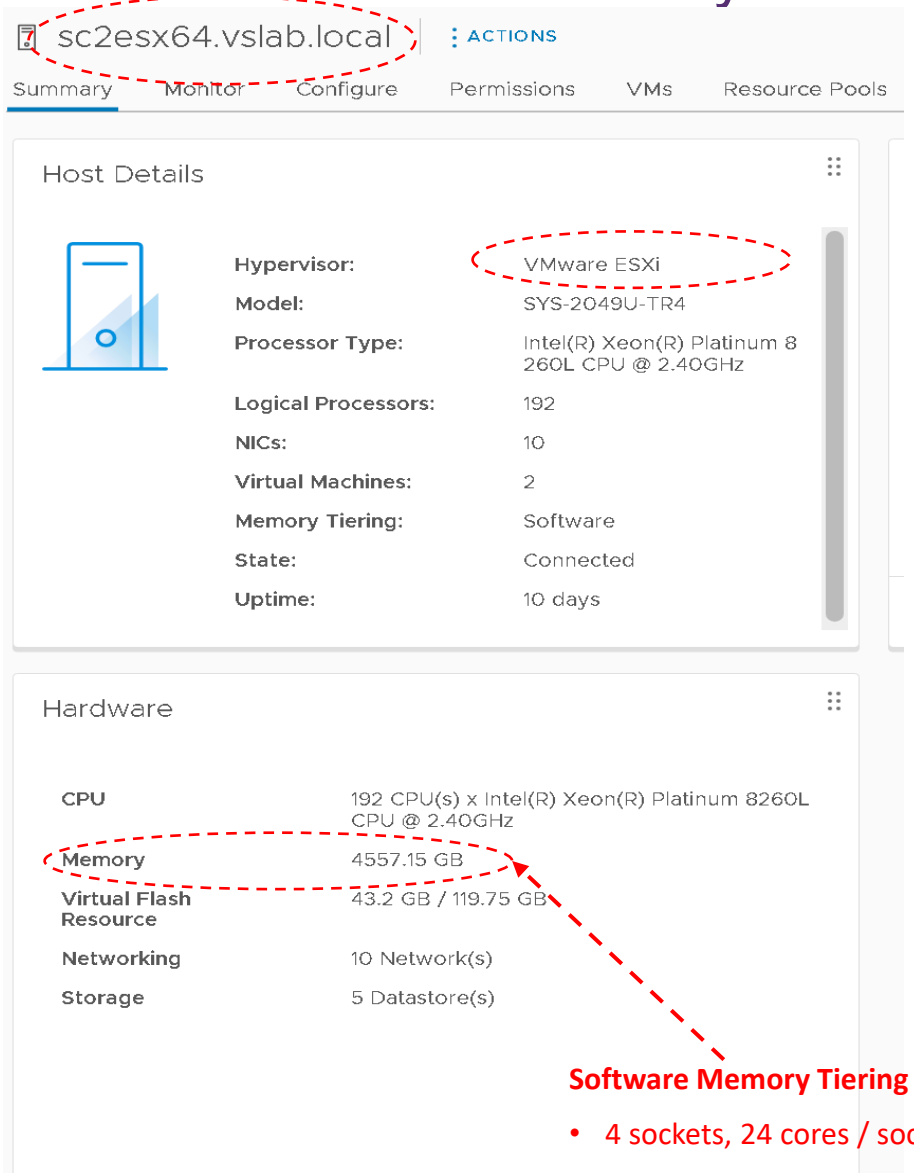

Software Memory Tiering Server

- 4 sockets, 24 cores / socket
- Total Memory 4.5 TB
  - 1.5 TB DRAM
  - 3 TB Tier 2 Memory

