

Ask the Expert

Introduction to Making VMware Perform
With SAS®

Tony Brown, Distinguished Software Performance Engineer





Tony Brown

Distinguished Software Performance Engineer

Tony is an engineer in the SAS Platform R&D Performance Lab, where he conducts software performance testing on a broad variety of hardware virtualization, servers, storage types and networks. He jointly tests and works with operating systems, file systems, hardware and third-party vendors to ensure SAS performance on their platforms and software. Tony globally supports SAS system customers with performance on reporting and analysis, statistical modeling, vertical solutions, SAS Grid, and systems using Cloud Analytics Service and SAS[®] Viya[®]. He provides SAS performance architecture advice for on-premises and cloud-based solutions.



Making VMware® Perform with SAS

Part 1 - Introduction

What is VMware®?

Topics


- What is Virtualization?
- VMware Products and Layout
 - Hardware, vSphere®, ESXi®, vCenter®, Management Tools, VMDKs®, vLans®
- What is a Hypervisor?
- Tenancy
- Benefits and Uses
- VMs vs. Containers
- Dangers
 - Thin Provisioning Shared Resources, VM Placement & NUMA, VMDK Performance, vLan Overloading, vMotion

What is VMware?

Virtualization


Dedicated Physical Hardware

- Single OS
- No Separation
- Not Shareable
- Not Movable
- Unused Resources
- Manually Managed



Shared Physical Hardware

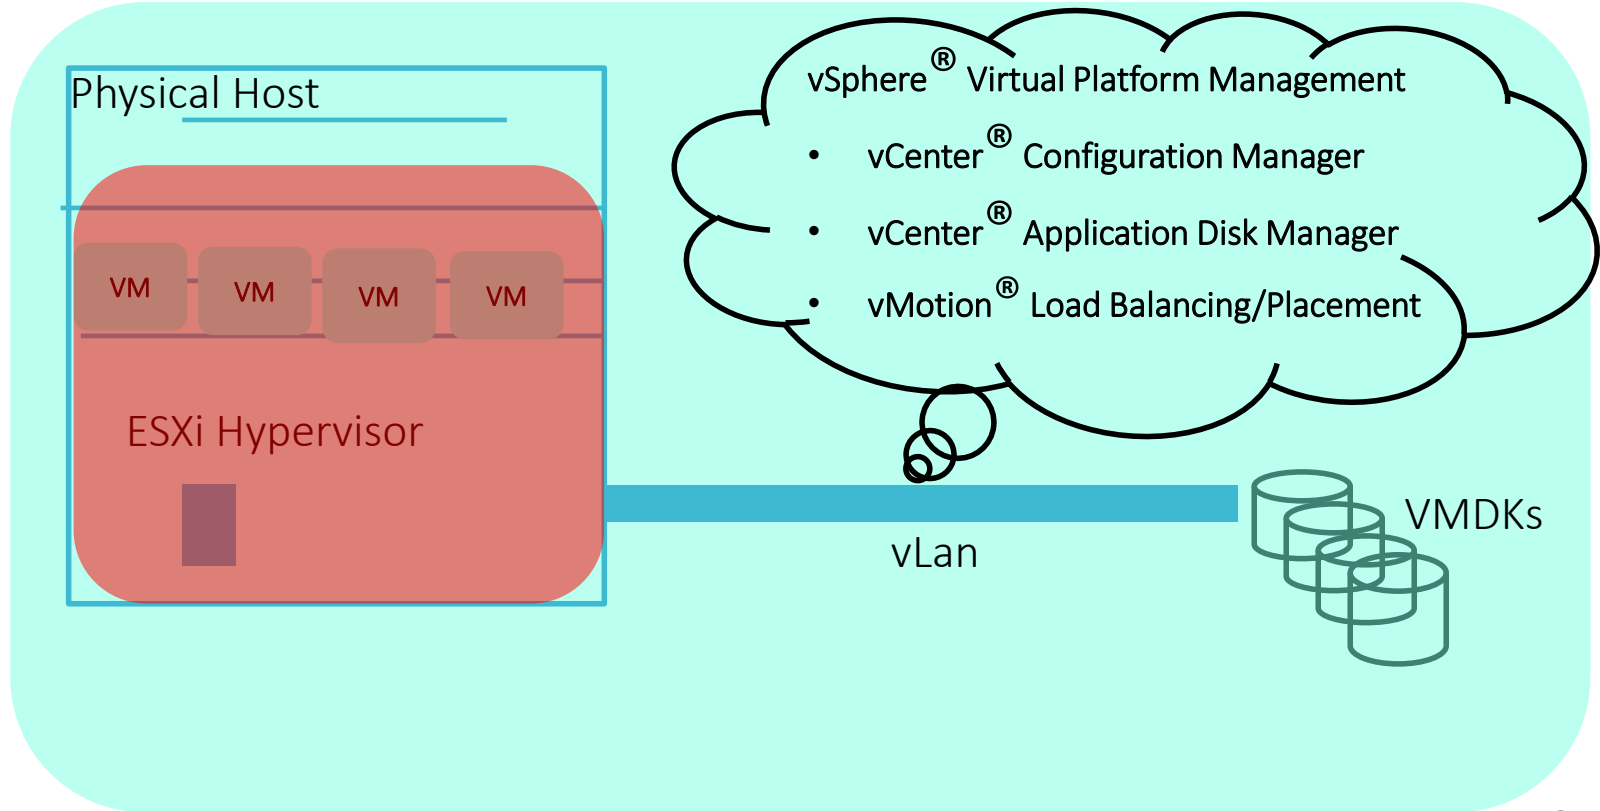
- Multiple OS
- Separation of Apps
- Shareable
- Movable
- Fully Utilized Resources
- Automatically Managed



