STORAGE DEVELOPER CONFERENCE
SDC 22
Fremont, CA
September 12-15, 2022
BY Developers FOR Developers
A SNIA Event

Addressing Capacity and Cost Challenges for Oracle Workloads Using VMware Software Memory Tiering

Sudhir Balasubramanian - Senior Staff Solution Architect - Oracle, VMware

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

Value Proposition and Vision

VMware Competencies

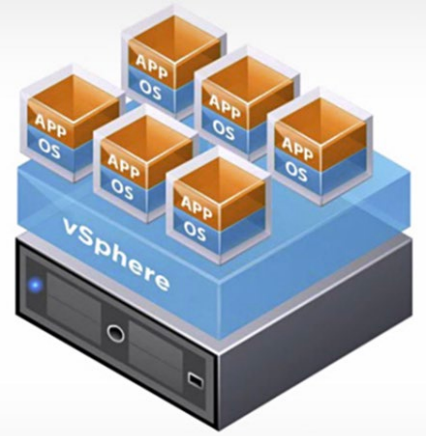

Virtualization ideal for transparent tiering

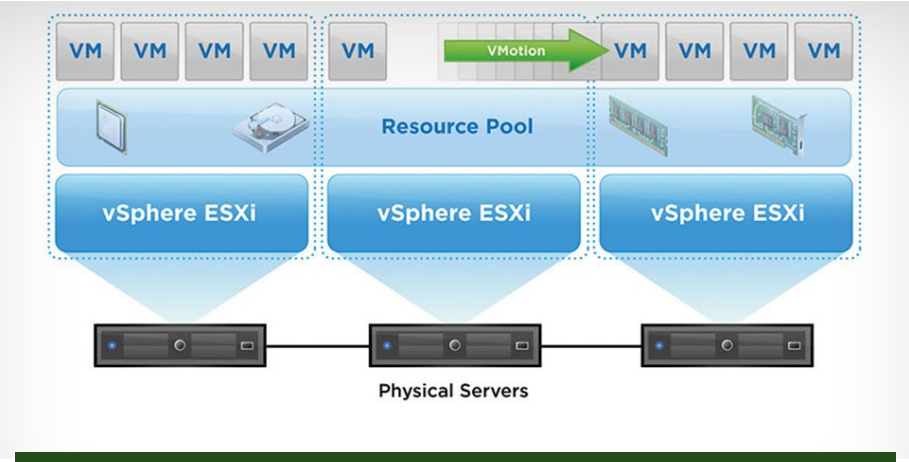

Cluster-wide DRS helps load balance and mitigate risks

Strong Ecosystem of partners

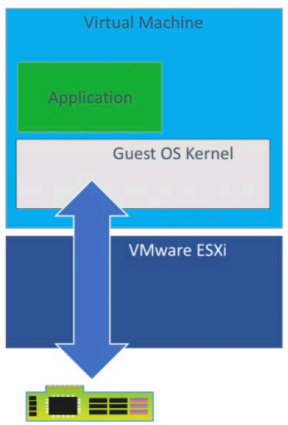

Passthrough devices

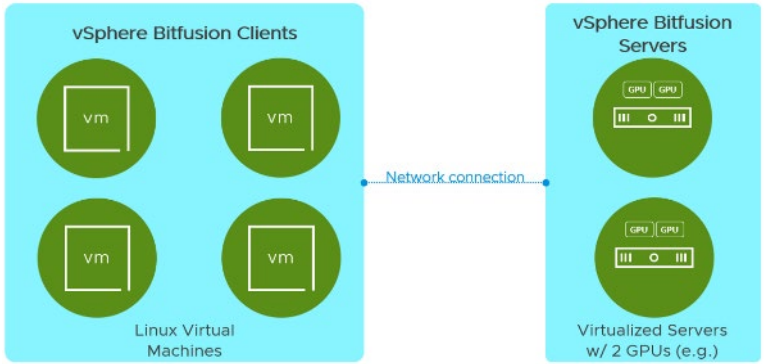

GPUs, sharing and Assignable hardware

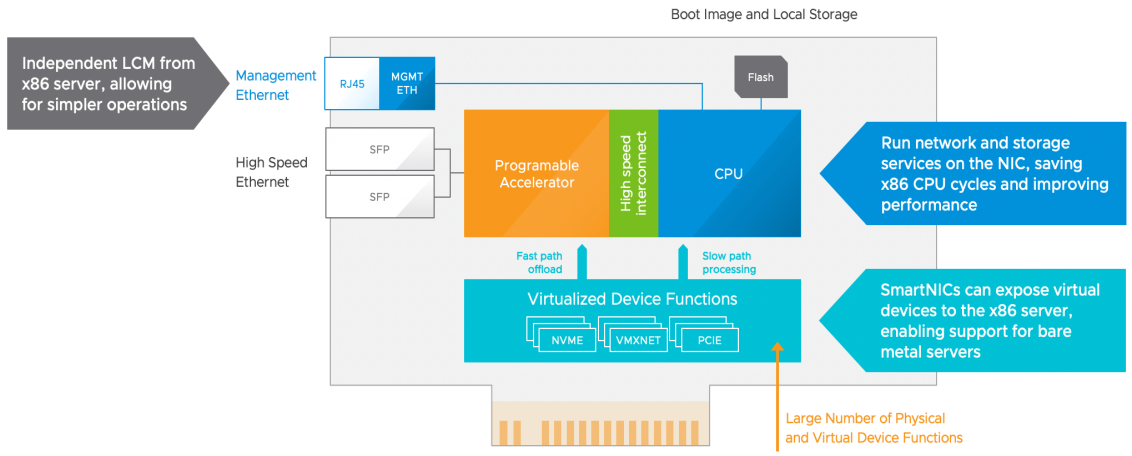

SmartNICs and Accelerators

Digital Transformation of Businesses

Explosive growth in data

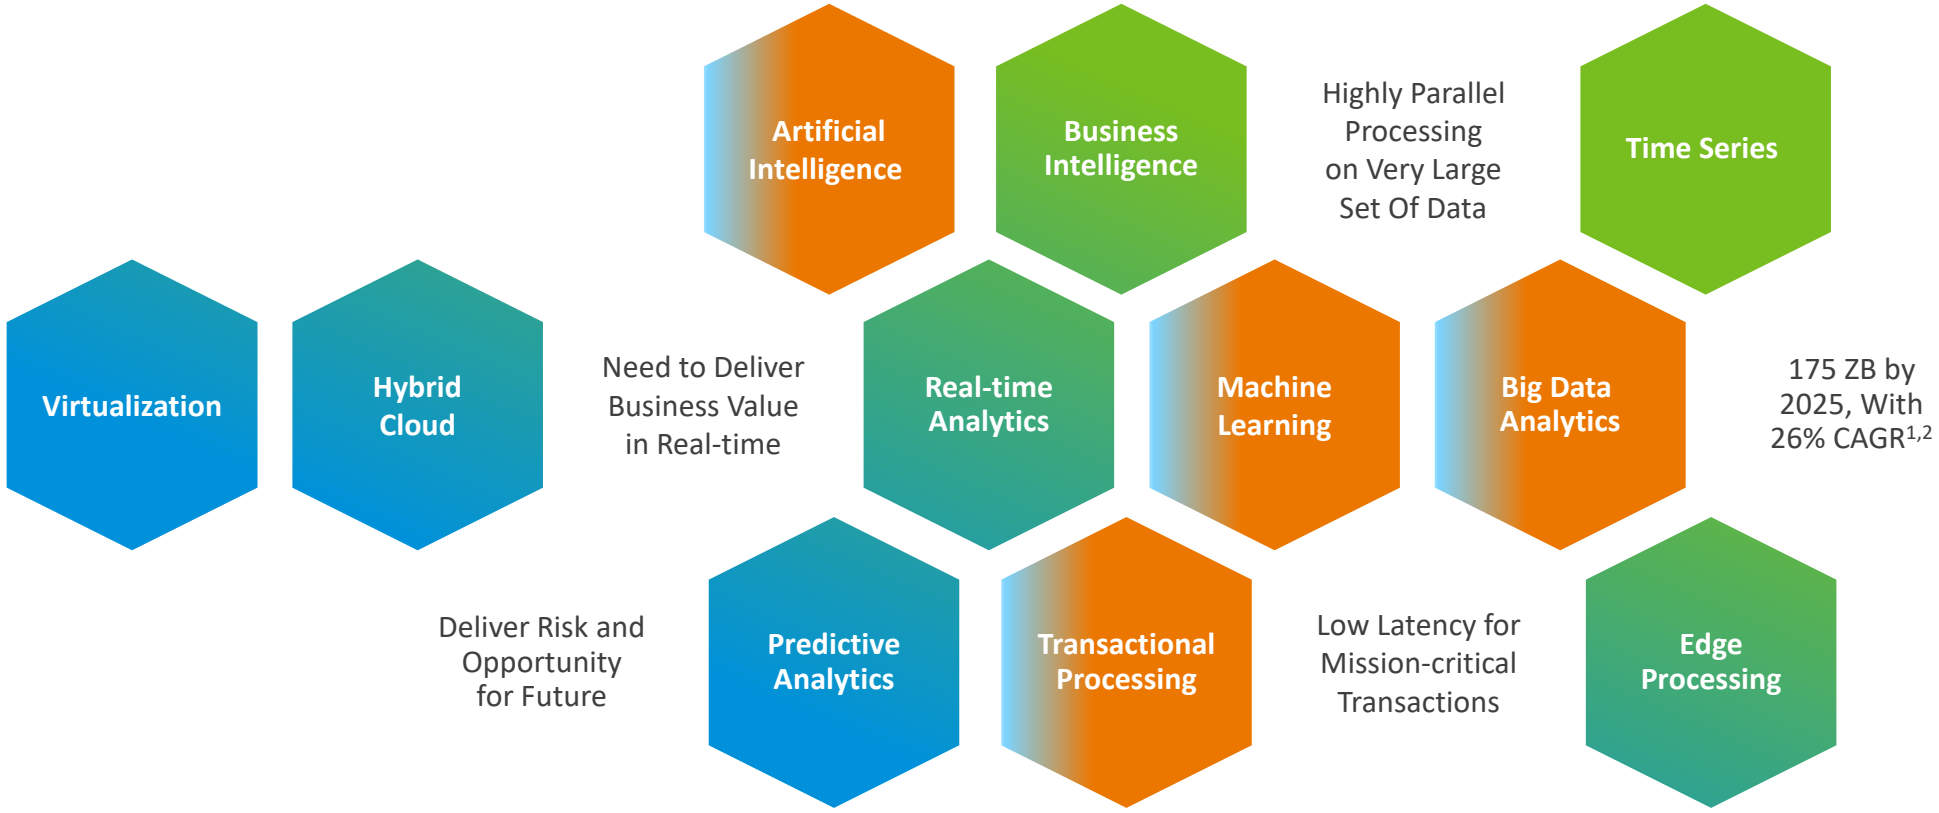

By 2025, IDC predicts **30%** of global data will be real time!

[1] NetworkWorld. "IDC: Expect 175 zettabytes of data worldwide by 2025." December 2018. networkworld.com/article/3325397/idc-expect-175-zettabytes-of-data-worldwide-by-2025.html.
[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.

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

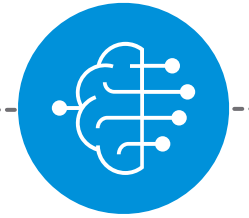

Digital Transformation

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

VMware's Big Memory Vision starts with Software Tiering

Software Tiering on ESX

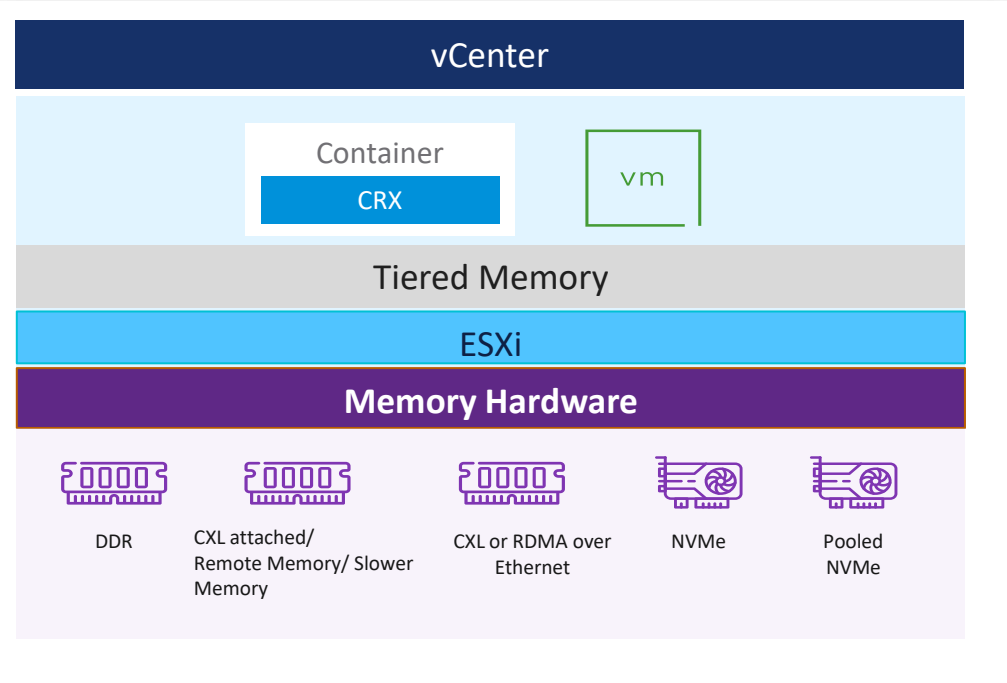

- Higher Density, more capacity
- Lower TCO
- Minimum Performance Degradation
- Transparent – Single system memory address
  - No Guest or Application changes
    - Run any Operating System
  - ESX internally handles page placement
- DRS and vMotion to mitigate risks
  - Tiering heuristics fed to DRS
- Ensures fairness across workloads
  - Consistent performance
- Zero Configuration changes
  - No special tiering settings
- Processor specific monitoring
  - vMMR monitors at both VM- and Host-levels

Software Tiering: What is it?