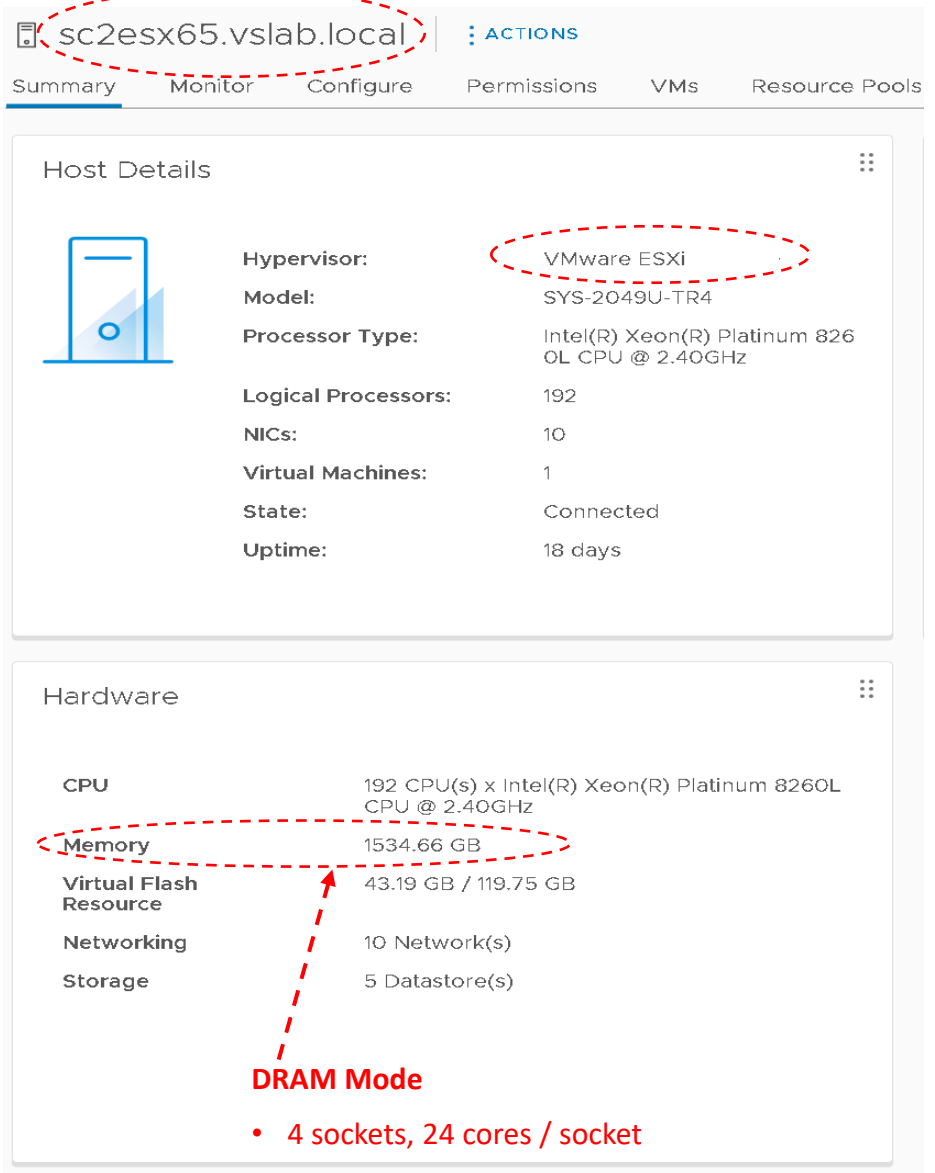

DRAM Mode

- 4 sockets, 24 cores / socket
- Total Memory 1.5 TB
  - 1.5 TB DRAM

VMware Software Memory Tiering - NUMA and Memory Details

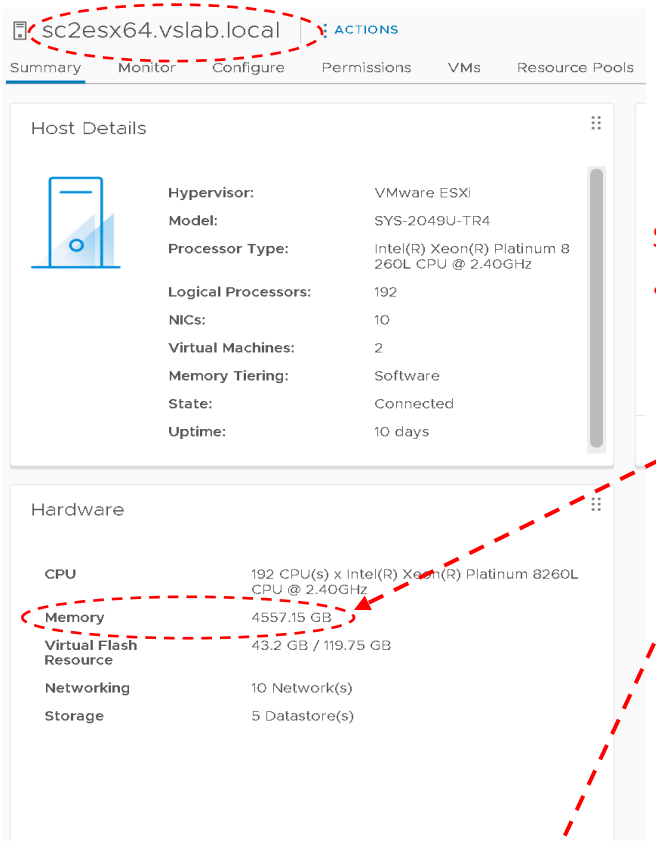

Software Memory Tiering Server

- **Total Memory 4.5 TB**
  - 1.5 TB DRAM