Thanks to a Hypervisor!

What is VMware?

VMware Products



What is VMware?

Hypervisors

- A hypervisor creates a software version of hardware. It creates a hard drive and log file to configure, monitor, track, and manage a virtual machine.
- There are two types of hypervisors:
 - Type 1 – The hypervisor is a custom OS and replaces the Linux OS on a bare metal Server. It is used to configure virtual machines (VMs) upon the Server, that have their own Guest OS, resource allocations, application space, and security, separate from other VMs. Each VM can have a different OS!
 - VM templates are created for easy replication, i.e. vSphere Client. Most Hypervisors you encounter in a server space are Type 1 – VMware, KVM, Cloud Hypervisors (Google, Amazon, Azure, etc.). The VMware hypervisor is ESXi (Extended Sky X Integrated).
 - Type 2 – The Type 2 hypervisor resides as a guest application on a Server, or User Host, allowing the creation of a separate VM on that host (like vMWorkstation® for Windows®). If you have a VM on your personal computer, this is what governs it.

What is VMware?

Benefits of the VMware

- Ability to template VM Definitions for central management and re-use, scaling, across a large enterprise, leveraging a fixed physical host volume definition for high volume hardware discount (one size fits all)
- Flexibility in creating multiple template definitions for disparate application classes/workloads, hardware classes and models
- Provision of security from a central control point (VShpere)
- Ability to place and move VMs around to maximize resources, avoid hotspots
- Automated Maintenance from a central control point
- Highly beneficial for:
 - Cloud, Devops, Test Systems, Security Operations and Testing, Incompatible Software Applications provisioning on different OS, Security, and workload footprints.
- Manage the entire farm with the push of a button, not just one cow.

What is VMware?