Host Local memory tiering

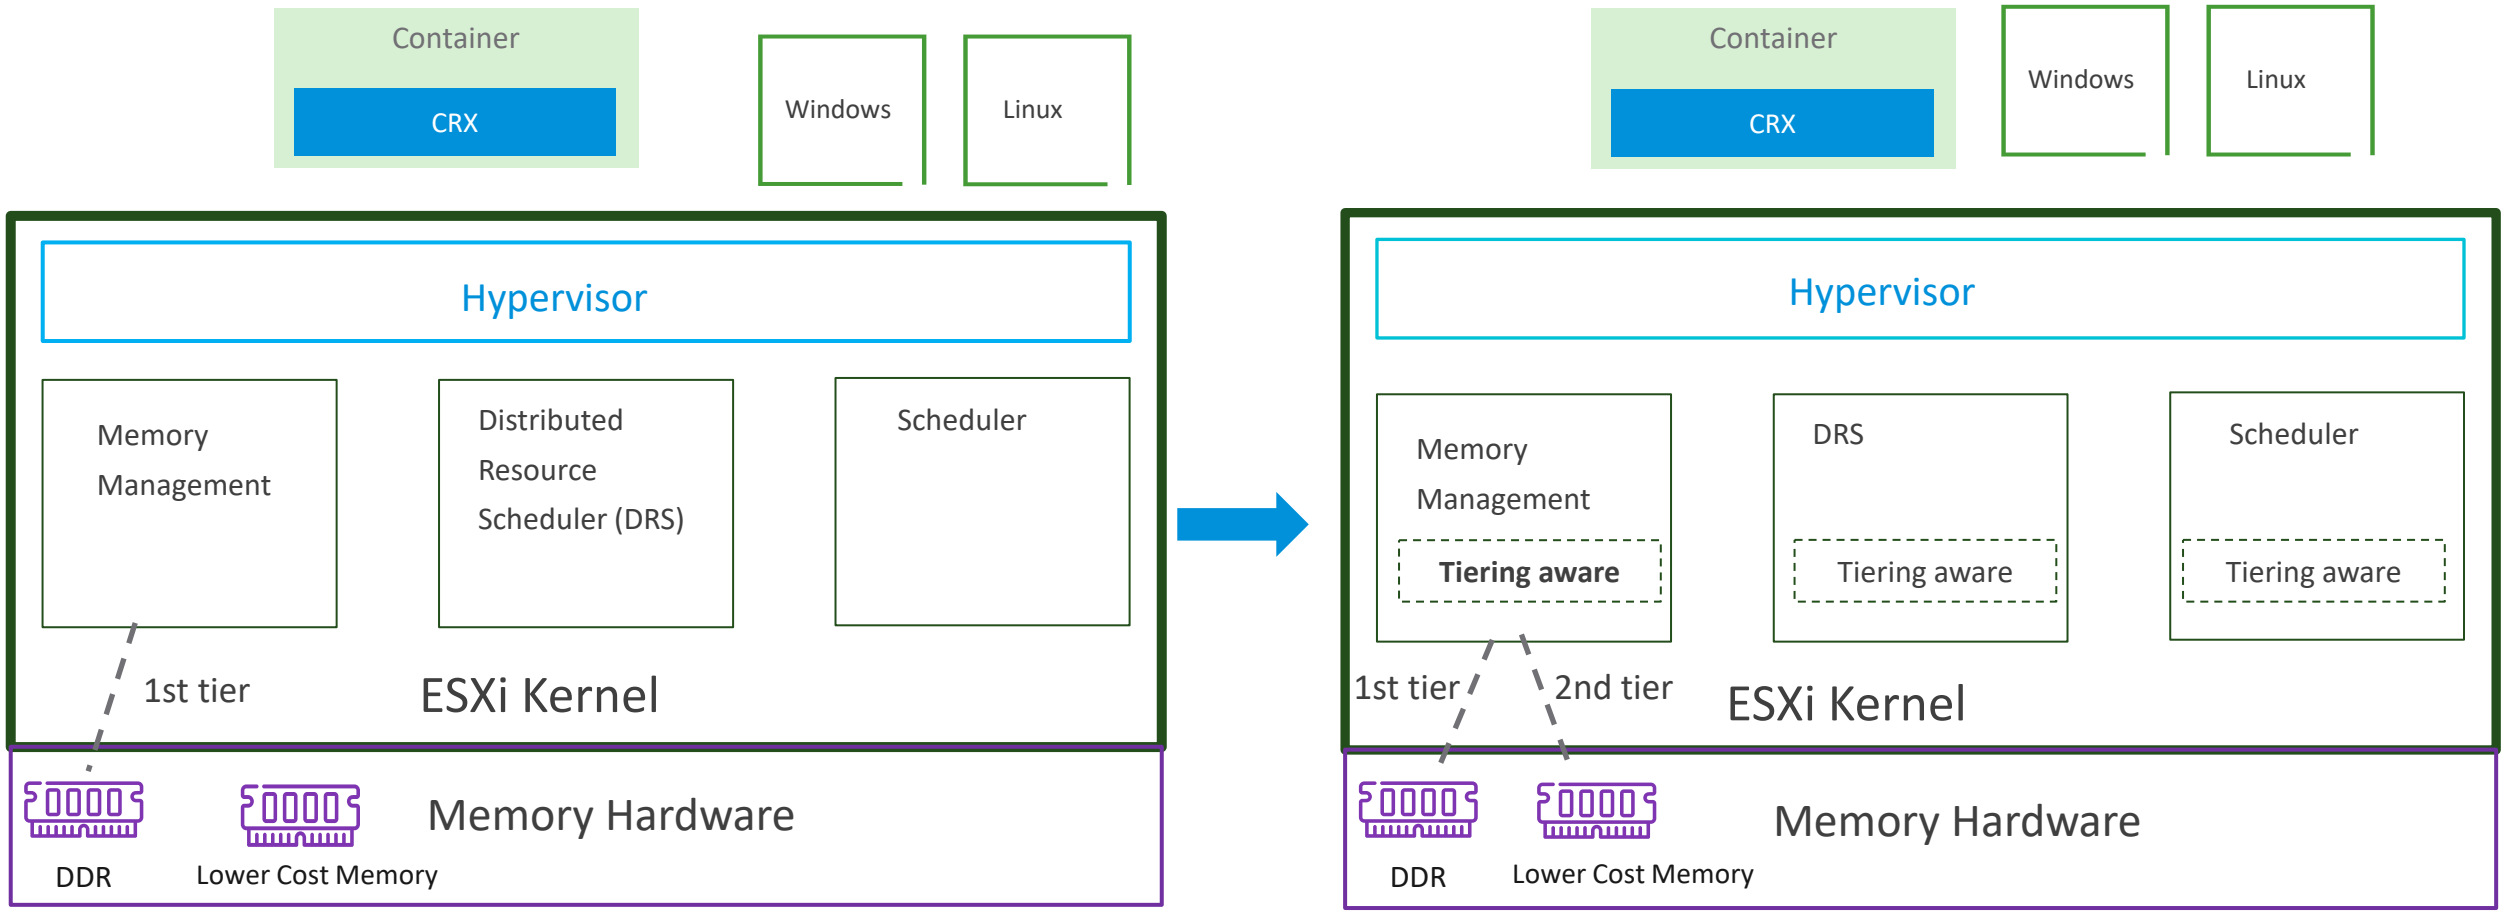

Software Tiering: Relevance for the future?