  - 3 TB Tier 2 Memory

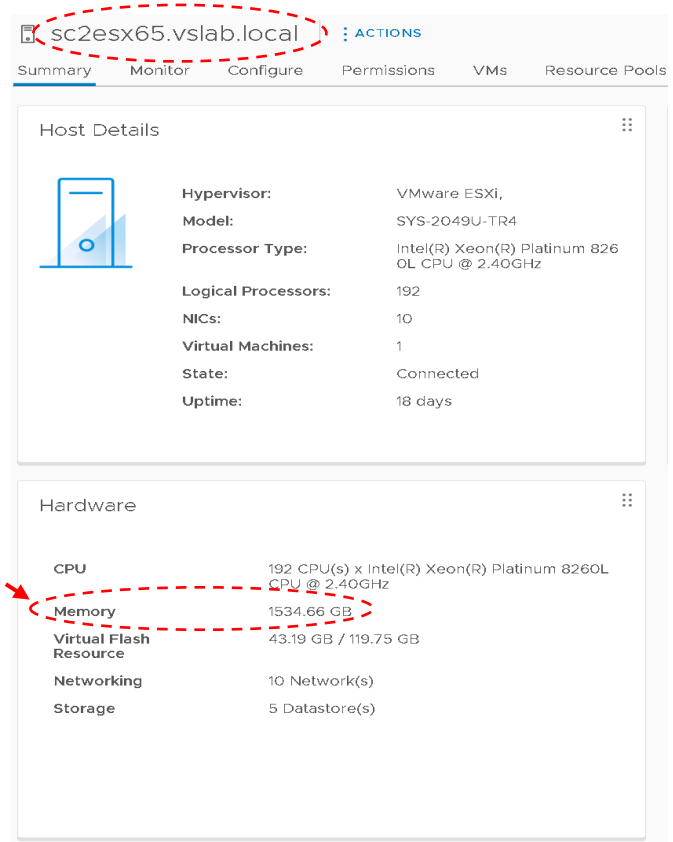

DRAM Mode Server

- **Total Memory 1.5 TB**
  - 1.5 TB DRAM

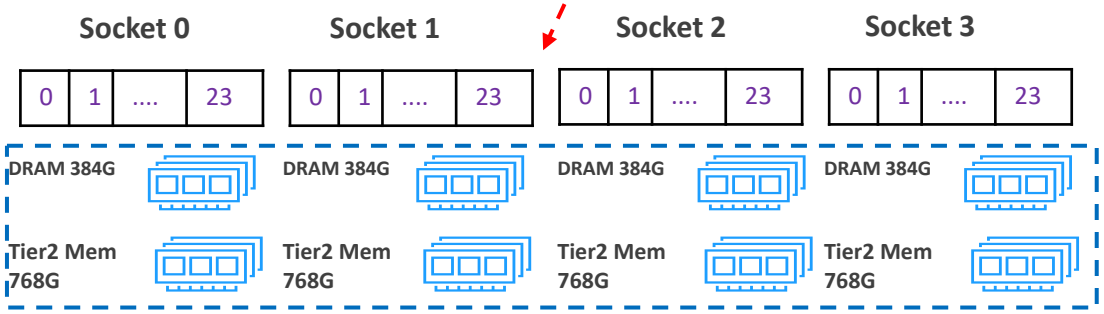

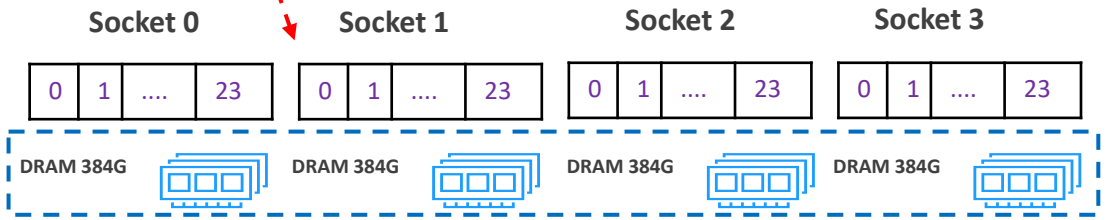

VMware Software Memory Tiering & DRAM only VM Details

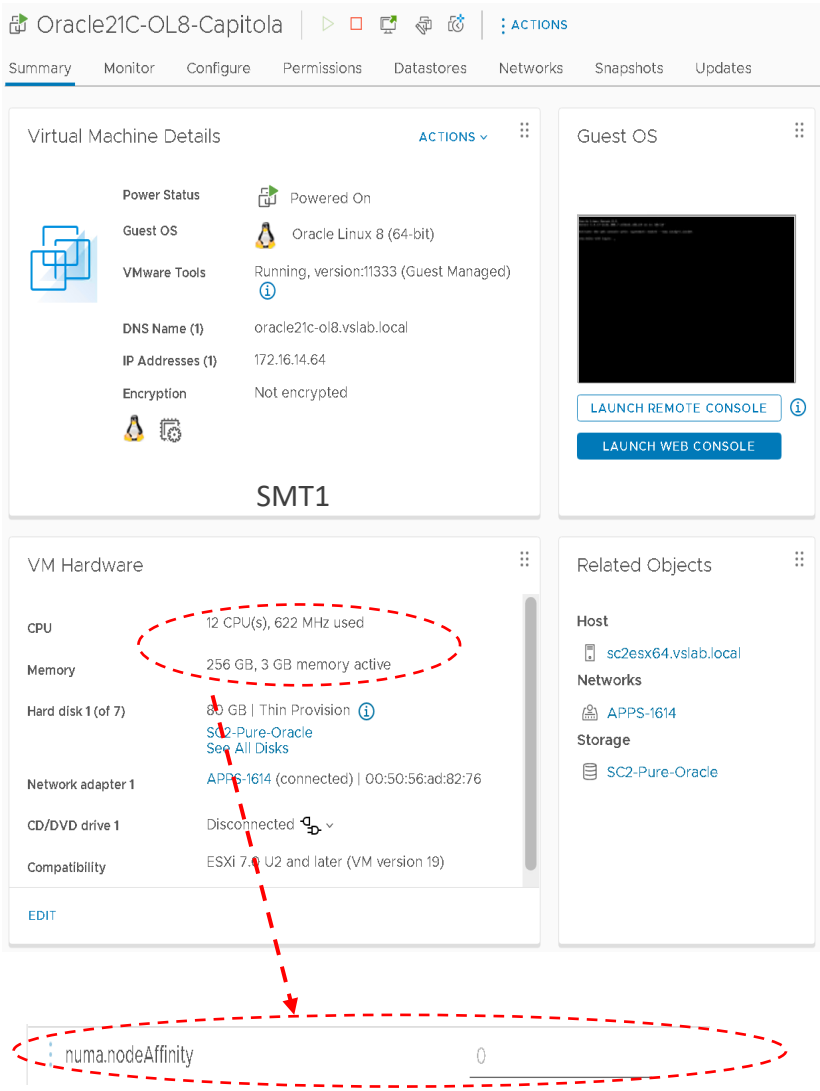