Tenancy

- VMs can be set up as multi-tenant or single tenant
- Multi-Tenant
 - Multiple VMs with disparate applications can share the same underlying hardware infrastructure (**most common – SAS Least Preferred**)
 - If you reboot this VM(s) you can end up anywhere, sharing anything, with anyone
- Single Tenant
 - Single Tenant Dedicated Host – The VM Tenants own the entire host infrastructure, not sharing with other VMs (**SAS Most Preferred**)
 - If you reboot this VM(s) you will end up on the same physical machine
 - Single Tenant Instance – VMs are dedicated and isolated, with the same control over workload placement as Dedicated, but not isolated a single physical host! (**SAS Less Preferred**)
 - If you reboot this VM(s) you can end up on another physical machine

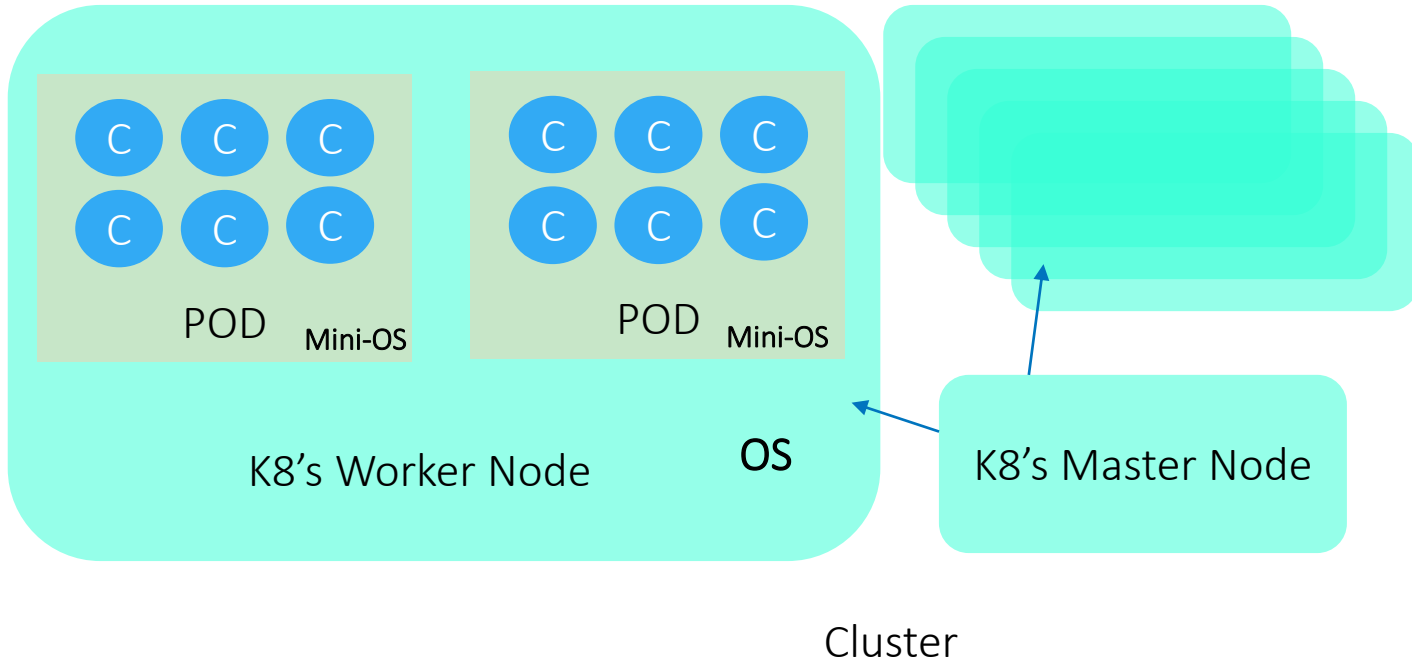
What is VMware?

VMs vs. Containers – What is the difference?

