Virtualizing Microsoft Skype for Business Server on VMware vSphere®

Skype for Business 2015 / vSphere 6.0

BEST PRACTICES GUIDE

vmware®

Table of Contents

List of Figures

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

1. Introduction

Microsoft Skype for Business Server is one of the leading unified communications and collaboration application platforms in the enterprise today. Skype for Business combines video, voice, chat, conferencing, file transfer, and instant messaging capabilities into one enterprise-class suite, enabling organizations to deliver these business-critical services to their users, partners, and customers at a reduced cost, with better performance, resilience, and agility. Skype for Business Server 2015, the current shipping version at the time of writing, replaces Microsoft Lync Server and is a cloud- and virtualization-ready suite of applications suitable for the VMware vSphere® virtualization platform.

While there is currently no official Microsoft guidance for virtualizing Skype for Business Server in the enterprise, Microsoft's technical requirements and recommendations for designing, implementing, and operating a Skype for Business infrastructure provide sufficient frames of reference for customers considering a Skype for Business infrastructure hosted on the vSphere platform.

This document provides technical guidance for VMware customers who are considering virtualizing their Microsoft Skype for Business Server 2015 on the vSphere virtualization platform. This document incorporates some of Microsoft's public guidance and best practices for a Skype for Business infrastructure, coupled with the results of VMware internal testing and validation efforts over several months, both in controlled environment and in production.

VMware encourages customers to adhere to the recommendations provided in this document so that their vSphere hosted Skype for Business infrastructure provides the required performance, stability, resilience, and ease of administration they have come to expect from running other mission- and business-critical applications on the industry's leading virtualization platform.

1.1 Purpose

This guide provides best practice guidelines for deploying Skype for Business Server 2015 on vSphere. The recommendations in this guide are not specific to any particular hardware, nor to the size and scope of any particular Skype for Business Server 2015 implementation. The examples and considerations in this document provide guidance but do not represent strict design requirements, because the flexibility of Skype for Business Server 2015 on vSphere allows for a wide variety of valid configurations.

1.2 Target Audience

This guide assumes a basic knowledge and understanding of vSphere and Skype for Business Server 2015.

- Architectural staff can use this document to gain an understanding of how the system will work as a whole as they design and implement various components.
- Engineers and administrators can use this document as a catalog of technical capabilities.
- vSphere administrators and support staff can use this document to gain an understanding of how Skype for Business Server 2015 might fit into a virtual infrastructure.
- Management staff and process owners can use this document to help model business processes to take advantage of the savings and operational efficiencies achieved with virtualization.

1.3 Scope

The scope of this document is limited to the following topics:

- **VMware ESXi™ Host Best Practices for Skype for Business Server 2015** – Best practice guidelines for preparing the vSphere platform for running Microsoft Skype for Business Server 2015. Guidance is included for CPU, memory, storage, and networking.
- **Using VMware vSphere vMotion®, VMware vSphere Distributed Resource Scheduler™ (DRS), and VMware vSphere High Availability (HA) with Skype for Business Server 2015** – Overview of vSphere vMotion, vSphere HA, and DRS, and guidance for usage of these vSphere features with Skype for Business Server 2015 virtual machines.
- **Microsoft Skype for Business Server 2015 Performance on vSphere** – Background information on Skype for Business Server 2015 performance in a virtual machine. This section also provides information on official VMware partner testing and guidelines for conducting and measuring internal performance tests.

This document is limited in focus to successfully and optimally deploying Skype for Business Server 2015 on the vSphere virtualization platform. Customers should consult Microsoft's guidance for specific Skype for Business Server 2015 deployment considerations for best results.

1.4 External References

This document includes references to external links on third-party web sites for the purposes of clarifying statements where necessary. The statements are accurate as of the time of writing. However, these third-party websites are not under the control of VMware, and the content available at those sites might change.

2. ESXi Host Best Practices for Skype for Business Server 2015

A well-designed vSphere hypervisor platform is crucial to the successful implementation of virtualized enterprise applications such as Skype for Business Server 2015. The following sections outline general best practices for designing vSphere for Skype for Business Server 2015.

2.1 CPU Configuration Guidelines

The latest release of vSphere (vSphere 6.0) has dramatically increased the scalability of virtual machines, enabling configurations of up to 128 virtual processors for a single virtual machine. With this increase, one option to increase performance is simply to create larger virtual machines. While this might appear to be an attractive scalability option, additional factors are involved in deciding how much processing power goes into a virtual machine, so simply creating "jumbo-sized" virtual machines to accommodate Skype for Business workloads might not be the appropriate option in most cases. This section reviews features that are available in vSphere with regard to virtualizing CPUs. Where relevant, this document discusses the impact of those features to Skype for Business Server 2015 and the recommended practices for using those features.

2.1.1 Physical and Virtual CPUs

VMware uses the terms *virtual CPU* (*vCPU*) and *physical CPU* (*pCPU*) to distinguish between the processors within the virtual machine and the underlying physical processor cores. Virtual machines with more than one virtual CPU are also called symmetric multiprocessing (SMP) virtual machines. The virtual machine monitor (VMM) is responsible for virtualizing the CPUs. When a virtual machine begins running, control transfers to the VMM, which is responsible for virtualizing guest operating system instructions.

2.1.2 Skype for Business Server 2015 Capacity Sizing Guidelines

Microsoft provides guidelines for calculating the required compute resources for a single instance of Skype for Business Server 2015 (as an application) so that the infrastructure does not experience unintended performance degradation due to incorrect sizing. These guidelines are the same whether the Skype for Business Server 2015 is virtualized or installed on physical servers.

Please see **Server requirements for Skype for Business Server 2015** for the latest recommendations.

Skype for Business architects and implementers should also consider using one (or both) of the following sizing tools provided by Microsoft to aid in their sizing exercise:

- **Skype for Business Server 2015 Capacity Calculator**
- **Skype for Business Bandwidth Calculator**

Note These tools provide reasonable estimates, but do not account for all possible use cases and deployment scenarios.

The success of the recommendation provided by these tools in satisfying an organization's specific requirements is determined by the accuracy of the input provided.