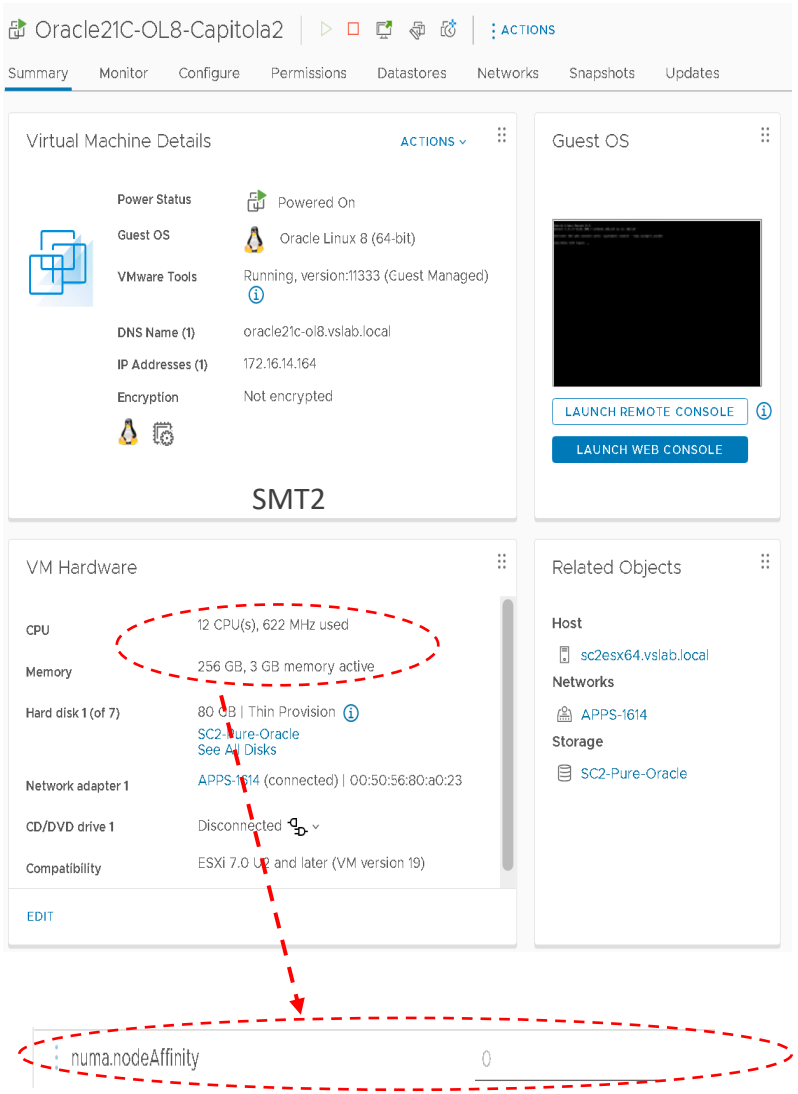

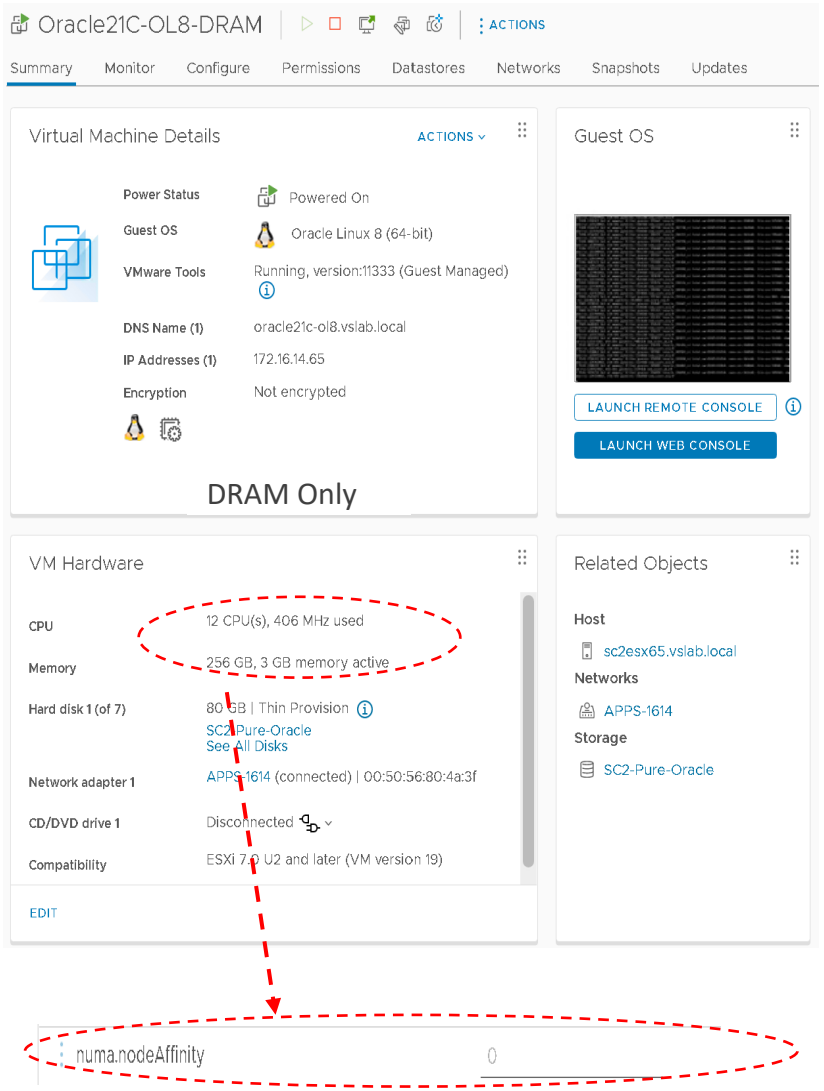

Software Memory Tiering VM's – SMT1 and SMT2

DRAM Only VM

Oracle Database – Details

- OEL 8.5 UEK ,Oracle 21.5
- Oracle Standalone DB , ASM & ASMLIB
- Oracle on VMware Best Practices Followed