- A VM uses a hypervisor to be virtualized (and placed) on Hardware
 - The VM runs on its own Guest OS, has a virtual hardware definition, and its own heavyweight libraries, software resources, etc.
 - These are more like a chunk of a physical machine, running independently
 - Resources are permanent
- A Container uses an orchestrator to be virtualized (and placed) on an Operating System
 - There is no full Guest operating system, applications libraries, and software instances are lightweight in the container, and use only a portion of the underlying OS and hardware resources via the orchestrator.
 - The same Container definitions can be provisioned across servers, laptop, cloud, on-prem
 - Resources are ephemeral
- You can put Containers on top of VMware (if you like navigating complex multi-virtual stacks to solve problems)

What is VMware?

VMs vs. Containers – What is the difference?



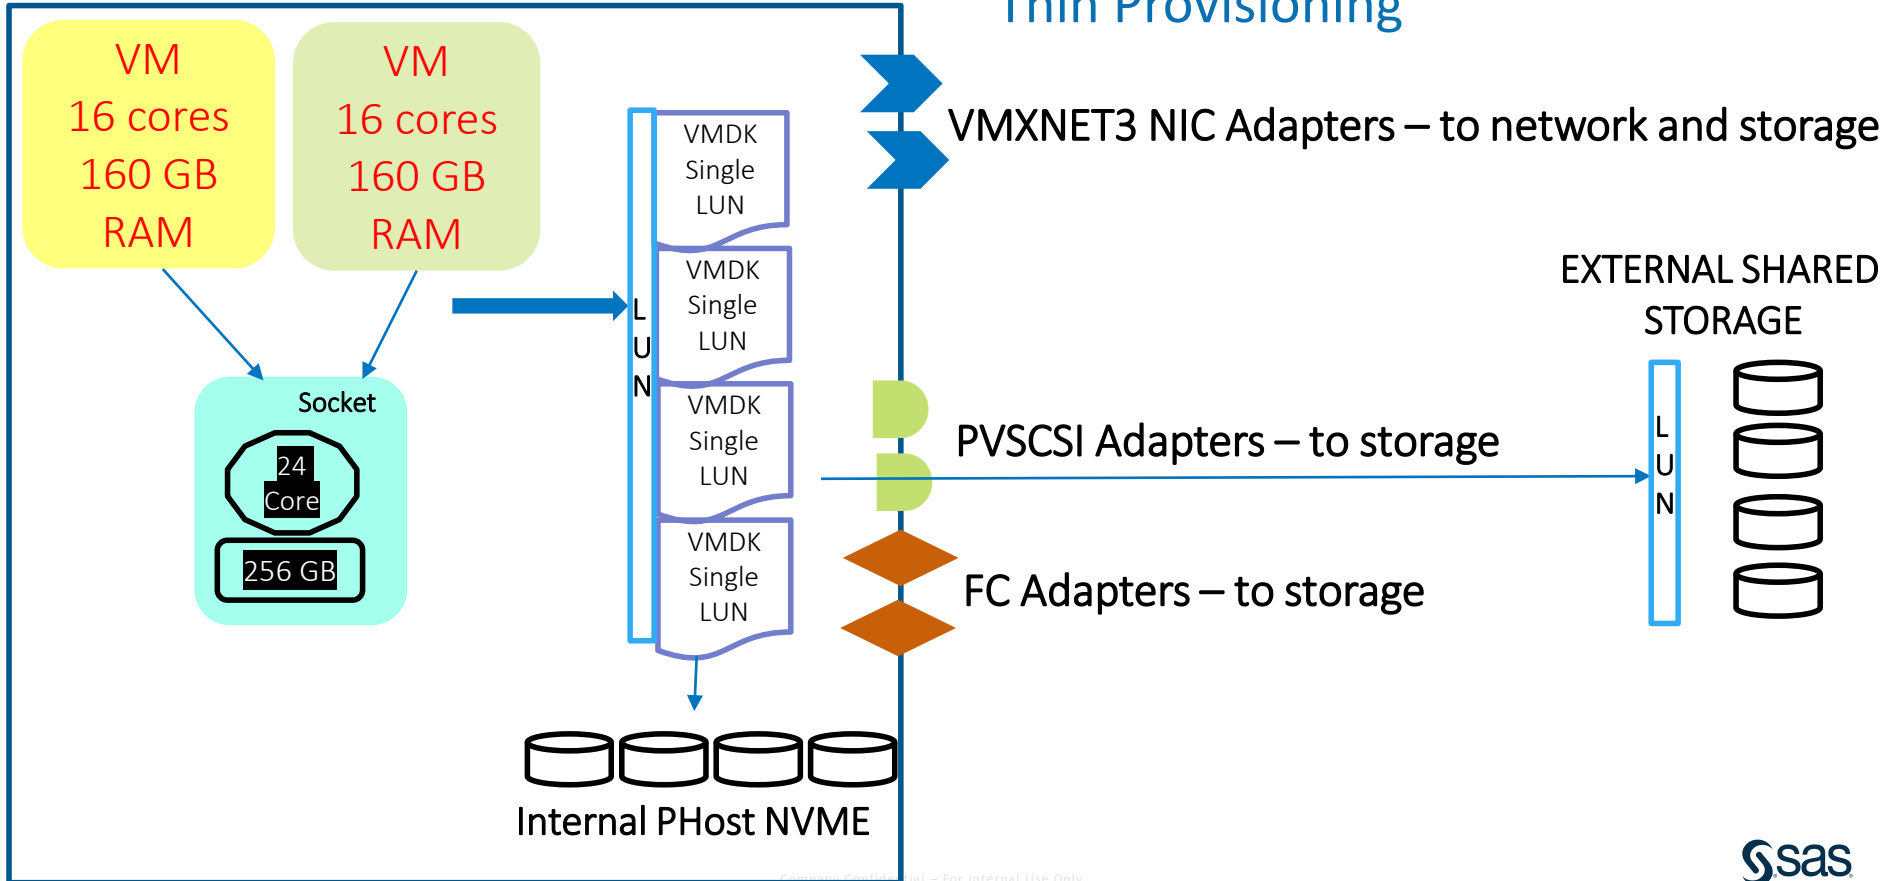
What is VMware?