Disaggregation, Sharing and Pooling

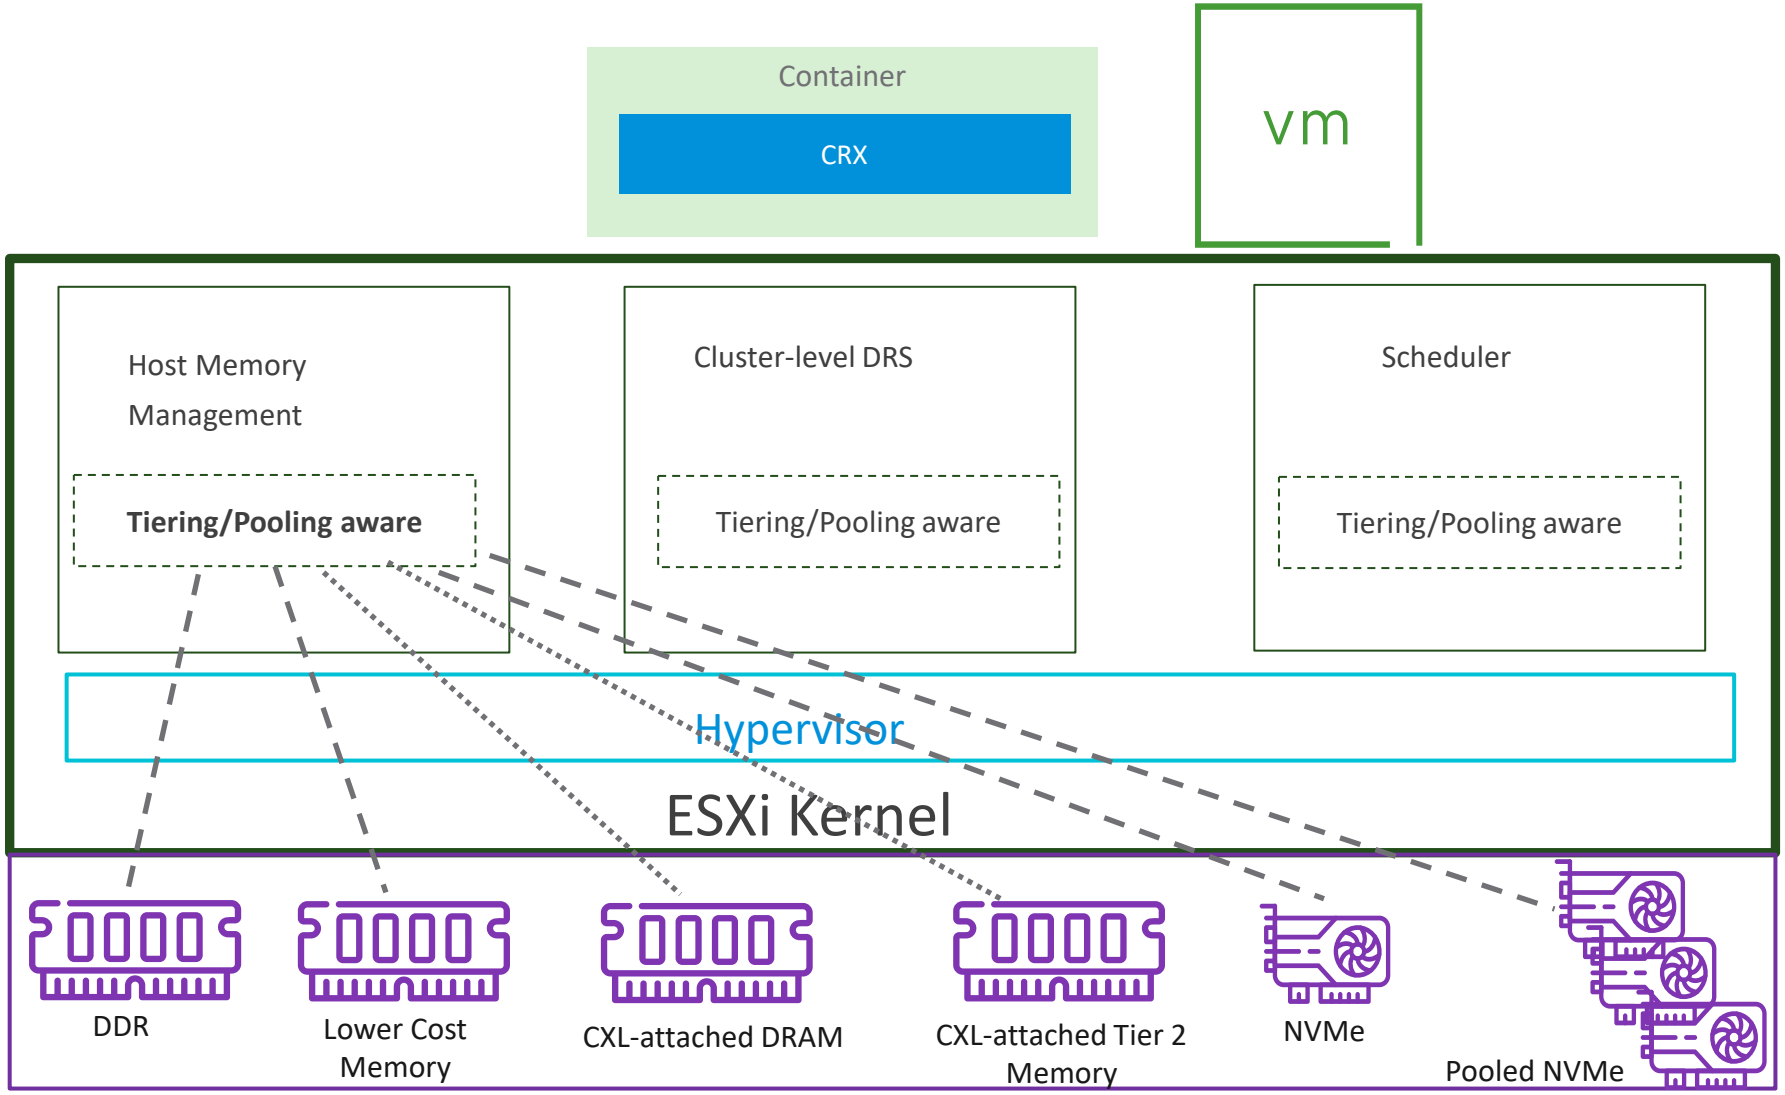

Various Tiering Approaches

Future

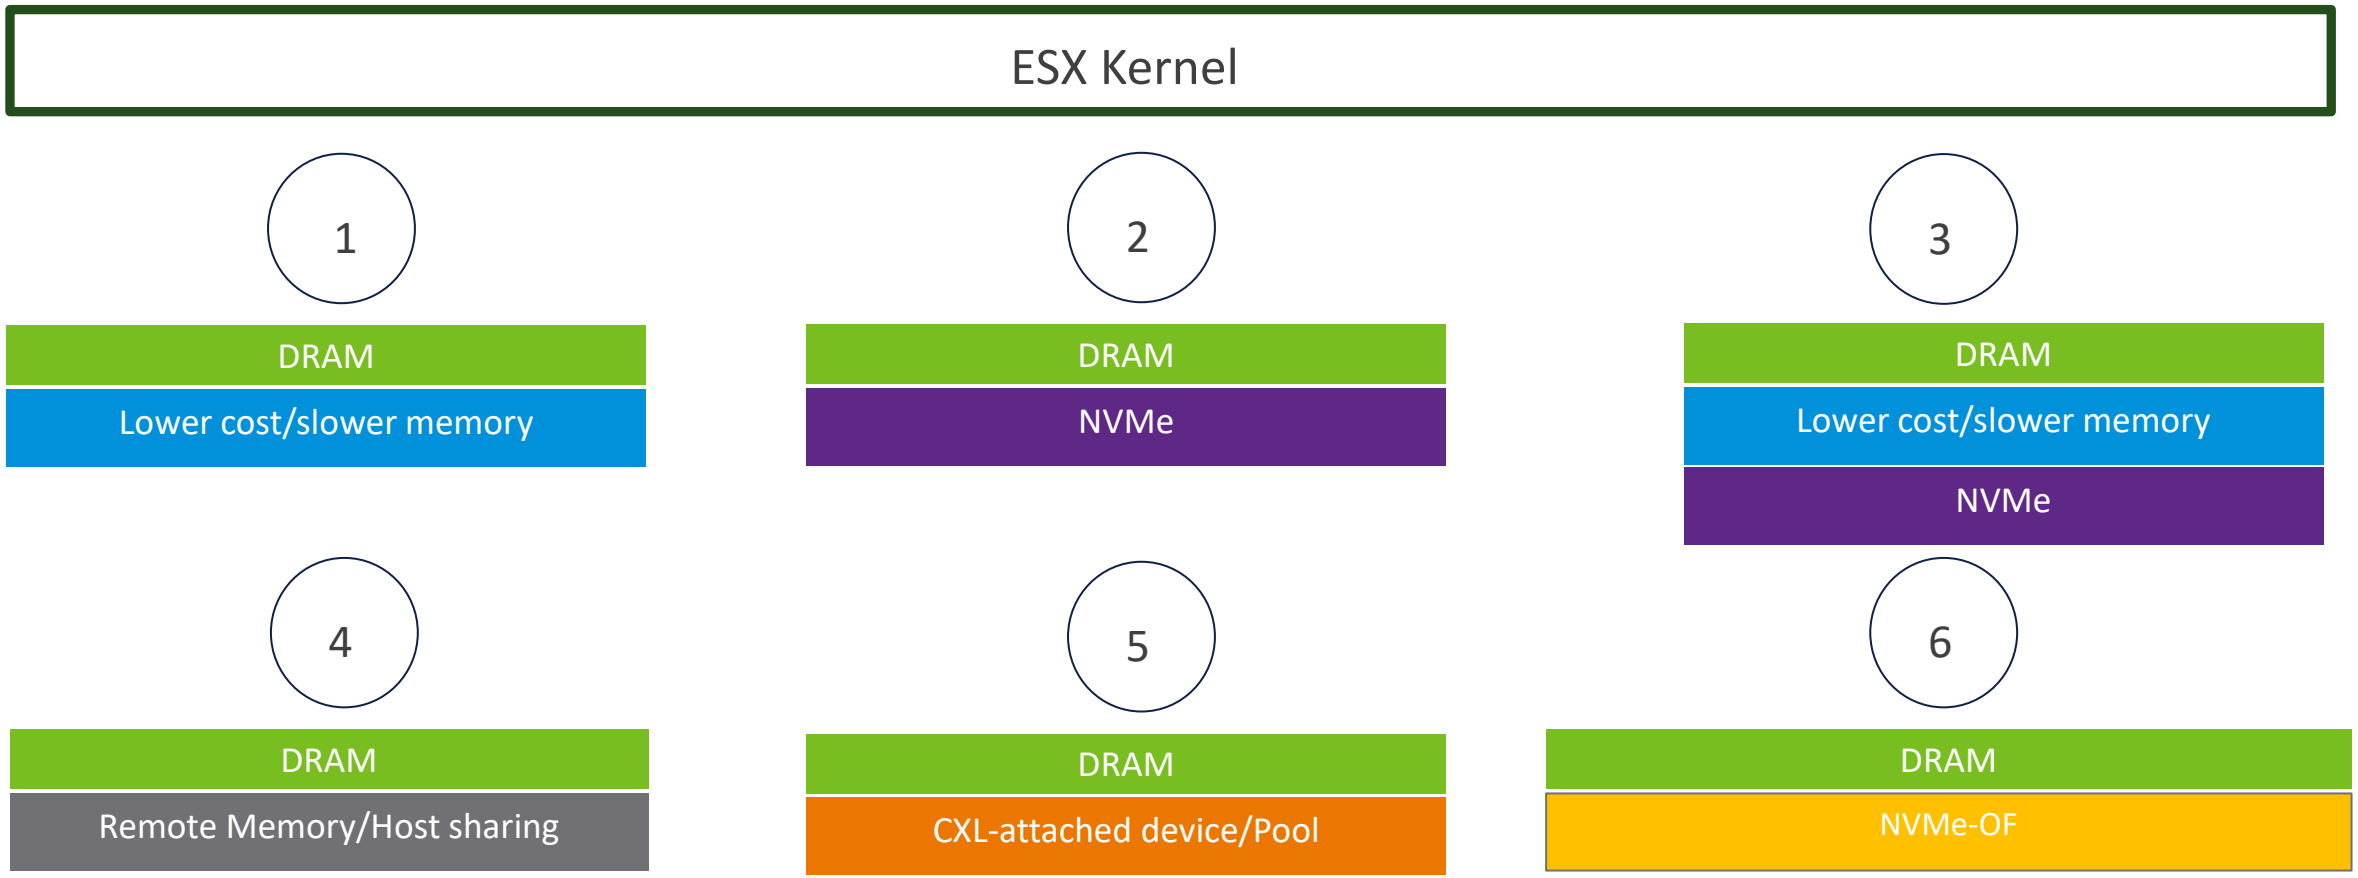

Software Tiering with CXL 2.0

Future

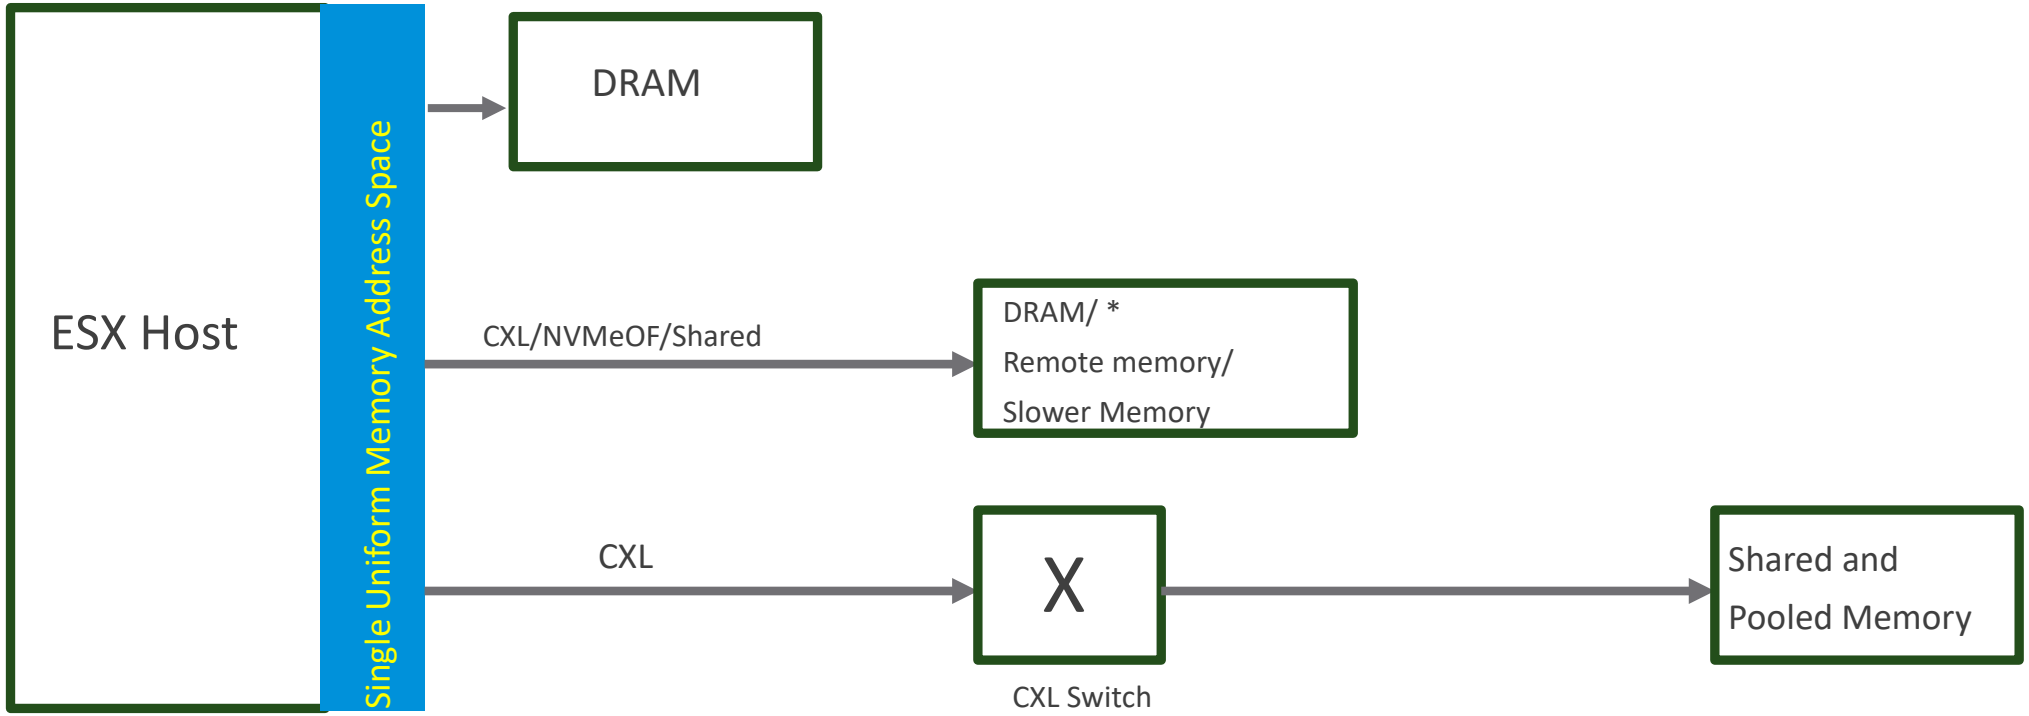

**Prototyping with CXL1.1*

Software Memory Tiering in vSphere

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

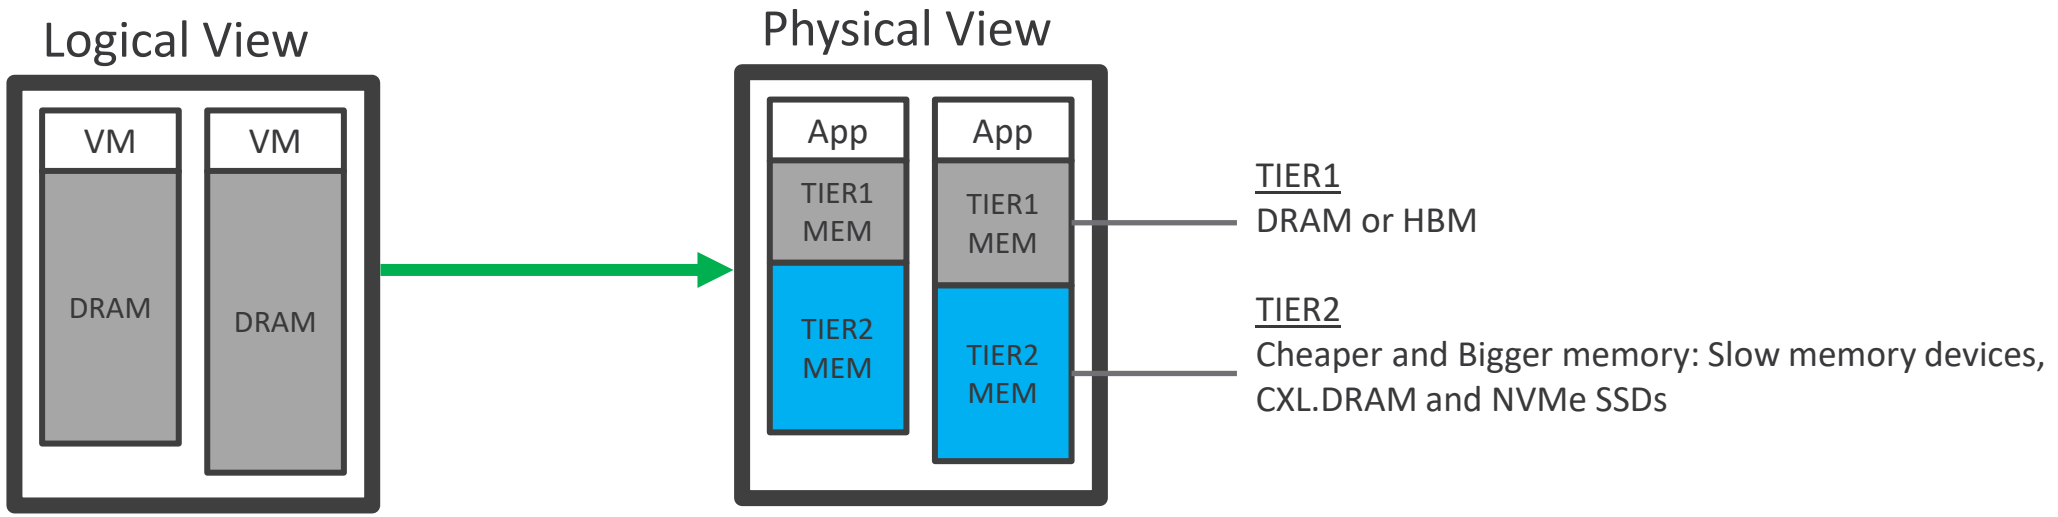

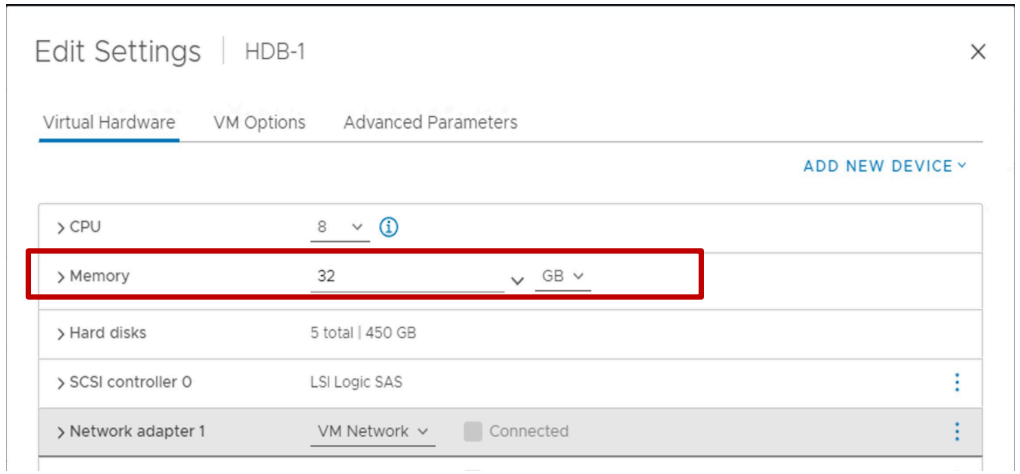

- ✓ vSphere decides what tiers to use and when
- ✓ Memory available to the host is sum of all memory tiers

How It All Fits Together in a vSphere Cluster?

Managed part of end-to-end vSphere workflow!