Oracle Memory

- VM
  - vCPUS=12, RAM = 256G
- Oracle
  - SGA =192G, PGA=30G

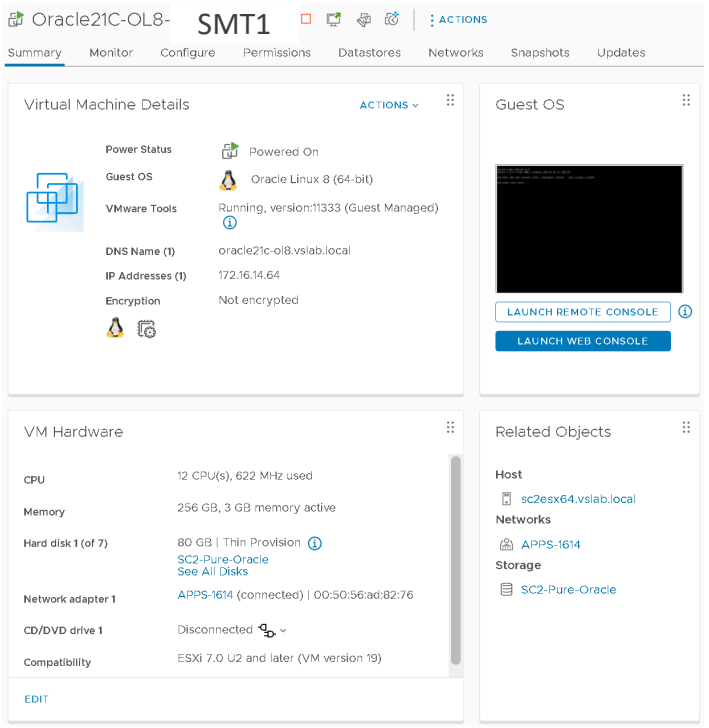

SW Memory Tiering VM – SMT1

Oracle Memory

- VM
  - vCPUS=12, RAM = 256G
- Oracle
  - SGA =192G, PGA=30G

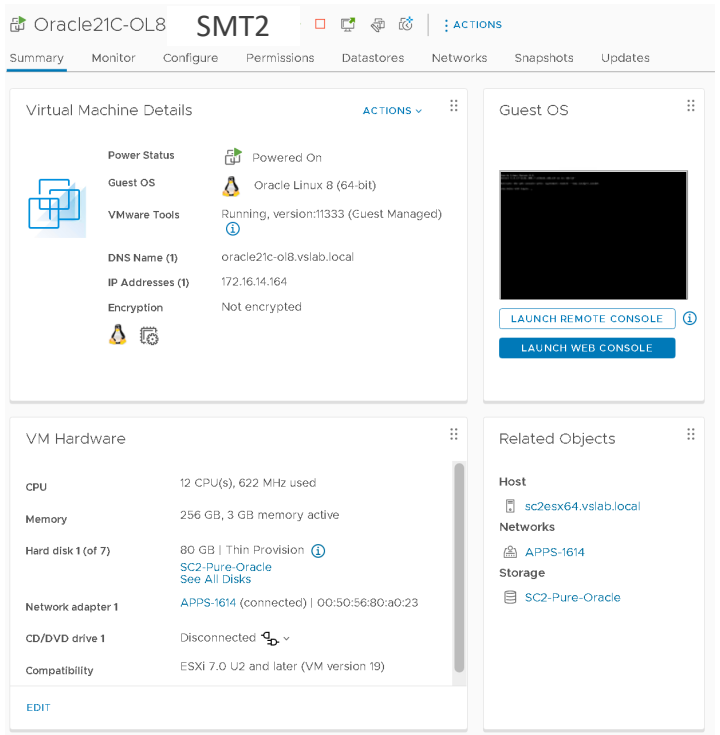

SW Memory Tiering VM – SMT2

Oracle Memory

- VM
  - vCPUS=12, RAM = 256G
- Oracle
  - SGA =192G, PGA=30G

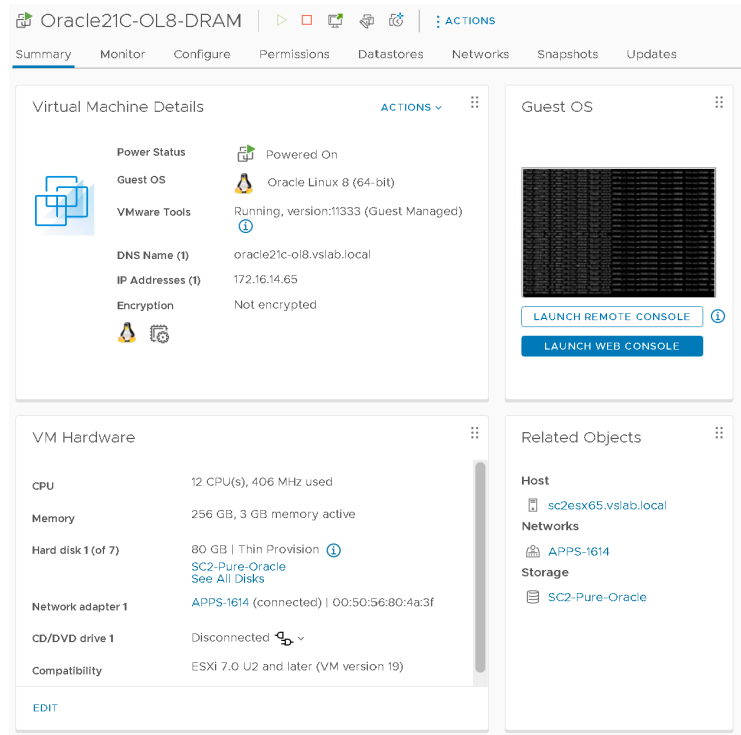

DRAM Mode VM1

Goal – Run '2 SW Memory Tiering' VM's on SMT Server on 1 NUMA node v/s '1 DRAM VM' on DRAM only Server on 1 NUMA node – Can we double our workload performance with lower TCO ?