The recommended compute capacity estimate does not make a distinction among the various platforms on which the Skype for Business infrastructure can be hosted – physical servers, virtualized, or cloud-based.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

2.1.3 VMware vSphere Virtual Symmetric Multiprocessing

VMware vSphere Virtual Symmetric Multiprocessing enhances virtual machine performance by enabling a single virtual machine to use multiple physical processor cores simultaneously. The most recent version of vSphere (ESXi 6.0 Update 2, as of the time of writing) supports allocating up to 128 vCPUs per virtual machine. The biggest advantage of a Virtual SMP system is the ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput (for example, the number of transactions per second). Only workloads that support parallelization (including multiple processes or multiple threads that can run in parallel) can benefit from Virtual SMP.

Be aware that when making your sizing decision, the fact that you could allocate up to 128 vCPUs to a VM should of less importance in this consideration. VMware strongly recommends that you allocate resources to a VM based on the actual needs of the applications hosted on the VM.

The ESXi scheduler uses a mechanism called *relaxed co-scheduling* to schedule processors. *Strict co-scheduling* requires all vCPUs to be scheduled on physical cores simultaneously, whereas relaxed co-scheduling monitors time skew between vCPUs to make scheduling or co-stopping decisions. A leading vCPU might decide to co-stop itself to allow for a lagging vCPU to catch up. Consider the following points when using multiple vCPUs:

- Virtual machines with multiple vCPUs perform well in the latest versions of vSphere, as compared with older versions where strict co-scheduling was used.
- Regardless of relaxed co-scheduling, the ESXi scheduler prefers to schedule vCPUs together, when possible, to keep them in sync. Deploying virtual machines with multiple vCPUs that are not used wastes resources and might result in reduced performance of other virtual machines.

  For detailed information regarding the CPU scheduler and considerations for optimal vCPU allocation, see the section on ESXi CPU considerations in *Performance Best Practices for VMware vSphere 6.0 (*http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf?vmw_so_vex_mande_12%20-%20G6.1006951).
- VMware recommends allocating multiple vCPUs to a virtual machine only if the anticipated Skype for Business Server 2015 workload can truly take advantage of all the vCPUs.
- VMware supports up to 2:1 virtual-to-physical CPU allocation for Skype for Business Server 2015 in a virtual environment. VMware recommends that, for the initial sizing of performance-critical Skype for Business Server 2015 virtual machines (production systems), the total number of vCPUs assigned to all the virtual machines be no more than the total number of physical cores on the ESXi host machine, not hyper-threaded cores. By following this guideline, you can gauge performance and utilization within the environment until you are able to identify potential excess capacity that could be used for additional workloads.
- Although larger virtual machines are possible in vSphere, VMware recommends reducing the number of virtual CPUs for a virtual machine if monitoring of the actual workload shows that the Skype for Business Server 2015 application is not benefitting from the increased virtual CPUs.

2.1.4 CPU Reservations

Setting a CPU reservation sets a guaranteed CPU allocation for the virtual machine. This practice is generally not recommended because reserved resources are not available to other virtual machines, and flexibility is often required to manage changing workloads. However, certain Skype for Business component servers are very sensitive to latency and jitters, which can adversely impact the quality of services provided by those servers. The vSphere hypervisor provides an easy and efficient mechanism to provide exclusive CPU access to these types of workloads, reducing (and in some cases eliminating) the occurrence of such latencies and jitters, and improving the quality of service for these components. This feature will be discussed in greater detail later in this document

2.1.5 Virtual Cores and Virtual Sockets

vSphere now supports configuration of the number of virtual cores per virtual socket in the VMware vSphere Web Client. This feature provides two functions:

- When used with virtual non-uniform memory access (vNUMA)-enabled virtual machines, this setting can be used to present specific NUMA topologies to the guest operating system.
- More commonly, this feature allows a guest operating system to utilize all of its assigned vCPUs in the case of an operating system that is limited to a certain number of CPU sockets.

On vSphere, virtual CPUs can be allocated to a virtual machine by socket, or by number of cores per socket. Controlling vCPUs by number of cores can help you overcome the limitations in previous versions of the Windows operating system and certain versions of Microsoft SQL Server. As an example, Windows Server 2008 R2 Standard Edition and Microsoft SQL Server Standard Editions are limited to a maximum of four CPU sockets. In a physical environment, physical servers with more than four sockets utilize only the total cores in four of those sockets. In previous versions of vSphere, configuring a Windows Server 2008 R2 Standard Edition virtual machine with more than four vCPUs resulted in the operating system and Microsoft SQL Server Standard Edition seeing only four vCPUs. This was configurable within the virtual machine VMX file and is now configurable in the GUI. By configuring multiple cores per socket, the guest operating system and the application can see and utilize all configured vCPUs, up to their technical limits of addressable logical CPUs.

While controlling vCPU allocation through cores per socket helps overcome this limitation, the configuration may impede the efficiency of the vSphere CPU scheduler’s resource optimization algorithms. Without adequate planning and care, this configuration can also present an incorrect NUMA topology to the underlying operating system, leading to a performance penalty and impaired VM mobility within an ESXi cluster.

Ordinarily, because Skype for Business Server 2015 requires a minimum OS version of Windows Server 2012, which does not suffer from the same socket limitations of prior Windows versions, VMware would have recommended that customers allocate vCPUs to a Skype for Business Server 2015 virtual machine using the **Socket** option. VMware generally recommends that customers leave the cores per socket count at one at all times for important (mission-critical) application VMs. Unfortunately, certain dependencies in a typical Skype for Business deployment make this standard recommendation unsuitable, even in a VM running the Windows OS version not susceptible to these limitations.

The following figure shows the configuration interface for controlling vCPU allocation for a VM.