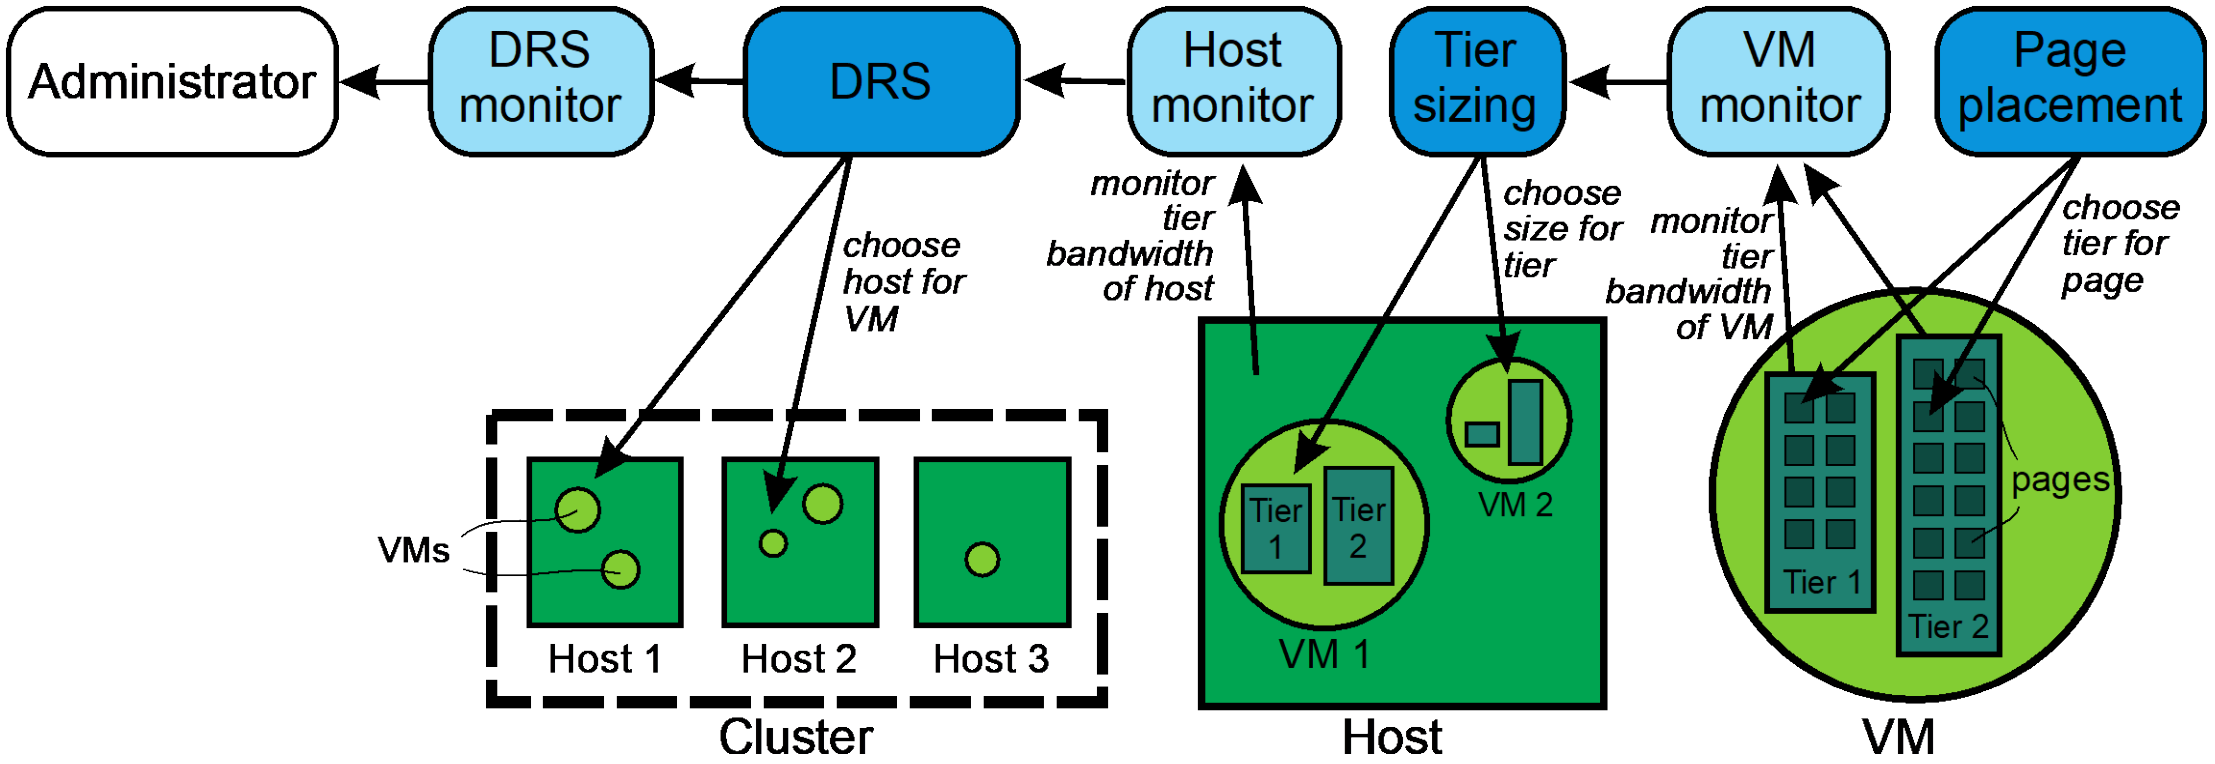

Optimizing Oracle Workloads –Use Case

Oracle Workloads Using VMware Software Memory Tiering - Concept

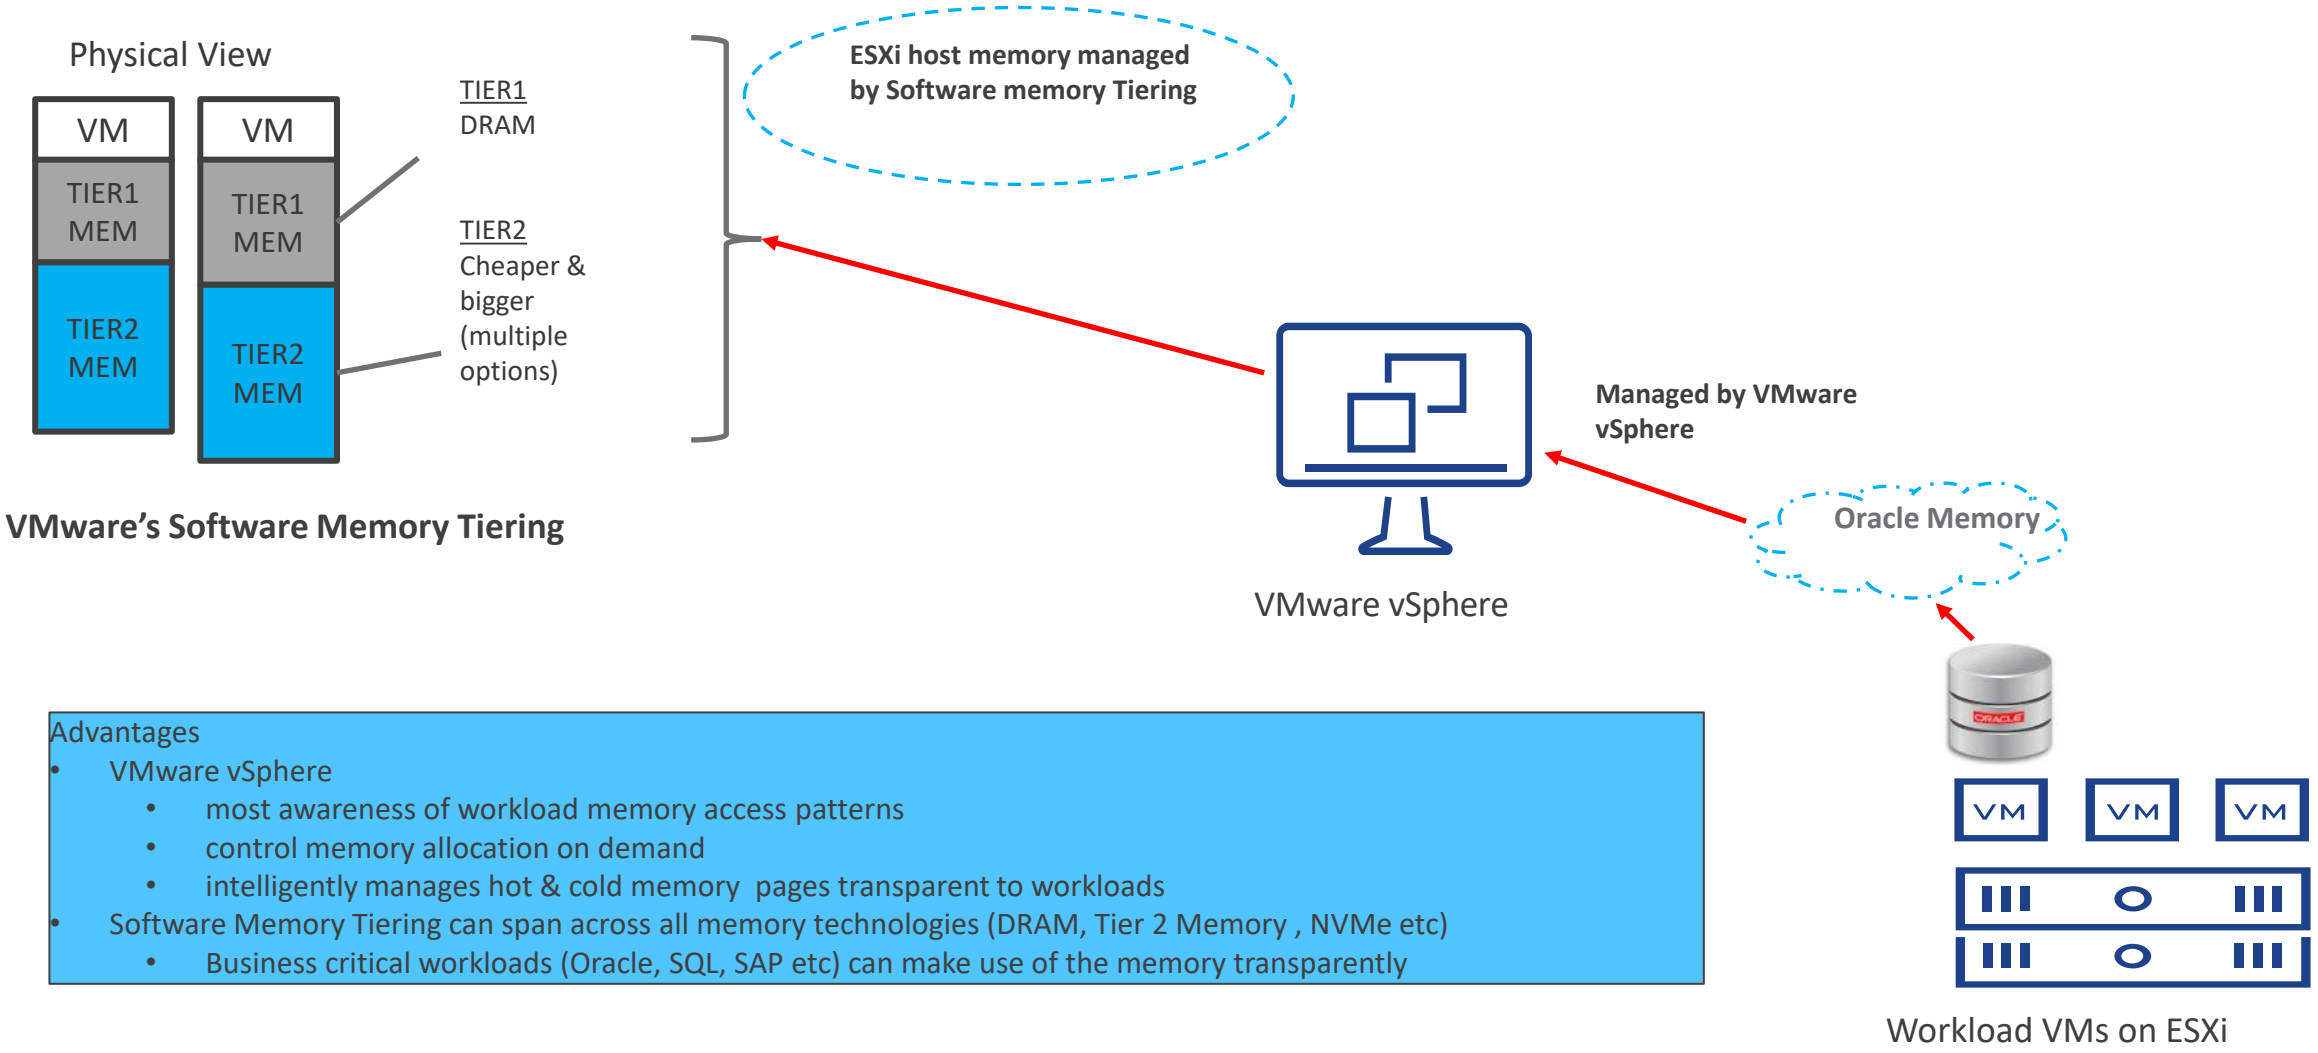

Advantages

- 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

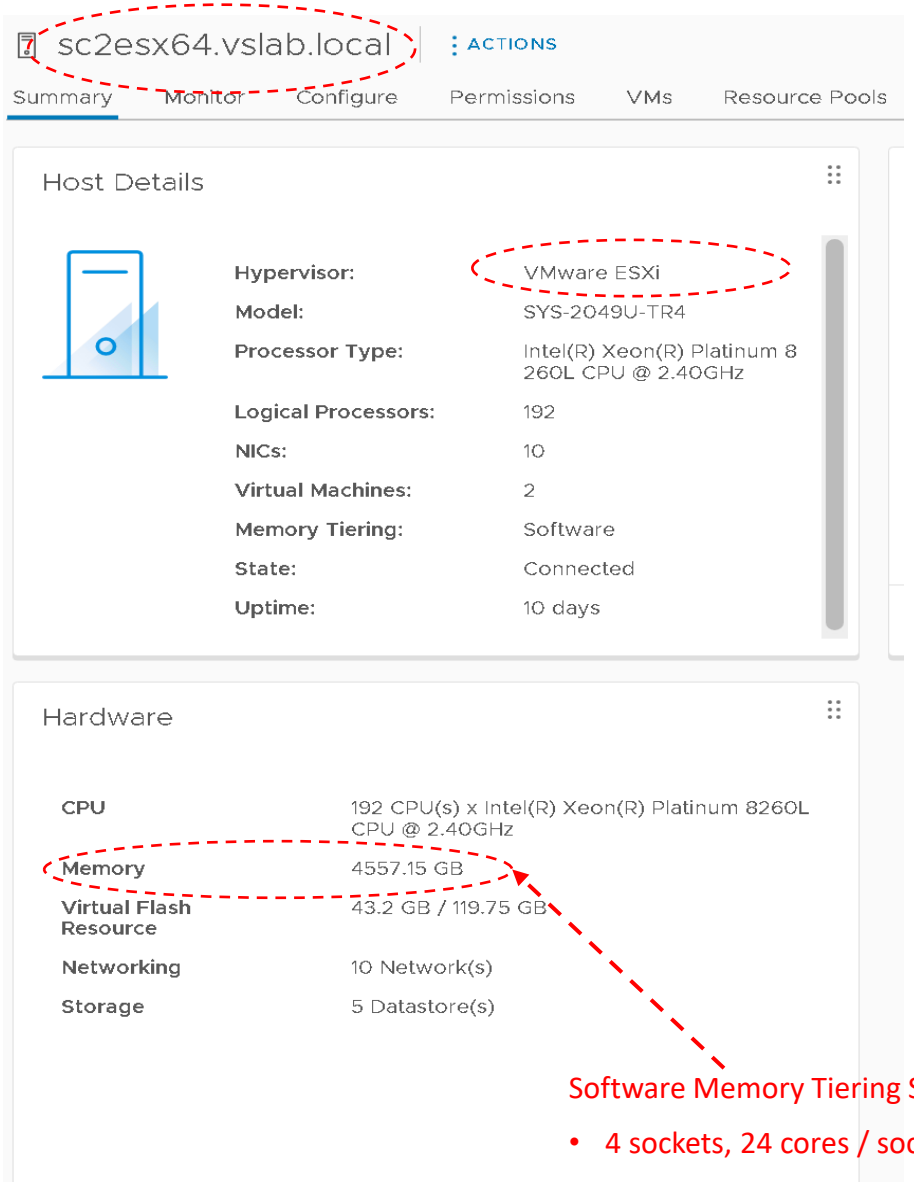

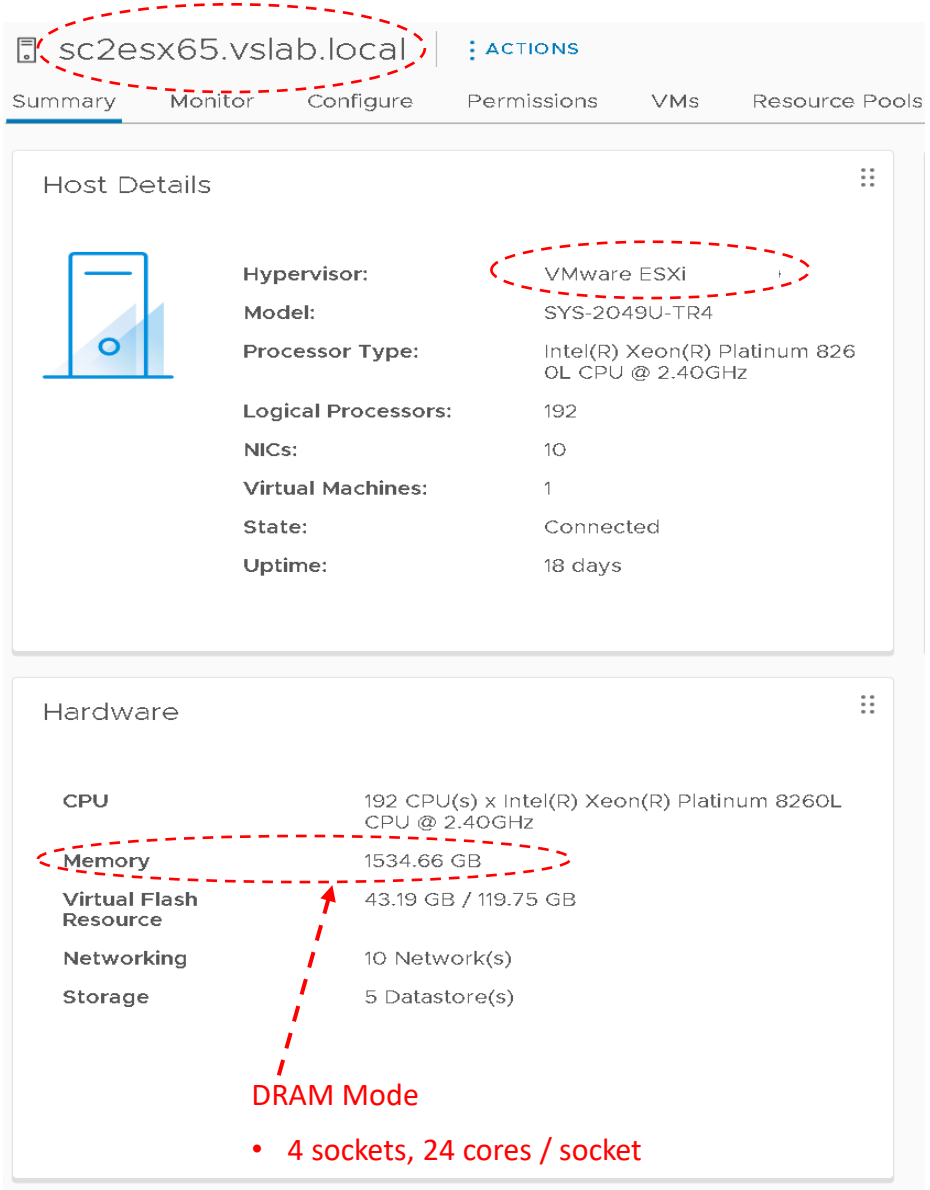

Software Memory Tiering Server

- 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

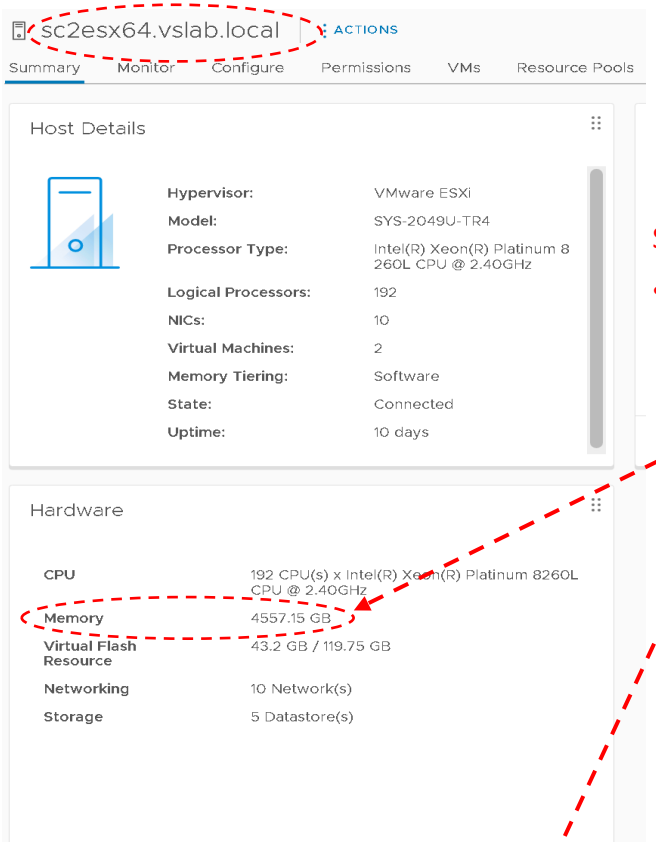

Software Memory Tiering Server

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

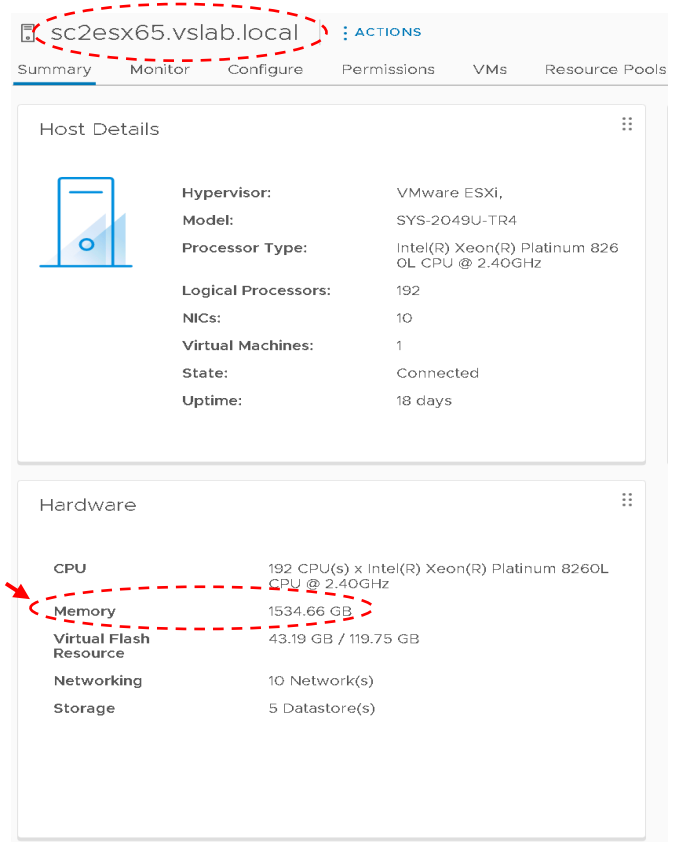

DRAM Mode Server

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