Oracle Workload on SW Memory Tiering & DRAM Mode - Metrics

Executes (SQL) per second

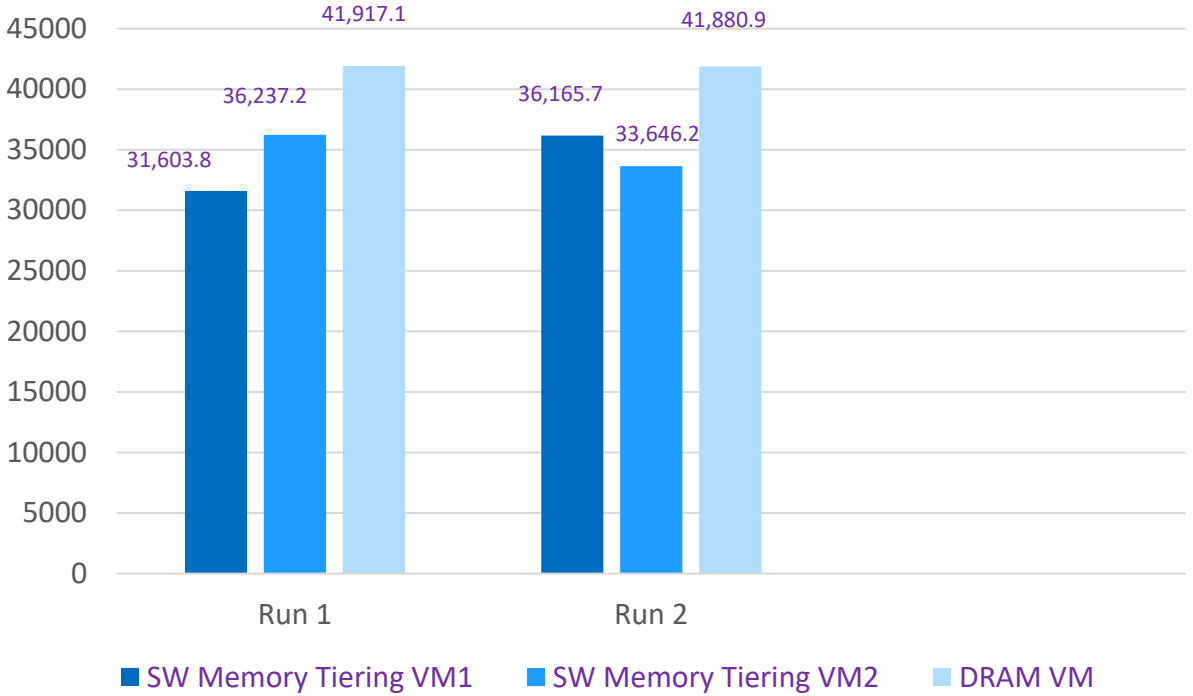

- Load Generator chosen as SLOB 2.5.4.0
  - UPDATE_PCT=0 - READ only test - performance comparison between SW Memory Tiering v/s DRAM Mode
  - RUN_TIME=1200 secs(20mins)

- Test Results
  - **Executes(SQL) / second**
  - Run 1
    - **SW Mem Tier VM1** – 33,603.8/sec , **SW Mem Tier VM2** – 36,237.2/sec
    - **DRAM Mode VM** - 41,917.1/sec

  - Run 2
    - **SW Mem Tier VM1** – 36,165.7/sec, **SW Mem Tier VM2** – 33,646.2/sec
    - **DRAM Mode VM** - 41,880.9/sec