Figure 1. New Virtual Machine CPU Configuration

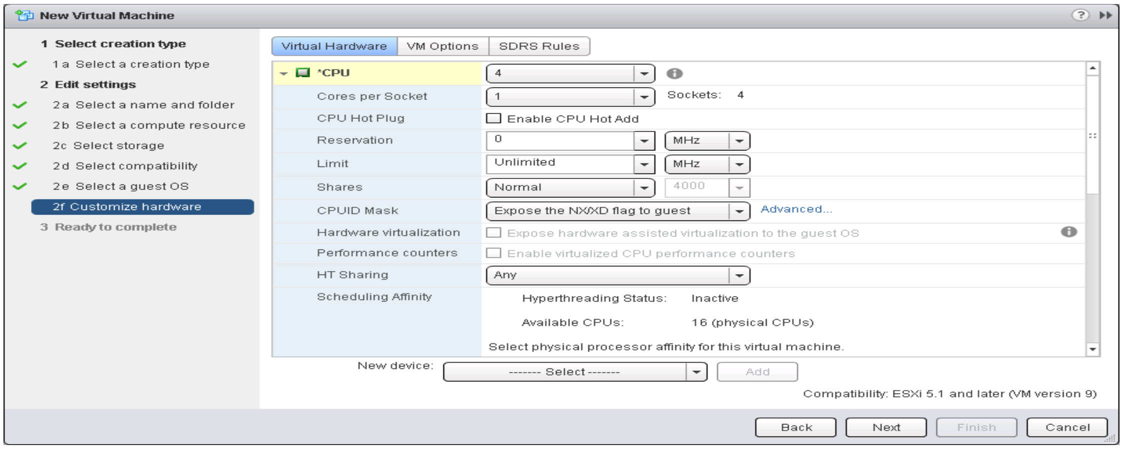

Virtual machines, including those running Skype for Business Server 2015, should be configured with multiple virtual sockets/CPUs. Use the **Cores per Socket** option only if the guest operating system requires the change to make all vCPUs visible. Virtual machines using virtual NUMA might benefit from this option, but the recommendation for these virtual machines is to use virtual sockets (CPUs in the web client).

NOTES:

- Please see Section 2.1.5.1 (Skype for Business Roles and Certain Microsoft SQL Server Editions Limitations) for discussion of important exceptions to the foregoing recommendation.
- Although not generally available as at the time of this writing, Windows Server 2016 alters the previously listed considerations. According to Microsoft's published guidance, both Windows Server 2016 Datacenter and Windows Server 2016 Standard versions will be licensed based on the physical cores allocated to a system, with specific minimum core licenses required. Please see Microsoft's Windows Server 2016 and System Center 2016 - Pricing and Licensing FAQs (May 2016) for more information (PDF).
- Be aware that VMware's recommendation to allocate CPUs to VMs using the socket rather than the Cores per Socket option is based solely on the technical performance and efficiency considerations discussed above.

Although Skype for Business Server 2015 is not a NUMA-aware application and performance tests have not shown any significant performance improvements by enabling vNUMA, the Windows Server 2012 OS is NUMA-aware, and Skype for Business Server 2015 (as an application) does not experience any performance, reliability, or stability issues attributable to vNUMA.

2.1.5.1. Skype for Business Roles and Certain Microsoft SQL Server Editions Limitations:

Some Skype for Business 2015 roles are dependent on Microsoft SQL Express as the backend database engine for their functionalities. Microsoft SQL Express versions are limited in the number of compute resources they can utilize. Typically, this is the lesser of a single CPU socket OR up to four CPU cores.

In addition, a separate, dedicated Full version of Microsoft SQL Server is typically deployed along with Skype for Business component servers in an enterprise-grade Skype for Business Server 2015 deployment. Some versions of the Microsoft SQL Server Standard and Enterprise Editions also have an upper limit on the number of Sockets that they can utilize. For an example, Microsoft SQL Server 2012 Standard edition supports up to four (4) and the Enterprise edition supports up to eight (8).

Applying the standard VMware vCPU resource allocation recommendations of allocating all vCPUs by sockets (and leaving the **"Cores per Socket"** value at the default **"1"**) means that the Windows operating system detects **x** number of sockets (each with 1 core) and reports this to the application. In this configuration, the following behavior will be observed if the "Cores per Socket" count is 1:

- Microsoft SQL Express will be able to utilize only 1 socket (effectively being limited to only one CPU core)
- Microsoft SQL Server Standard Edition will be able to utilize up to four sockets (effectively limiting it to only four CPU cores)
- Certain versions of Microsoft SQL Server Standard Edition will be able to utilize up to eight sockets (effectively limiting it to only eight CPU cores)

The net effect of the situation described above is that, on a highly-utilized Skype for Business role server (e.g. a Front-end Server), the backend database engine becomes a possible cause of performance bottleneck, due to its inability to utilize more than 1 CPU core. Similarly, customers who wish to allocate logical CPUs greater than the number of sockets supported by their version of Microsoft SQL Server will be unable to do so, creating a potential performance bottleneck as a result.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

VMware, therefore, recommends the following methodology for allocating virtual CPUs to any Skype for Business Component Server component VM which uses the embedded Microsoft SQL Express version, and for the standalone versions of Microsoft SQL Server which are limited in addressable CPU sockets:

- Determine the optimal number of vCPUs required by VM
- Verify the physical NUMA topology of the ESXi Host hardware
- Allocate the required number of vCPUs using a combination of “Socket” and “Cores per Socket” counts which reflects the actual physical “Socket” count.

This is illustrated in the table below:

Table 1 - vCPU Allocation with Considerations Microsoft SQL Server Sockets Limitations

ESXi Host Topology (sockets x cores)	vCPUs Required for Skype Component Server VM	Recommended Allocation Mix (Sockets x Cores)
4 x 10	32	4 x 8
4 x 10	16	4 x 4
2 x 8	16	2 x 8
2 x 14	12	2 x 6
4 x 8	12	2 x 6

2.1.6 Hyper-Threading

Hyper-threading technology (also referred to as symmetric multithreading, or SMT) allows a single physical processor core to behave like two logical processors, so that two independent threads are able to run simultaneously on one physical CPU core. Unlike having twice as many processor cores that can roughly double performance, hyper-threading can provide anywhere from a slight to a significant increase in system performance by keeping the processor pipeline busier. For example, an ESXi host system enabled for SMT on an 8-core server sees 16 threads that appear as 16 logical processors. This does not imply that the host has the capacity of 16 physical processors.

Previous guidance provided by Microsoft regarding CPU sizing and the use of hyper-threading led to some confusion among customers who were considering virtualizing their Microsoft business critical applications. Microsoft has updated these recommendations to clarify that statements relating to hyper-threading do not apply to the virtualization platforms. See **Ask The Perf Guy: What’s The Story With Hyperthreading and Virtualization?** for more information.

vSphere uses hyper-threads to provide more scheduling choices for the hypervisor. Hyper-threads provide additional targets for *worlds*, a schedulable CPU context that can include a vCPU or hypervisor management process. For workloads that are not CPU bound, scheduling multiple vCPUs onto a physical core’s logical cores can provide increased throughput by increasing the work in the pipeline. The CPU scheduler schedules to a whole core over a hyper-thread, or partial core, if CPU time is lost due to hyper-thread contention. Consequently, VMware recommends enabling hyper-threading on the ESXi host if the underlying hardware supports the configuration.