What are the Dangers?

- The primary Goals of Virtualization are to maximize hardware and software resources, provide a (near) lights-out service platform, and adjust and scale resources to meet workloads
- SAS is a thick application involving CPU, Memory, and IO Intensive resource usage. It relies on resources to be instantly available to assure adequate performance
- VM Systems are often thinly provisioned (the VMs definitions are set up with more resources than the underlying hardware can deliver). Then the underlying hardware gets fully loaded, someone doesn't have what they think – and it quickly shows.
- VMware is typically set up with memory and CPU sharing across VMs in the same cluster. On top of thin provisioning, this can cause unexpected resource shortages if your VM's resources are borrowed.

What is VMware?

Thin Provisioning



What is VMware?

What are the Dangers?

- VMware Farms most often uses shared, network attached storage. SAS IO requirements quite often cannot be satisfied with that speed/bandwidth, especially if the host NICs aren't adequate. This may result in having to create local storage on the host (some shops don't have any), or allocating FC SAN or dedicated non-shared network storage to get around IO insufficiency.
- vLans are shared across the Farm cluster, and SAS can experience noisy neighbors (or be a noisy neighbor) with our large IO requirements and SAS Grid cluster inter-process communications being affected.
- VM Guests are regulated by ESXi, and we cannot always apply our typical Linux OS Tuning that is often crucial to SAS Performance.
- While most VM Shops endeavor to obey NUMA alignment in VM placement on multi-socket hosts, sometimes they miss. If vNUMA isn't used properly VMs can span sockets, causing performance issues with memory utilization.


What is VMware?