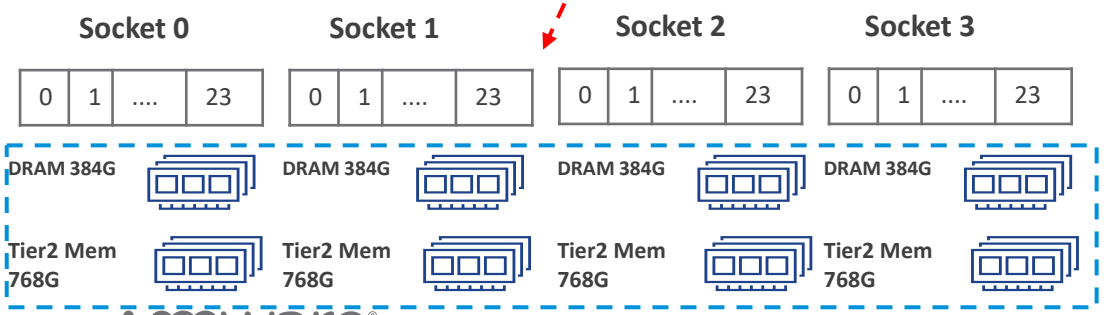

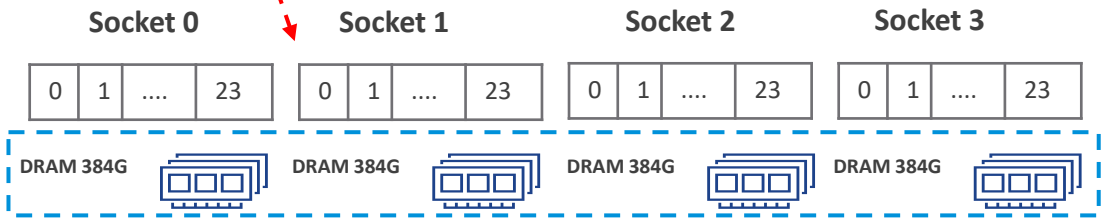

VMware Software Memory Tiering – VM Details

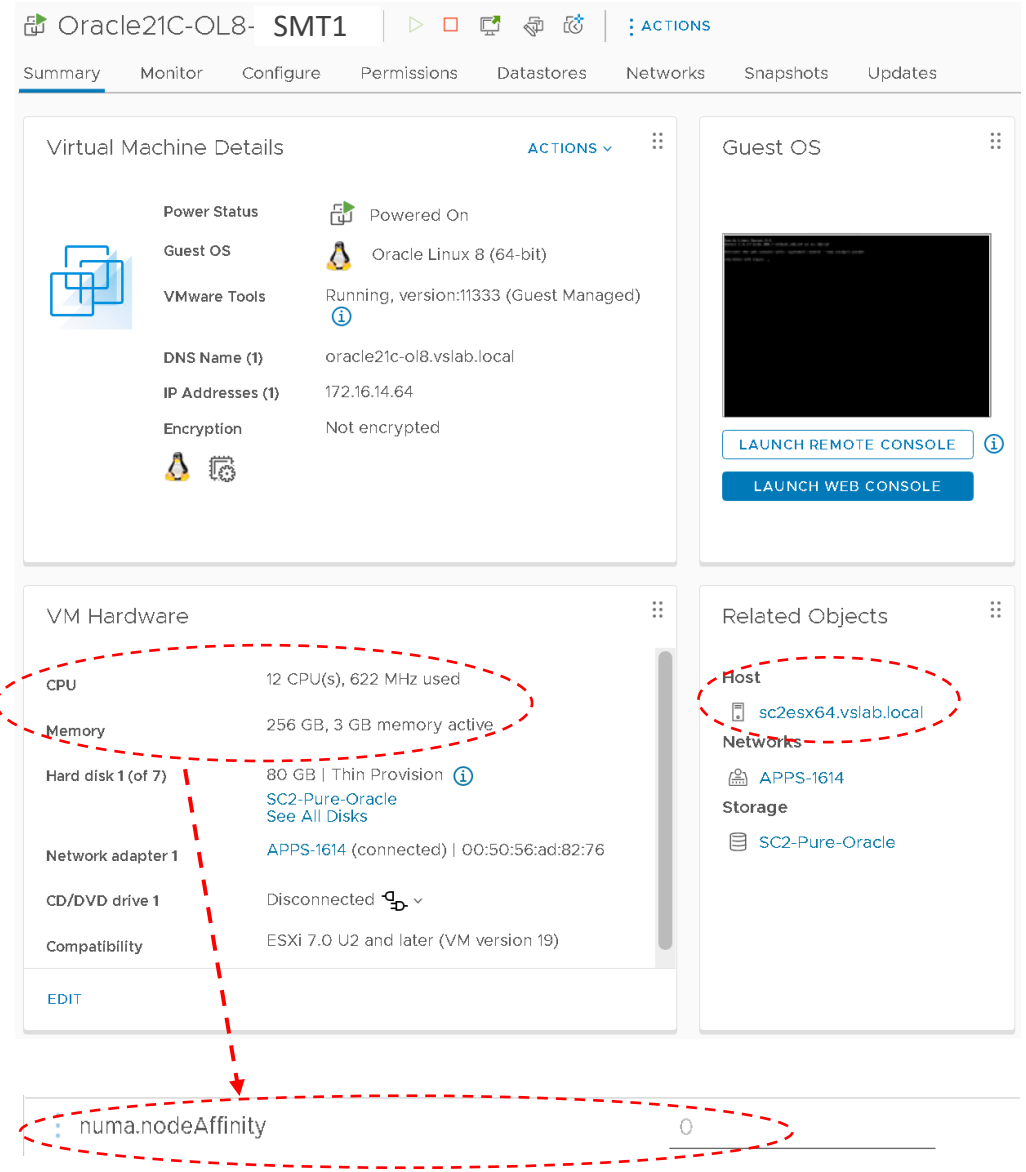

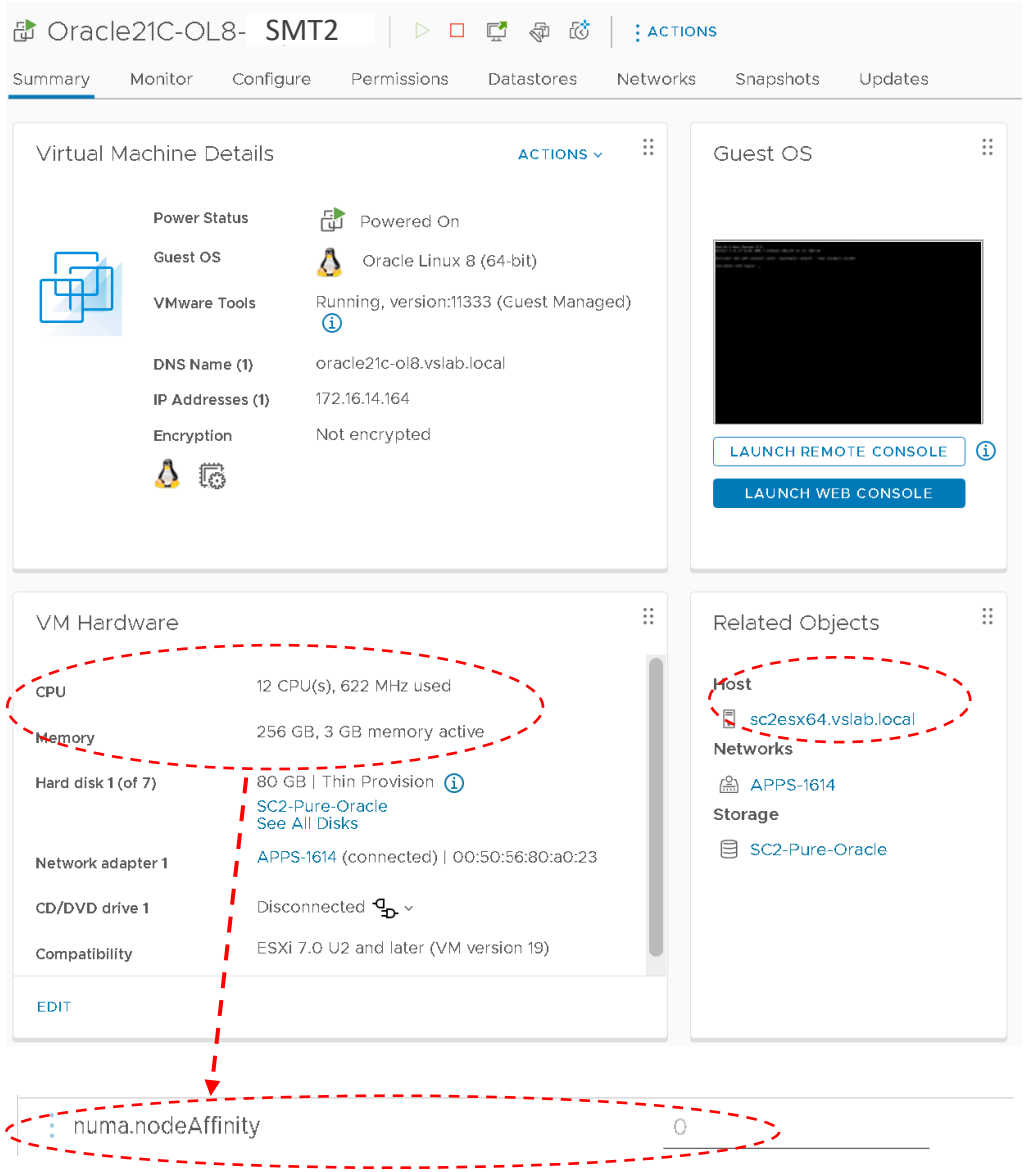

DRAM Mode – Virtual Machine Details

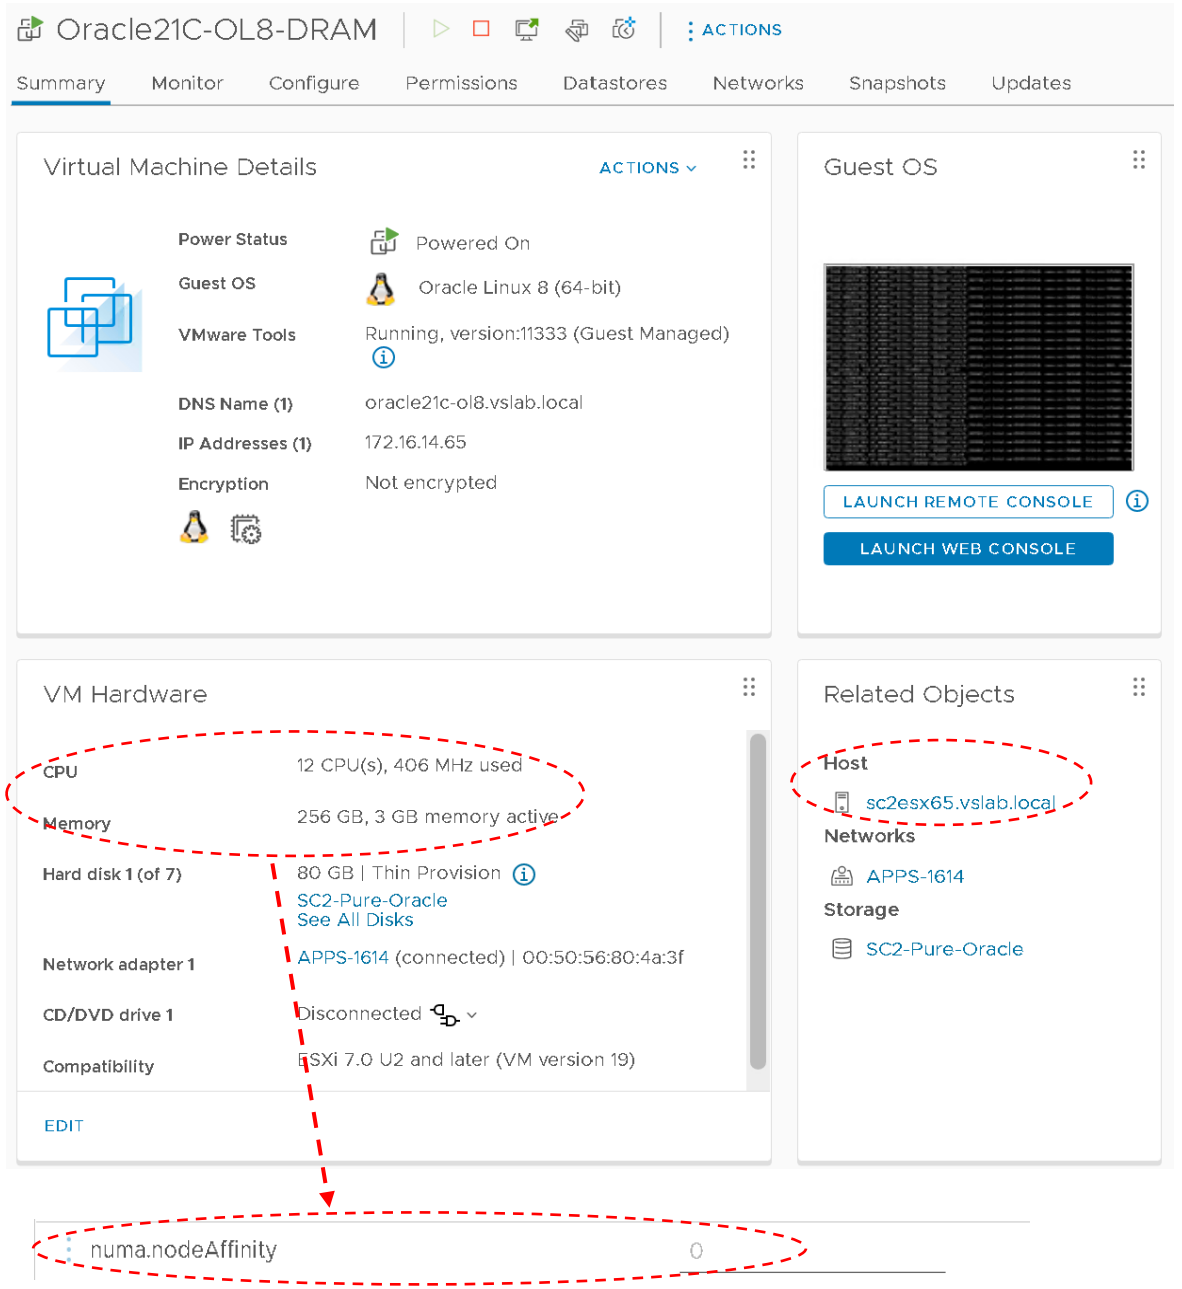

Oracle Database – Details

- 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

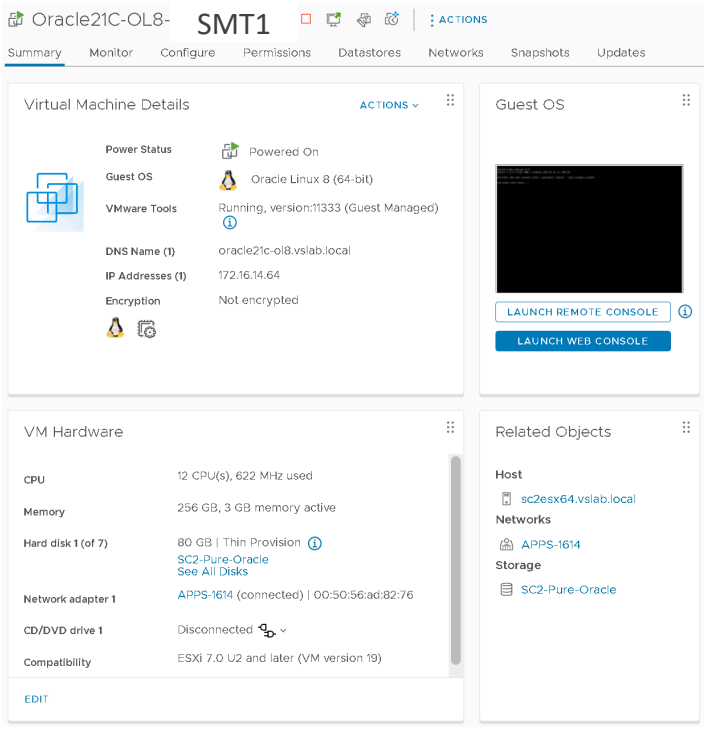

SW Memory Tiering VM1

Oracle Memory

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

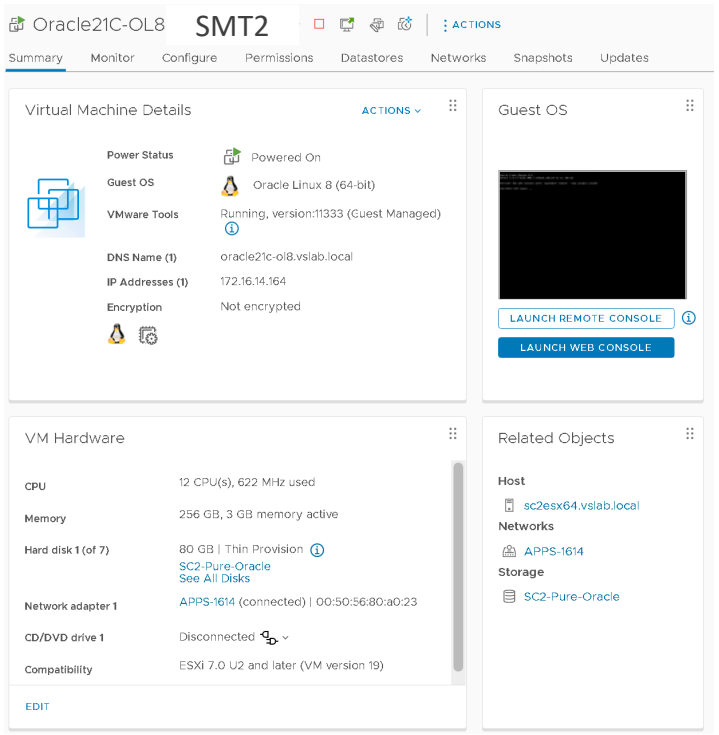

SW Memory Tiering VM2

Oracle Memory

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