2.1.7 Non-Uniform Memory Access

In non-uniform memory access (NUMA) systems, a processor or set of processor cores have memory that they can access with very little latency. The memory and its associated processor or processor cores are referred to as a NUMA node. Operating systems and applications designed to be NUMA-aware can make decisions as to where a process might run relative to the NUMA architecture. This allows processes to access memory local to the NUMA node rather than having to traverse the interconnect, incurring additional latency. Modern versions of the Windows operating systems are NUMA-capable and NUMA-aware. Although we are unaware of specific guidance from Microsoft regarding Skype for Business Server 2015 and NUMA-awareness, extensive VMware testing (in both controlled and production environments) has not revealed any adverse impact on Skype for Business workloads attributable to the exposure of the NUMA features and capabilities of the vSphere virtualization platform.

If the NUMA capability is enabled on an ESXi host at the hardware level, ESXi provides mechanisms for exposing this capability to the operating system running inside a virtual machine on ESXi, enabling such virtual machines take advantage of NUMA. The first mechanism is transparently managed by ESXi while it schedules a virtual machine's virtual CPUs on NUMA nodes. By attempting to keep all of a virtual machine's virtual CPUs scheduled on a single NUMA node, memory access can remain local. For this to work effectively, size the virtual machine to fit within a single NUMA node. This placement is not a guarantee because the scheduler migrates a virtual machine between NUMA nodes based on the demand.

The second mechanism for providing virtual machines with NUMA capabilities is vNUMA. When vNUMA is enabled, a virtual machine is presented with the NUMA architecture of the underlying hardware. This allows NUMA-aware operating systems and applications to make intelligent decisions based on the underlying host's capabilities. By default, vNUMA is automatically enabled for virtual machines with nine or more vCPUs allocated on vSphere. Consider sizing Skype for Business Server 2015 virtual machines to fit within the size of the physical NUMA node for best performance.

The following figure depicts an ESXi host with two NUMA nodes, each comprising four physical cores and 32 GB of memory. The virtual machine that has been allocated with 4 vCPUs and 24 GB of memory can be scheduled by ESXi onto a single NUMA node.

Figure 2. NUMA Architecture Sizing

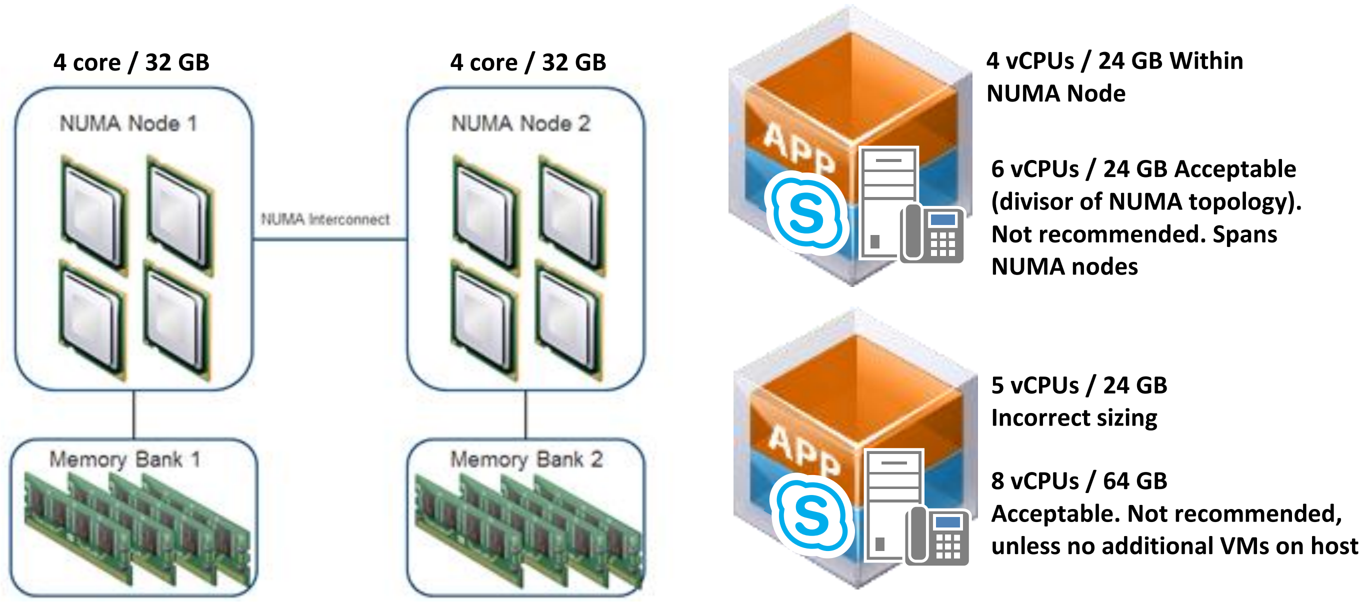

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

The virtual machine that has been allocated 6 vCPUs and 24 GB of memory must span NUMA nodes and might incur some memory access latency. The associated latency can be minimized or avoided through the use of the appropriate combination of vNUMA control options in the virtual machine's advanced configuration options. See **Specifying NUMA Control** in the *vSphere Resource Management* guide.

Customers should adhere to the following sizing approaches when allocating compute resources (CPU and memory) to a virtual machine on the vSphere platform:

- Where possible, size the virtual machine to fit a single NUMA boundary.
- If the virtual machine requires compute resources that exceed the capacity of a single NUMA node available on the physical host, size the virtual machine to fit multiple NUMA boundaries.
- If the virtual machine requires compute resources that are less than the capacity of a single NUMA node available on the physical host, size the virtual machine to fit an even fraction (divisor) of a single NUMA boundary.

For large environments, you can test each configuration to determine whether additional latency associated with remote memory addressing warrants creating additional, smaller virtual machines.

Verify that all ESXi hosts have NUMA enabled in the system BIOS. In some systems, NUMA is enabled by disabling node interleaving.

2.1.8 Virtual NUMA

If a server NUMA capability is enabled at the hardware layer, modern Windows operating systems are capable of automatically detecting and presenting the topology to its resident applications and processes. In a vSphere infrastructure, the hypervisor will expose the physical NUMA architecture to the guest operating system inside a virtual machine if the virtual machine is allocated more than eight virtual CPUs. The topology so presented to the guest VM is referred to as *virtual NUMA*, or *vNUMA*.