Logical Reads (blocks) per second

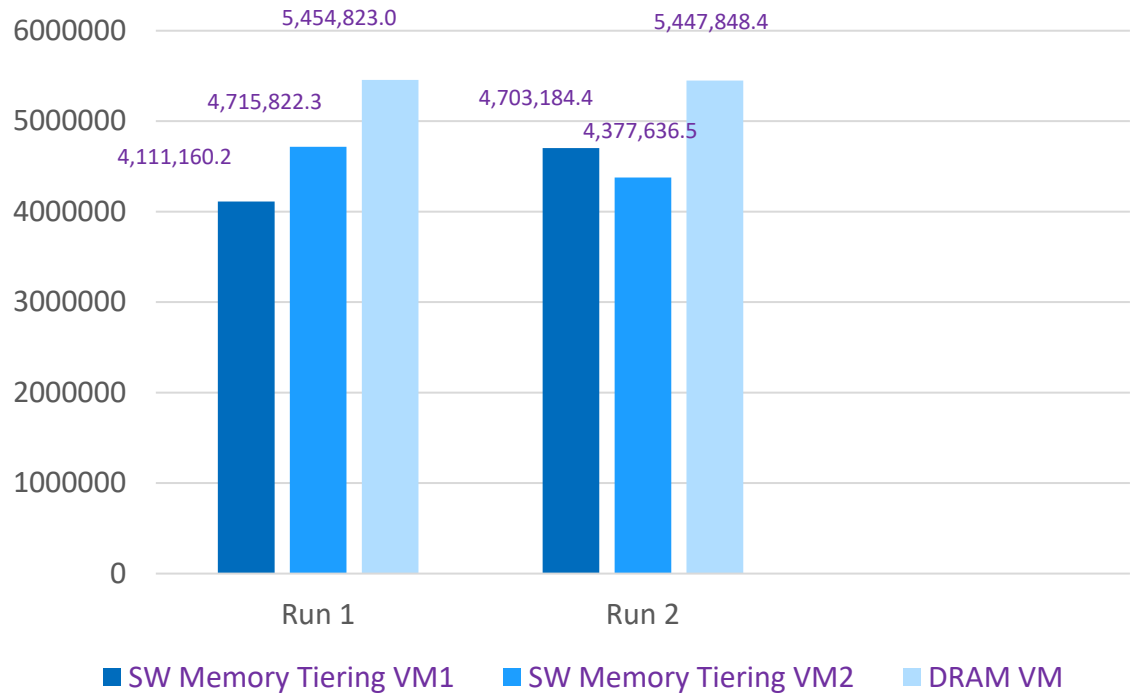

- Test Results
  - **Logical Reads (blocks) per second**
  - Run 1
    - **SW Mem Tier V**M1 – 4,111,160.2/sec
    - **SW Mem Tier VM2** – 4,715,822.3/sec
    - **DRAM Mode VM** - 5,454,823.0/sec

  - Run 2
    - **SW Mem Tier VM1** – 4,703,184.4/sec
    - **SW Mem Tier VM2** – 4,377,636.5/sec
    - **DRAM Mode VM** - 5,447,848.4/sec

SW Memory Tiering & DRAM Mode – OS & Oracle Metrics - Summary

Run	Metric	SW Mem Tier VM1	SW Mem Tier VM2	SW Mem Tier VM Aggregate	SW Mem Tier VM Average	DRAM VM	Difference (%)
Run 1	Executes(SQL) / second	33,603.80	36,237.20	**69,841.00**	34920.5	**41,917.10**	16.69
Run 2	Executes(SQL) / second	36,165.70	33,646.20	**69,811.90**	34,905.95	**41,880.90**	16.65
Run 1	Logical Reads (blocks) per second	4,111,160.20	4,715,822.30	**8,826,982.50**	4,413,491.25	**5,454,823.00**	19.09
Run 2	Logical Reads (blocks) per second	4,703,184.40	4,377,636.50	**9,080,820.90**	4,540,410.45	**5,447,848.40**	16.66

Run 1

SW Mem Tier VM1

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.16	12.50	75.7	2.3	0.0	21.2

SW Mem Tier VM2

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.15	413.04	86.8	2.5	0.0	9.8

DRAM Mem Tier VM1

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.09	141.57	88.0	2.5	0.0	8.8

Run 2

SW Mem Tier VM1

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.25	312.22	86.6	2.6	0.0	10.0

SW Mem Tier VM2

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.19	60.59	80.8	2.3	0.0	16.1

DRAM Mem Tier VM1

CPUs	Cores	Sockets	Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
12	12	12	0.22	234.71	87.9	2.5	0.0	8.8

More %Idle in case of SW Memory Tiering VM's – Better CPU Utilization

We were able to run '2 SW Memory Tiering' VM's on SMT Server on 1 NUMA node as compared to '1 DRAM VM' on DRAM Server on 1 NUMA node, even with the current VM size and DRAM capacity constraints , with increased workload performance with lower TCO as compared to DRAM only mode - VMware 'Software Memory Tiering' capability

Key Takeaways

Summary

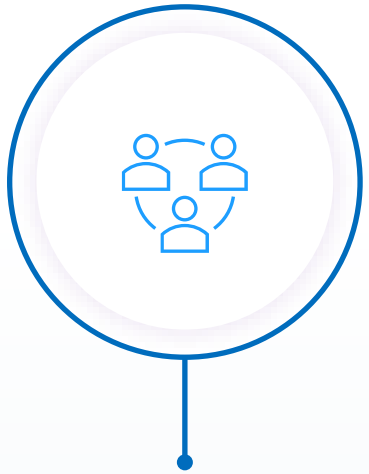

Intel and VMware have been collaborating to bring tiering to address cost, scale, and density challenges