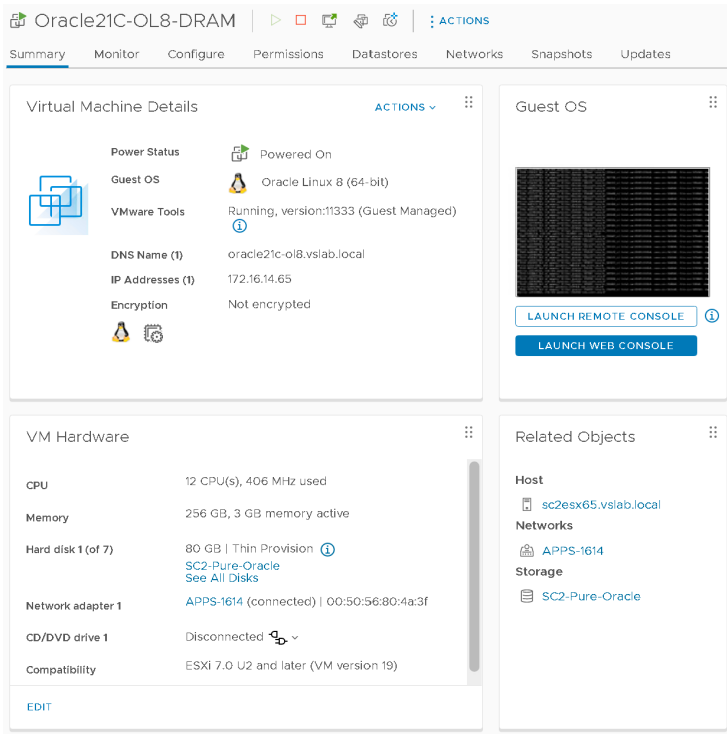

DRAM Mode VM1

Oracle Workload on SW Memory Tiering & DRAM Mode - Metrics

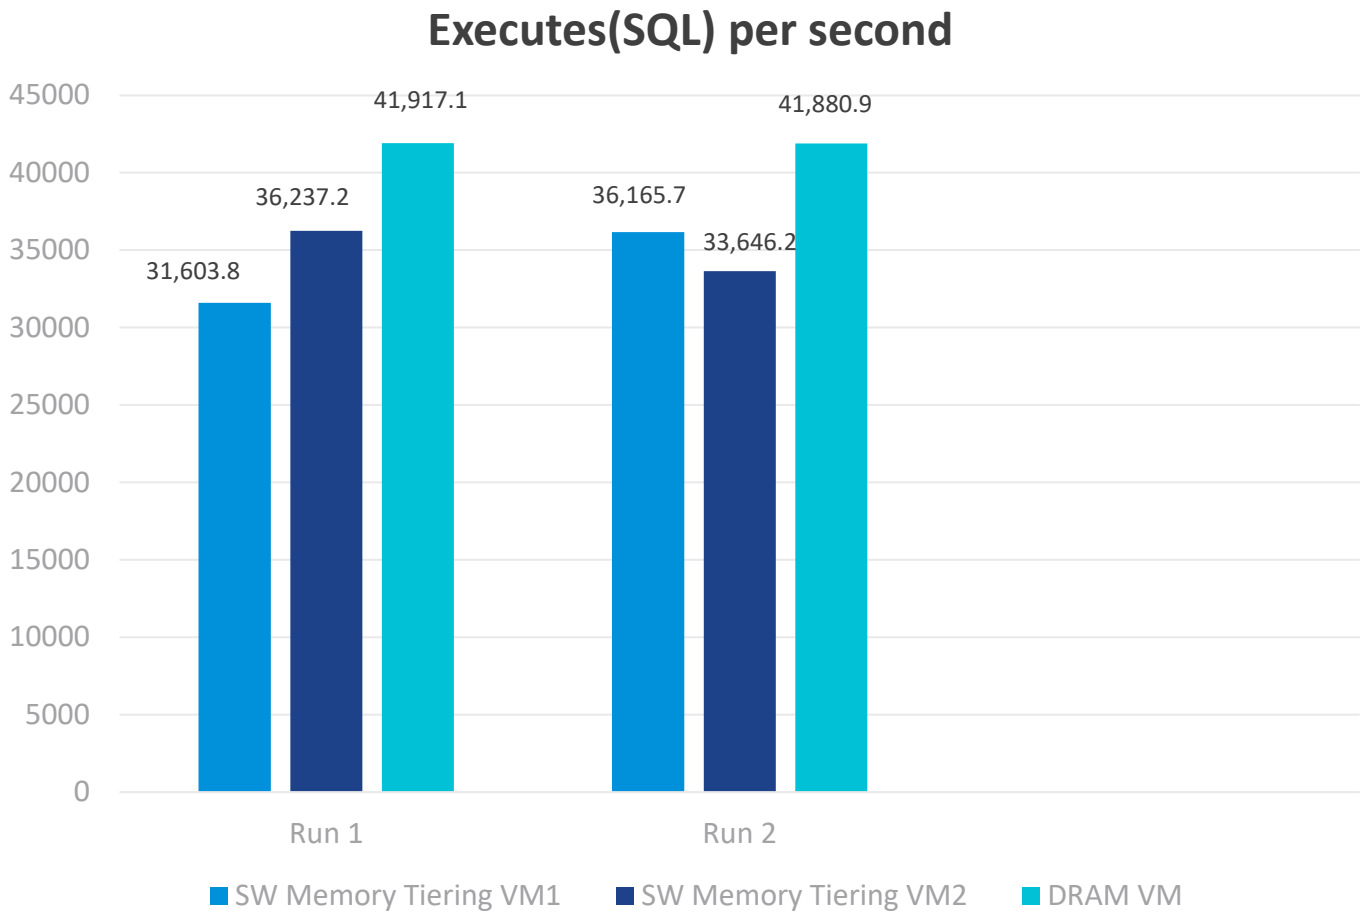

- Load Generator chosen as SLOB 2.5.4.0
  - UPDATE_PCT=0
    - READ only test
    - performance comparison between SW Memory Tiering v/s DRAM Mode
  - RUN_TIME=1200 secs(20mins)

- Test Results
  - Executes(SQL) / second
  - Run 1
    - SW Mem Tier VM1 – 33,603.8/sec
    - SW Mem Tier VM1 – 36,237.2/sec
    - DRAM Mode VM - 41,917.1/sec
  - Run 2
    - SW Mem Tier VM1 – 36,165.7/sec
    - SW Mem Tier VM1 – 33,646.2/sec
    - DRAM Mode VM - 41,880.9/sec