Crossing NUMA Boundaries

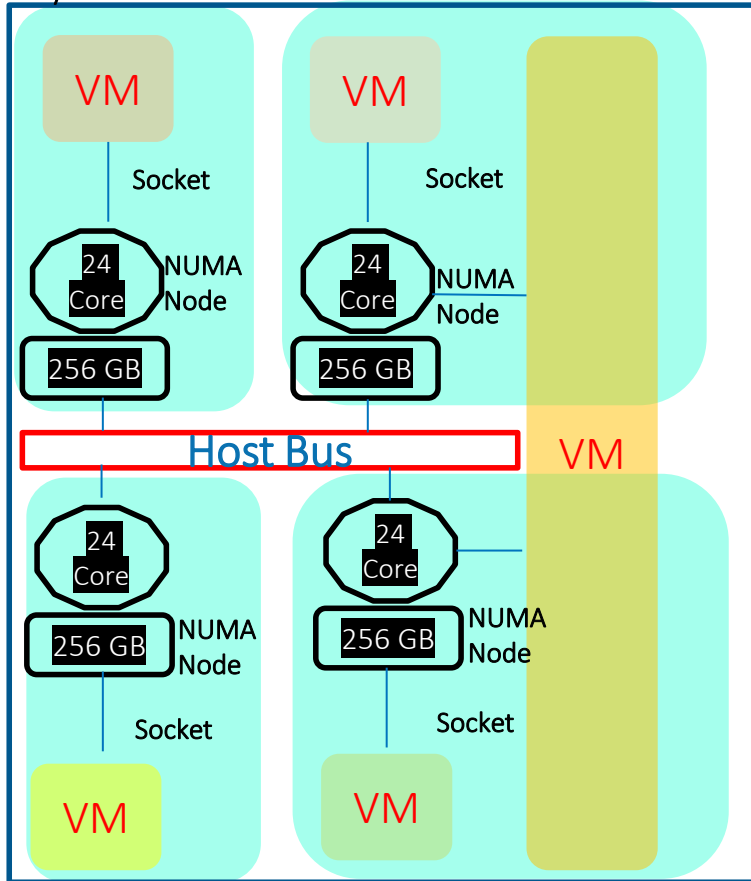
Reaching to access Memory Space when Memory pages are not local....

- Reaching across 2 Sockets
- Reaching Sockets Across the Bus
- Reaching Sockets on another Phost

Slow Mini-Cooper Pulling a Travel Trailer Across the Sahara Sand Dunes....