Note The number of vCPUs required for exposing NUMA topologies to a virtual machine is configurable in a vSphere environment by using the **numa.vcpu.min = "x"** advanced configuration, where "x" represents the minimum number of vCPUs at which an administrator would like to expose vNUMA to the Windows operating system.

Although the Windows OS can detect vNUMA, the benefits of such awareness depend entirely on the ability of the hosted applications to detect and utilize the enhancements provided by this capability. Microsoft SQL Server is a NUMA-aware application that can take advantage of performance benefits of vNUMA.

2.1.9 vNUMA, vSphere vMotion and Dissimilar ESXi Hosts in a vSphere Cluster

It is possible for a single vSphere cluster to contain ESXi hosts that have different CPU generations, sizes, and topologies. Because vSphere vMotion is a cluster-level configuration, enabling vSphere vMotion in such an environment is possible only through the use of an enhanced vSphere feature called Enhanced vMotion Compatibility (EVC).

EVC simplifies vSphere vMotion compatibility issues across CPU generations. EVC automatically configures server CPUs with Intel FlexMigration or AMD-V Extended Migration technologies to be compatible with older servers. After EVC is enabled for a cluster of ESXi hosts, all hosts in that cluster are configured to present identical CPU features and provide CPU compatibility for vSphere vMotion. The features presented by each host are determined by selecting a predefined EVC baseline. VMware

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

vCenter Server® does not permit the addition of hosts that cannot be automatically configured to be compatible with the EVC baseline.

While EVC enables dissimilar physical hosts to be added to the same ESXi cluster, customers should be aware that there is no normalization of the underlying NUMA topologies presented by the disparate hosts in the cluster. For example, it is possible to mix a host with a 4x6 NUMA topology with a host which has a 2x10 topology. EVC does not cause the 2x10 host to present a 4x6 NUMA architecture or vice versa. A virtual machine evaluates the vNUMA topology presented by the ESXi host only at boot-up time. The virtual machine does not reevaluate this topology while it is in a running state, and the topology persists when the virtual machine is rebooted.

For EVC and vSphere vMotion, the implication of the mixture of this topology evaluation process is the likelihood that a vNUMA-enabled virtual machine could originally start up on a host with one NUMA topology and end up on another host with a different topology in the vSphere cluster described in our previous example. This vNUMA architecture is then relied upon by the guest OS and its applications in their threads and process scheduling. If the virtual machine were to be migrated to the 4x6 ESXi at some later point, there will be an imbalance between the presented and actual topologies. This condition will result in gradual, but persistent and appreciable, decline in the virtual machine's performance for as long as the condition persists. The resolution is a reboot of the virtual machine, enabling it to reevaluate its host's actual NUMA topology so that it could present the correct topology to the OS and applications.

VMware recommends that customers be conscious of this behavior and to be diligent in their monitoring and reporting if using EVC, vSphere vMotion, and vNUMA. Recent vSphere vMotion activities in such environments should be one of the possibilities investigated in the course of troubleshooting a sudden, unexplained performance degradation on an otherwise healthy virtual machine.

2.1.10 vNUMA and CPU Hot Plug

CPU hot plug is a feature that allows an administrator to easily scale up the number of CPUs allocated to a virtual machine while the virtual machine is in a powered-on state, and servicing client requests or executing instructions. This feature reduces the downtime and administrative efforts associated with responding to increasing resource demands on a running workload. While modern Windows operating systems support the hot addition of compute resources to a Windows workload, customers should be aware that many Windows applications are unable to immediately consume the hot-added resources without some administrative interventions. Skype for Business is one such application.

Enabling CPU hot add for a virtual machine on vSphere disables vNUMA for the virtual machine. Given the fact that Skype for Business does not benefit from the CPU hot add feature, VMware recommends that you do not enable CPU hot add for a Skype for Business Server 2015 virtual machine.

Even where an application is capable of utilizing CPU hot plug (for example, Microsoft SQL Server), VMware recommends that customers should avoid relying on CPU hot plug as a shortcut to right-sizing the virtual machine. VMware encourages customers to conduct thorough baseline testing of their critical workloads and to assign compute resources based on known needs. If a virtual machine is later determined to be undersized, customers should use the next scheduled maintenance window to correct the misconfiguration.

Another reason to avoid enabling CPU hot add for a virtual machine running NUMA-aware applications is that the performance benefits of vNUMA to such applications outweighs the administrative convenience provided by the CPU hot add feature.

2.2 Memory Configuration Guidelines

This section provides guidelines for memory allocation to Skype for Business Server 2015 virtual machines. These guidelines consider vSphere memory overhead and the virtual machine memory settings.

2.2.1 ESXi Memory Management Concepts

vSphere virtualizes guest physical memory by adding an extra level of address translation. Shadow page tables make it possible to provide this additional translation with little or no overhead. Managing memory in the hypervisor enables the following:

- Memory sharing across virtual machines that have similar data and same guest operating systems.
- Memory over-commitment – Allocating more memory to virtual machines than is physically available on the ESXi host.
- A memory balloon technique – Virtual machines that do not need all the memory they have been allocated give memory to virtual machines that require additional allocated memory without physically decreasing the amount of memory seen and addressable by the virtual machine.

For a more detailed discussion of vSphere memory management concepts, see the **Memory Virtualization Basics** and **Administering Memory Resources** sections in the *vSphere Resource Management* guide.

2.2.2 Virtual Machine Memory Concepts

The following figure illustrates the use of memory settings parameters in the virtual machine.

Figure 3. Virtual Machine Memory Settings

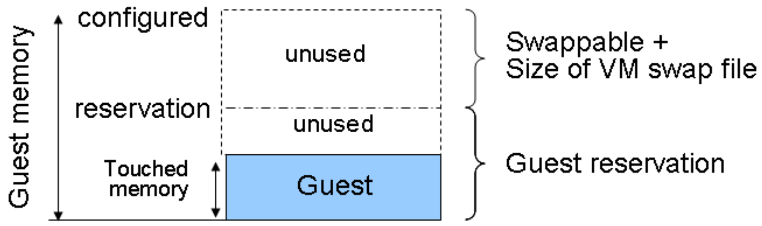

The vSphere memory settings for a virtual machine include the following parameters:

- Configured memory – Memory size of virtual machine assigned at creation.
- Touched memory – Memory used by the virtual machine. vSphere allocates guest operating system memory on demand.
- Swappable – Virtual machine memory that can be reclaimed by the balloon driver or by vSphere swapping. Ballooning occurs before vSphere swapping. If this memory is in use by the virtual machine (touched and in use), the balloon driver causes the guest operating system to swap. Also, this value is the size of the per virtual machine swap file that is created on the vSphere Virtual Machine File System (VMFS) file system.
- If the balloon driver is unable to reclaim memory quickly enough or is disabled or not installed, vSphere forcibly reclaims memory from the virtual machine using the VMkernel swap file.