Oracle Workload on SW Memory Tiering & DRAM Mode – More Metrics

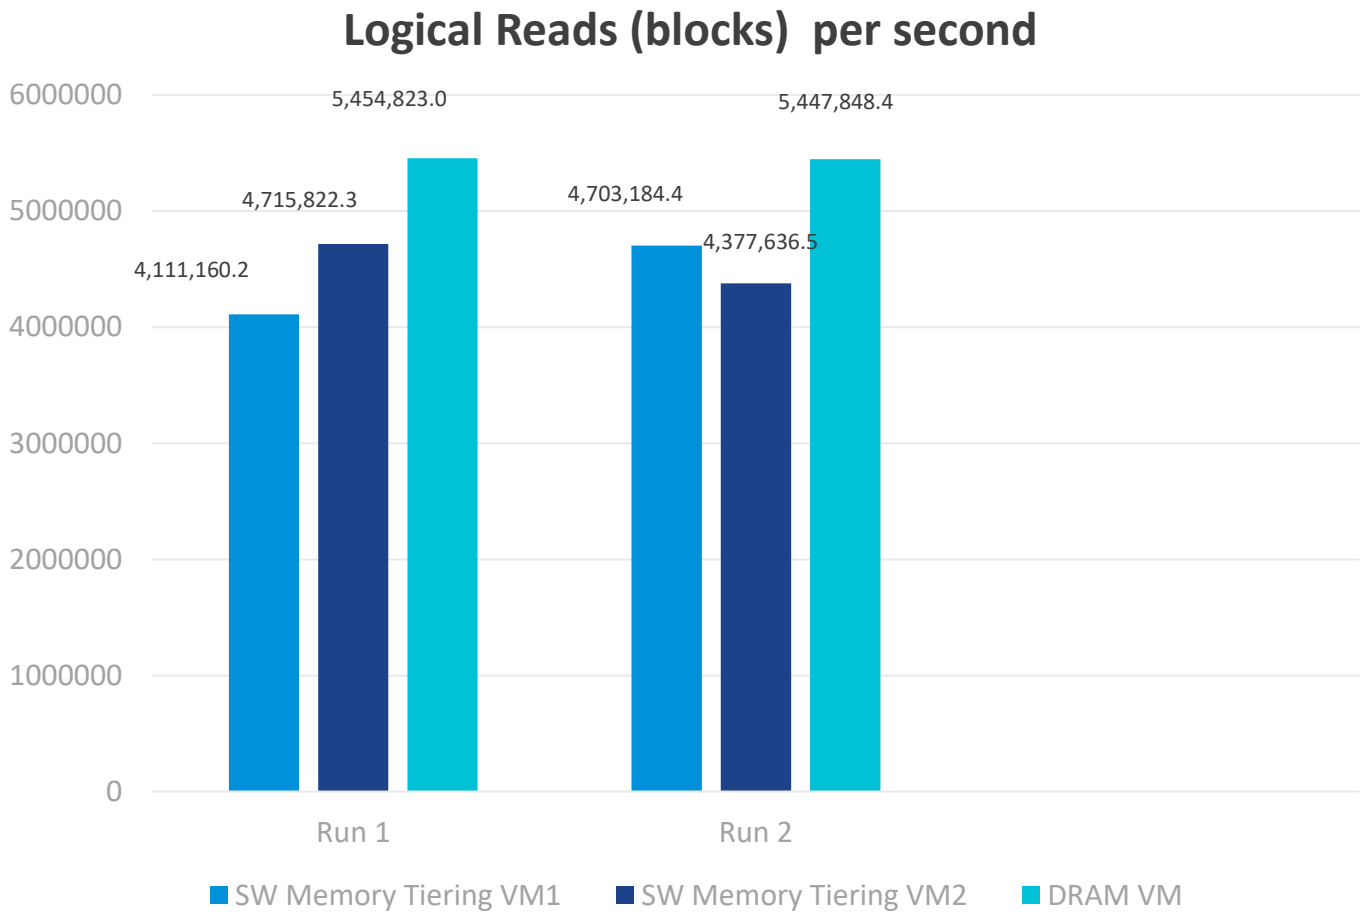

- Test Results
  - Logical Reads (blocks) per second
  - Run 1
    - SW Mem Tier VM1 – 4,111,160.2/sec
    - SW Mem Tier VM1 – 4,715,822.3/sec
    - DRAM Mode VM - 5,454,823.0/sec
  - Run 2
    - SW Mem Tier VM1 – 4,703,184.4/sec
    - SW Mem Tier VM1 – 4,377,636.5/sec
    - DRAM Mode VM - 5,447,848.4/sec

Oracle Workload on SW Memory Tiering & DRAM Mode – OS Metrics

Test Results – Guest Operating System Statistics

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

Oracle Workload on SW Memory Tiering & DRAM Mode – 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