Physical Host



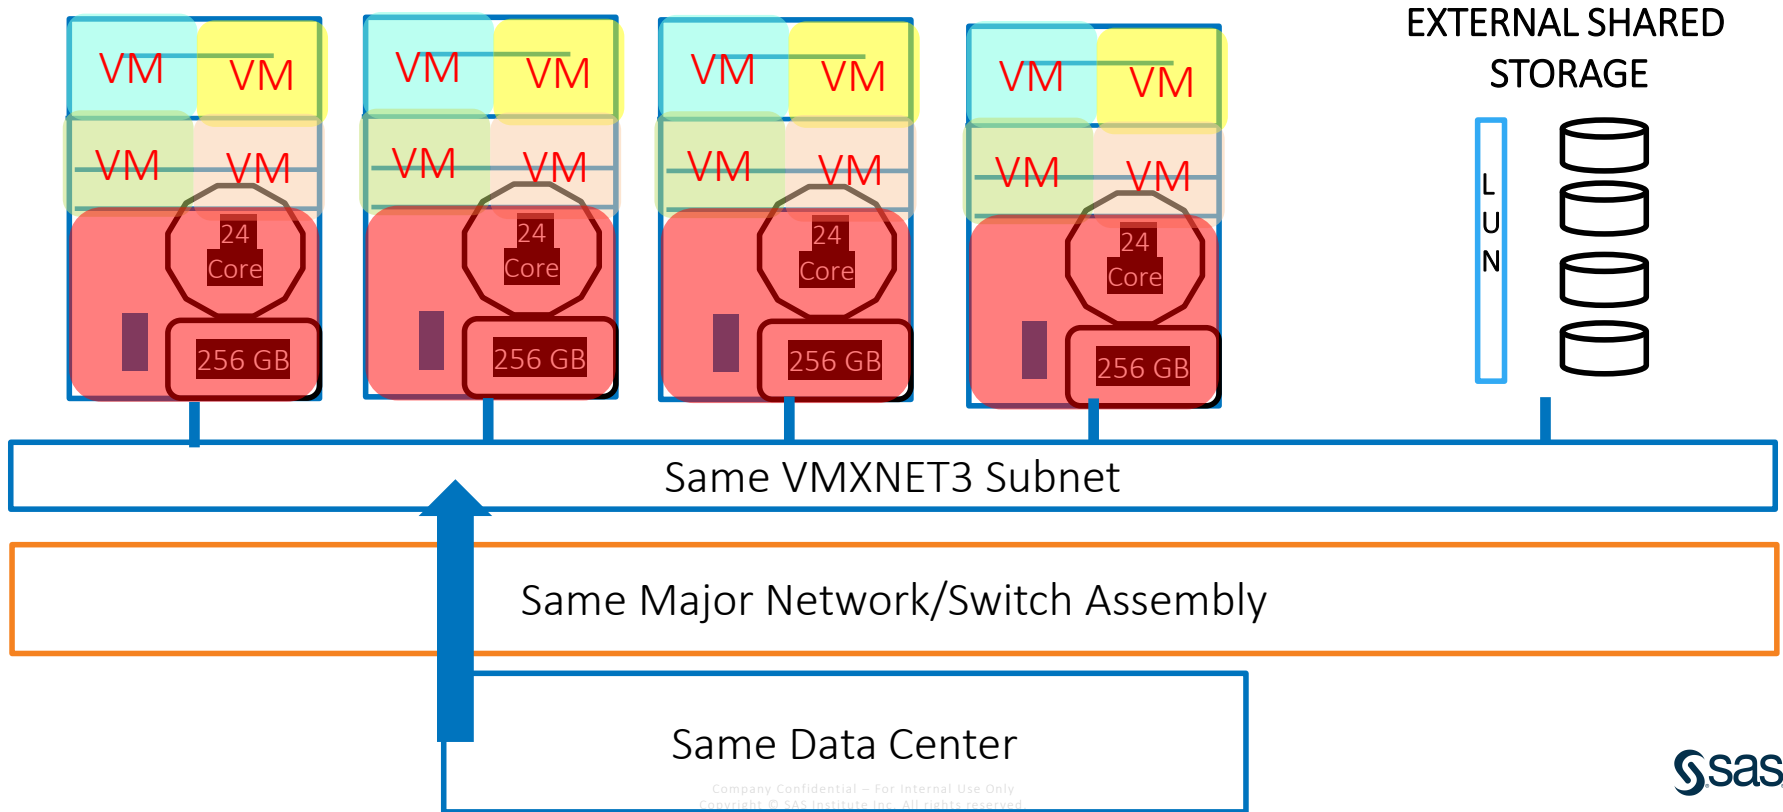
What is VMware?

What are the Dangers?

- Some vSphere Admins adjust host, network, or LAN settings for edge-case workloads. For example changing NIC MTUs from their default, for some, but not all servers. Things like this can cause packet destruction and the failure of SAS applications that rely on the network.
- VM placement for SAS machines in a Grid or Vertical Application Cluster are most stable and perform best with all VMS in the same local VMXNET3 subnet, and at least in the same geographic network cluster. Crossing clusters with SAS machines can cause serious performance and stability degradation.
- VM Storage must be reasonably co-located with SAS VMs, or the IO chain can become too long, and degraded, especially on highly shared storage and LAN/Switch resources.

What is VMware?

Resource Colocation



What is VMware?

- ??? Questions