2.2.3 Allocating Memory to Skype for Business Server 2015 Virtual Machines

Microsoft has developed a thorough sizing methodology for Microsoft Lync Server (the predecessor of Skype for Business) that continues to mature with recent versions of Skype for Business Server 2015. VMware recommends that customers use the Skype for Business sizing tools mentioned earlier in this document. The amount of memory required for a Skype for Business Server 2015 is driven by the

expected amount of load that will be generated AND the total number of Skype for Business Server 2015 component servers that will be supporting the load in the environment. Load characteristics include (but are not limited to) the following:

- Projected maximum number of unique users simultaneously connected.
- Projected ratio of mobile-to-non-mobile point of presence users connected.
- Projected maximum number of video teleconferencing (VTC) systems in concurrent calls.
- Redundancy and resilience required to satisfy availability, recoverability, and business continuity SLAs.
- Corporate, environmental, and/or regulator backup, archiving, security, and overall hygiene requirements.

Since some Skype for Business Server 2015 components are memory-intensive and performance is a key factor in a typical Skype for Business production environment, VMware recommends that in addition to utilizing the sizing tools and following other prescriptive guidance from Microsoft, customers should adhere to the following practices when determining the optimum memory allocation for a Skype for Business Server virtual machine on the vSphere platform:

- Do not overcommit memory on ESXi hosts running Skype for Business Server 2015 workloads. If memory over-commitment cannot be avoided, use vSphere memory allocation options to guarantee the required memory size to the Skype for Business component server virtual machines, or to limit memory access for other, non-essential virtual machines in the vSphere cluster.
- For production systems, it is possible to achieve this objective by setting a memory reservation equal to the configured/allocated memory size of the Skype for Business virtual machine.

Note Setting memory reservations might limit vSphere vMotion. A virtual machine can be migrated only if the target ESXi host has free physical memory equal to or greater than the size of the reservation.

Setting the memory reservation to the configured size of the virtual machine results in a per-virtual machine VMkernel swap file of near-zero bytes that consumes less storage and eliminates ESXi host-level swapping. The guest operating system within the virtual machine still requires its own page file.

Reservations are recommended only when it is possible that memory might become overcommitted on hosts running Skype for Business Server 2015 virtual machines, when SLAs dictate that memory be guaranteed, or when there is a desire to reclaim space used by a virtual machine swap file.

There is a slight, appreciable performance benefit to enabling memory reservations even if you do not expect memory to be overcommitted in the vSphere cluster.

The benefits of memory reservations are not manifest if there is no memory over-provisioning or contention within a given vSphere cluster.

- Do not disable the balloon driver (which is installed with VMware Tools™) or any other ESXi memory management mechanism.

Note Setting memory reservations might limit vSphere vMotion. A virtual machine can be migrated only if the target ESXi host has free physical memory equal to or greater than the size of the reservation.

Transparent page sharing (TPS) enables ESXi hosts to more efficiently utilize its available physical memory to support more workloads. TPS is useful in scenarios where multiple virtual machine siblings share the same characteristics (such as OS and applications). In this configuration, vSphere is able to avoid redundancy by sharing similar pages among the different Skype for Business virtual machines. This sharing is transparent to the applications and processes inside the virtual machine.

For security reasons, inter-virtual machine page sharing is disabled by default on current versions of vSphere. While a virtual machine continues to benefit from TPS in this configuration (the virtual machine is able to share pages internally among its own processes and components), a greater benefit can be realized by enabling inter-virtual machine page sharing. See **Sharing Memory Across Virtual Machines** in the *vSphere Resource Management* guide.

- Enable DRS to balance workloads in the ESXi host cluster. DRS and reservations can give critical workloads the resources they require to operate optimally. More recommendations for using DRS with mission critical application virtual machines that require guaranteed memory resource access are available in Section 3, Using vSphere Technologies with Skype for Business Server 2015.

2.2.4 Memory Hot Add, Oversubscription, and Dynamic Memory

vSphere exposes the ability to add memory to a virtual machine while it is powered on. Modern operating systems (including Windows) support this feature and are able to instantaneously detect and use the hot-added memory.

While vSphere also enables the ability to hot remove memory from a virtual machine, this capability is not supported by current versions of Windows, so it is unavailable to a Windows virtual machine on vSphere.

vSphere enables an administrator to allocate virtual memory to virtual machines beyond the physically available memory size of the ESXi host. This condition is called over-allocation, or oversubscription. Oversubscription is possible and non-intrusive and is an essential core benefit of virtualization, because repeated research and testing efforts have shown that virtual machines do not fully utilize their allocated resources at the same time or on a regular basis. If all virtual machines request their resources at the same time at any point in time, this creates resource over-commitment on the ESXi or cluster.

Transient resource over-commitment is possible within a virtual environment. Frequent or sustained occurrence of such incidents is problematic for critical applications such as Skype for Business Server 2015.

Dynamic memory is a Hyper-V construct that does not have a direct equivalent on vSphere. Even in an over-commitment scenario, the virtual machine on vSphere is never informed that its allocated amount of memory has been physically reduced. vSphere uses other memory management techniques for arbitrating contentions during a resource over-commitment condition.

Skype for Business Server 2015 allocates memory shares to its processes at startup time, based on the available memory resources reported by the operating system. This allocation is not reevaluated until the server or the related Skype for Business services and processes have been restarted. With this behavior of memory pre-allocation, adding memory to a running Skype for Business Server 2015 virtual machine provides no additional benefit unless the virtual machine was rebooted or Skype for Business Server 2015 services were restarted. Consequently, memory hot add is neither useable by (nor beneficial to) a Skype for Business Server 2015 virtual machine and is therefore neither recommended nor supported.