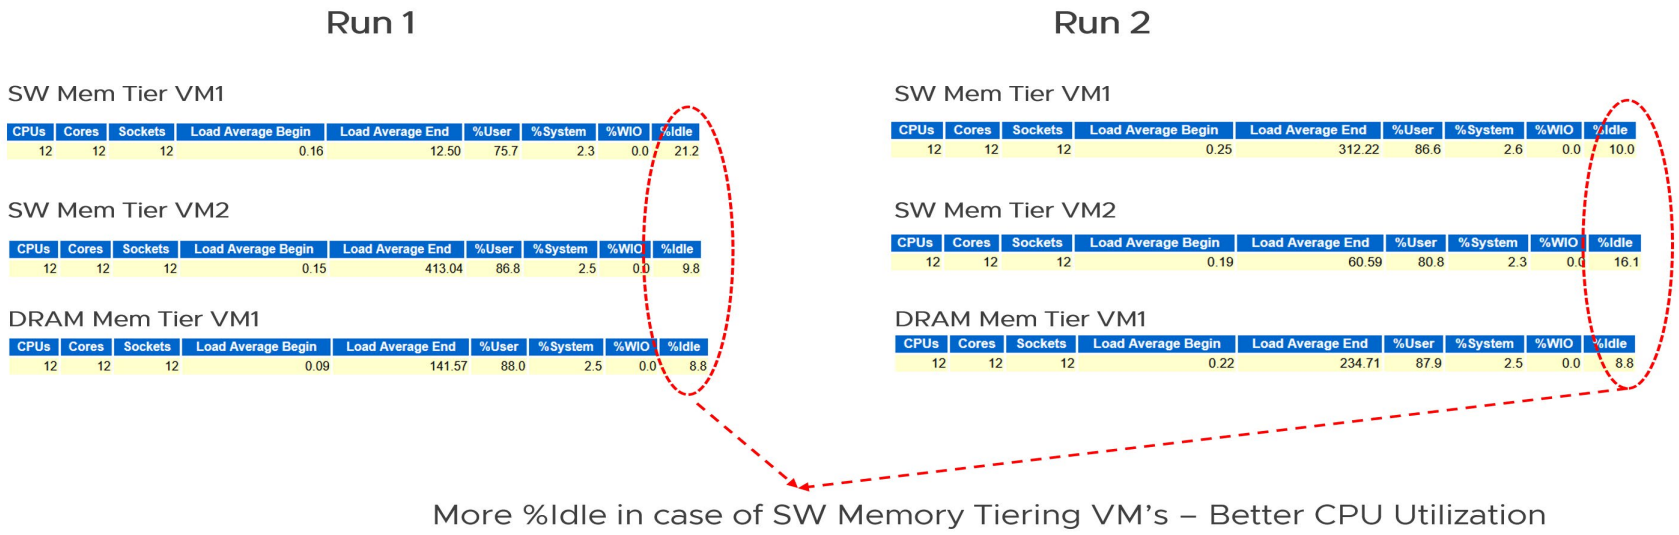

We were able to run '2 SW Memory Tiering' VM's on 1 NUMA node as compared to '1 DRAM VM' on 1 NUMA node, even with the current VM size and DRAM capacity constraints because of the 'Software Memory Tiering' capability

Key Takeaways

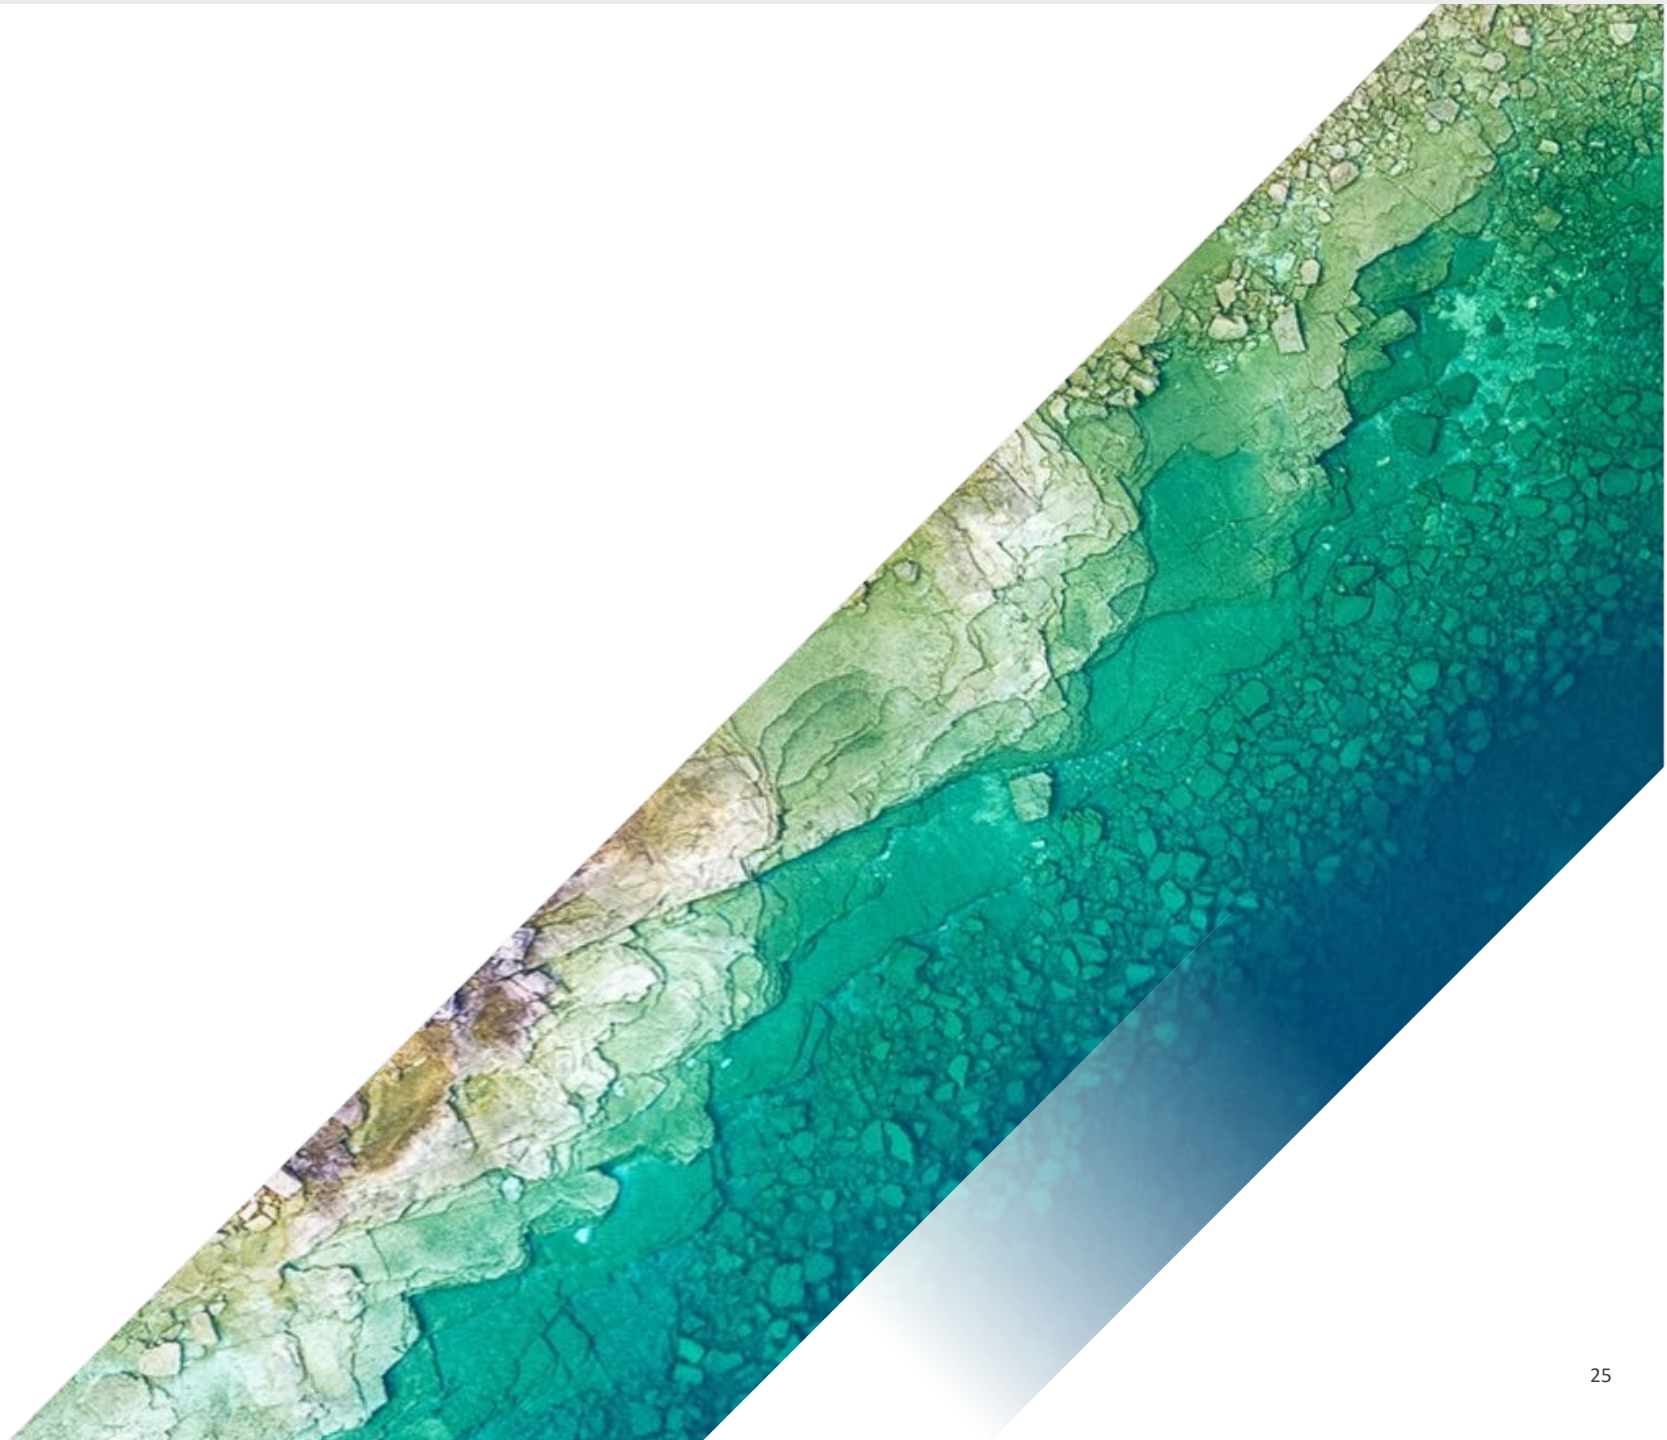

Summary

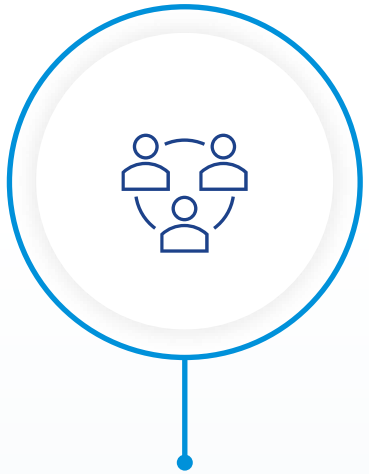

Intel and VMware have been collaborating to bring tiering to address cost, scale, and density challenges

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

- **26 + years Oracle hands on experience -** Principal Oracle DBA / Architect, Oracle RAC/Data Guard Expert, Experienced in EMC SAN Technologies
- 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 [2012-] Senior Staff Solution Architect & Global Oracle Practice Lead
- VMware VCA – Cloud ,VMware vBCA Specialist, **VMware vExpert**
- **Member of the Office of the Chief Technical Ambassador VMware (Alumni)**
- **Oracle ACE**
- **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/

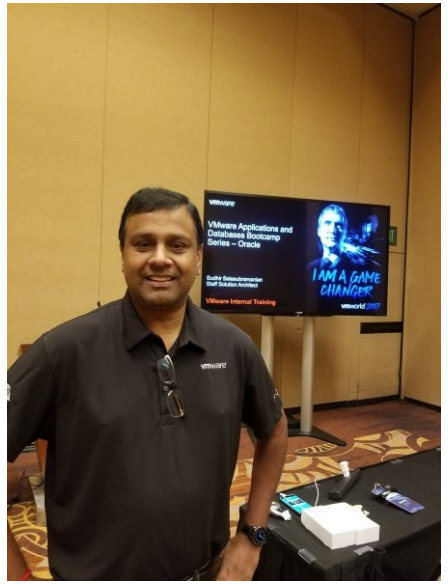

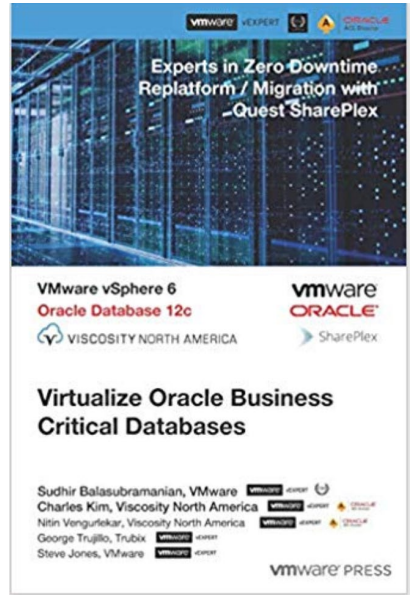

Arvind Jagannath – Product Line Manager for vSphere Platform

- 25+ years of experience in the industry working on networking, storage, embedded, and kernel development
- Leads infrastructure and core platform enablement for vSphere, working across the VMware ecosystem of server, IO, and storage partners
- Most recently drove platform product management at Cohesity and NetApp
- MBA from the University of Chicago, Booth school of Business and a Bachelors in Computer Science and Engineering

Please take a moment to rate this session.

Your feedback is important to us.

Questions ?