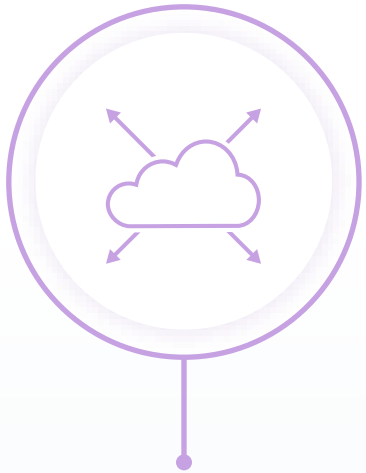

Business transformation is leading to larger datasets and real-time analytics that requires more performance and larger memory capacity

VMware performance results show software tiering is ready for the next phase of the big memory evolution. Applications benefit from being in memory. Mission critical applications like Oracle also benefit from such innovations

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

Oracle on VMware Collateral – One Stop Shop

- All Oracle on vSphere white papers including Oracle on VMware Hybrid Multi-Clouds (vSphere / vSAN / vVols / VMware Clouds) Best practices, Deployment guides, Workload characterization guide can all be found in the URL below

Oracle on VMware Collateral – One Stop Shop

https://via.vmw.com/Oracle-on-VMware

Sudhir Balasubramanian

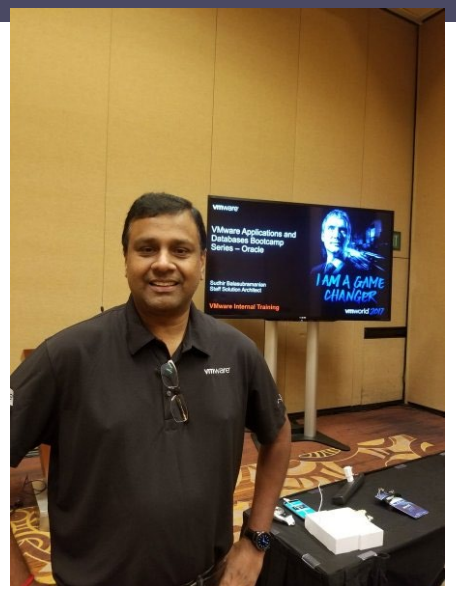

- **27 + years Oracle hands on experience** - Principal Oracle DBA / Architect, Oracle RAC/Data Guard Expert, Experienced in EMC SAN Technologies
- Currently - VMware [2012-] Senior Staff Solution Architect & Global Oracle Practice Lead
- Principal Oracle DBA/Oracle Architect (1995 – 2011) [Tata Consultancy Services (TCS), Sony Electronics, Newgen Results (Aspen) ,Teletech Corp, SAIC, Active Network, Sempra Energy Holdings]
- VMware VCA – Cloud ,VMware vBCA Specialist
- VMware vExpert & vExpert Application Modernization - https://vexpert.vmware.com/directory/1038
- Member of the Office of the Chief Technical Ambassador VMware (Alumni)
- Oracle ACE - https://apex.oracle.com/pls/apex/f?p=19297:3
- Leading Author of "Virtualizing Oracle Business Critical Databases on VMware SDDC"
- Recognized Speaker@ Oracle Open World, IOUG, VMworld, VMware Partner Exchange, EMC World, EMC Oracle Summit and Webinars
- Industry recognized expert in Oracle Virtualization technologies
- Blogs
  - http://vracdba.com/ | https://community.oracle.com/blogs/sudhirb
  - https://blogs.vmware.com/apps/author/sudhirbalasubramanian/
- Twitter : @vracdba [https://twitter.com/vRacDba]
- LinkedIn : https://www.linkedin.com/in/sudhirbalasubramanian/

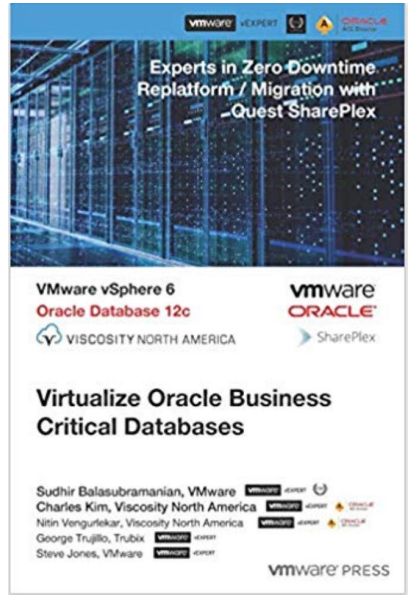

Arvind Jagannath

Sr. Product Line Manager for vSphere at VMware

Arvind Jagannath works in Product Management at VMware. With over 25 years of experience in the industry working on memory, networking, storage, embedded, and kernel development, he currently leads infrastructure and core platform enablement for vSphere, working across the VMware ecosystem of server/CPU, IO, and storage partners. Arvind most recently drove platform product management at Cohesity and NetApp. Arvind holds an MBA from the University of Chicago, Booth school of Business and a Bachelors in Computer Science and Engineering from India.

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Please take a moment to rate this session.

Your feedback is important to us.