In a Skype for Business infrastructure, the only exception to the preceding paragraph is the Microsoft SQL Server application supporting the environment. Microsoft SQL Server is capable of detecting and leveraging additional compute resource as they become available to the operating system. Although VMware strongly recommends that Skype for Business architects and administrators right-size all the

virtual machines they intend to use in a given Skype for Business infrastructure, it is important to point out that hot adding memory to a Microsoft SQL Server virtual machine is both supported and beneficial.

Although we are unaware of public guidance from Microsoft regarding memory hot add or memory oversubscription for virtualized Skype for Business Server 2015, historical Microsoft support statements for the virtualization of Microsoft Lync Server indicated that oversubscription and dynamic allocation of memory for such virtual machines is not recommended. To avoid confusion, refer to the preceding paragraphs to understand why these requirements are neither relevant nor applicable to Skype for Business Server 2015 workloads virtualized on the vSphere platform.

Oversubscription is different from over-commitment. Oversubscription is benign and does not impact the virtual machines. Over-commitment is the adverse extension of oversubscription and should be avoided in all cases. However, if you expect that your ESXi cluster might occasionally experience resource contention as a result of memory over-commitment, VMware recommends that you judiciously reserve all of the memory that you have allocated to your Skype for Business Server 2015 virtual machines.

Because ESXi does not support hot unplug (hot removal) of memory from a Windows virtual machine, the only way (outside of unsupported programmatic interference) to reduce the amount of memory presented to a virtual machine running Skype for Business Server 2015 is by powering off the virtual machine and changing the memory allocation. When the virtual machine powers on again, the OS will see the new memory size and Skype for Business Server 2015 will reallocate the available memory to its processes and threads. This is not dynamic memory.

2.3 Storage Virtualization

VMFS is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files, and VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre Channel, iSCSI, and network-attached storage (NAS) shared-storage protocols.

It is preferable to deploy virtual machine files on shared storage to take advantage of vSphere vMotion, vSphere HA, and DRS. This is considered a best practice for mission-critical application server deployments that are often installed on third-party, shared-storage management solutions.

VMware storage virtualization can be categorized into three pillars of storage technology, as illustrated in the following figure. The storage array is the physical storage pillar, comprising physical disks presented as logical storage volumes in the next pillar. Storage volumes, presented from physical storage, are formatted as VMFS datastores or with native file systems when mounted as raw device mappings (RDMs). Virtual machines consist of virtual disks or raw device mappings that are presented to the guest operating system as SCSI disks that can be partitioned and formatted using any supported file system.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Figure 4. VMware Storage Virtualization

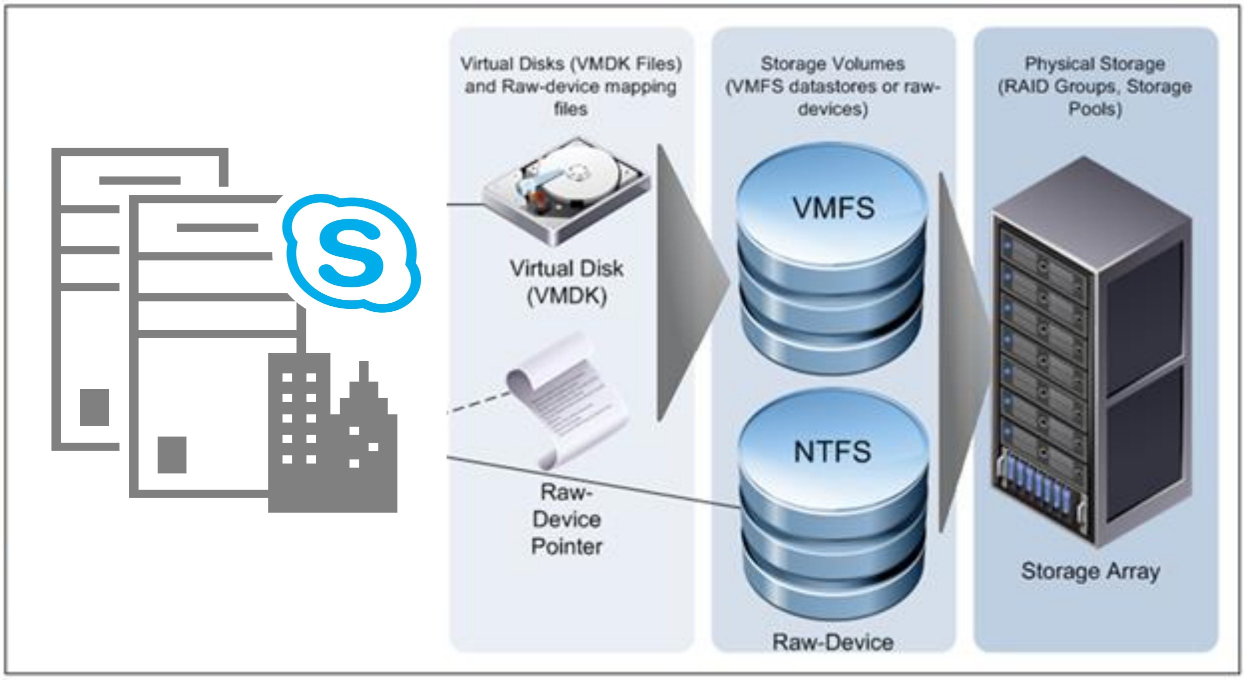

Skype for Business component servers do not typical require significant disk I/O capacity because they are typically not disk-bound. While it is tempting to conclude that this low IOPS requirement means that the disk IO needs of a typical Skype for Business component server can be satisfied with the IO throughput of commodity-class, direct-attached storage (DAS), customers should be aware that the associated operational and administrative overhead required to efficiently support DAS and JBODs in an enterprise-level production Skype for Business Server 2015 infrastructure might outweigh the potential financial savings that could be derived from such configuration options.

While VMware supports the use of converged storage solutions for virtualizing Skype for Business Server 2015 on vSphere, VMware recommends that customers using such solutions (such as VMware Virtual SAN™) thoroughly benchmark and validate the suitability of such solutions and engage directly with the applicable vendors for configuration, sizing, performance, and availability guidance.

Even with the reduced I/O requirements of a Skype for Business Server 2015 instance, without careful planning, sub-optimal storage access, availability, IO throughput, and latencies can still manifest in a Skype for Business Server 2015 infrastructure.

VMware recommends that you set up a minimum of four paths from an ESXi host to a storage array. To accomplish this, the host requires at least two host bus adapter (HBA) ports.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Figure 5. Storage Multi-Pathing Requirements for vSphere

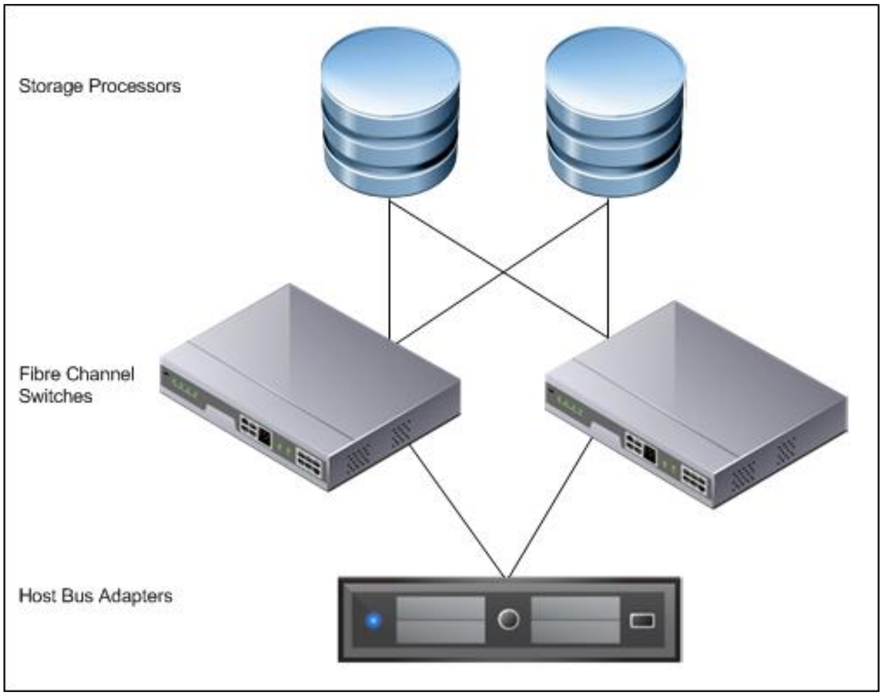

The terms used in the preceding figure are:

- HBA – A device that connects one or more peripheral units to a computer and manages data storage and I/O processing.
- Fibre Channel (FC) – A gigabit-speed networking technology used to build storage area networks (SANs) and to transmit data.
- Storage Processor (SP) – A SAN component that processes HBA requests routed through an FC switch and handles the RAID/volume functionality of the disk array.

2.3.1 Raw Device Mapping

ESXI also supports Raw Device Mapping (RDM). RDM allows a virtual machine to directly access a volume on the physical storage subsystem and can be used only with Fiber Channel or iSCSI. RDM provides a symbolic link or mount point from a VMFS volume to a raw volume. The mapping makes volumes appear as files in a VMFS volume. The mapping file, rather than the raw volume, is referenced in the virtual machine configuration. Connectivity from a Windows virtual machine to the raw volume is direct, and all data is stored using the native file system, NTFS. In the case of a failure of the VMFS datastore holding the RDM mapping file, a new mapping file can be created. Access to the raw volume and its data is restored, and no data loss occurs.

There is no performance difference between a VMFS datastore and an RDM. The following are the only conditions that impose a requirement for RDM disks for VMs in current versions of vSphere:

- If the backup solution performs hardware-based VSS snapshots of VMs (or otherwise requires direct storage access).

- When the virtual machine will be participating in a Windows Server Failover Clustering configuration that requires the clustered VMs to share the same disks (for example, Microsoft SQL Server AlwaysOn Failover Clustering Instances).

Because Skype for Business Server 2015 component servers do not required clustering (with the possible exception of the backend Microsoft SQL Server), the only scenario for the use of RDM disks for a virtualized Skype for Business component server on vSphere is when it is required by the backup solution vendor.

The decision to use VMFS or RDM for Skype for Business Server 2015 data should be based on technical requirements. The following table summarizes the considerations when making a decision between the two.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Table 2. VMFS and Raw Disk Mapping Considerations for Skype for Business Server 2015

VMFS	RDM
• Volume can contain many virtual machine disk files, reducing management overhead. • Increases storage utilization and provides better flexibility and easier administration and management. • Supports existing and future vSphere storage virtualization features. • Fully supports VMware Site Recovery Manager™. • Supports the use of vSphere vMotion, vSphere HA, and DRS. • Supports VMFS volumes (datastore) and virtual disks/VMDK files up to 62 TB each. • Supports up to 60TB of open VMDK storage per host (see http://kb.vmware.com/kb/1004424).	• Requires more LUNs, making it is easier to reach the limit of 255 LUNs that might be presented to an ESXi host. • Might be required to leverage array-level backup and replication tools (VSS) integrated with Skype for Business Server 2015 backend SQL Server databases. • Facilitates data migration between physical and virtual machines using the LUN swing method. • Fully supports Site Recovery Manager. • Supports vSphere vMotion (physical mode RDM first supported for vSphere vMotion in ESXi 6.0), vSphere HA, and DRS. • Supports presenting volumes of up to 64 TB to the guest operating system (physical-mode RDM only).

2.3.2 In-Guest iSCSI and Network-Attached Storage

Similar to RDM, in-guest iSCSI initiator-attached LUNs provide dedicated storage to a virtual machine. Storage presented using in-guest iSCSI is formatted natively using NTFS within the Windows guest operating system and bypasses the storage management of the ESXi host. Presenting storage in this way requires that additional attention be paid to the networking infrastructure and configuration at both the vSphere and the physical network levels. Some of the benefits from using in-guest iSCSI attached storage include the ability to allocate more than 256 LUNs to virtual machines in a single ESXi cluster, and retaining the ability to use array-level backup and replication tools.

When using in-guest iSCSI to present storage to a Skype for Business Server 2015 virtual machine, confirm that the iSCSI NIC is exempted from Windows Server Failover Clustering (for the backend SQL Server) or any other non-storage-related processes or components. Similarly, VMware recommends that customers use jumbo frames or 10 GB networks to support such storage configuration options.

2.3.3 Virtual SCSI Adapters

VMware provides two commonly used virtual SCSI adapters for Windows Server 2012 – LSI Logic SAS and VMware Paravirtual SCSI (PVSCSI). The default adapter when creating new virtual machines running Windows on vSphere is LSI Logic SAS.

The VMware Paravirtual SCSI adapter is a high-performance vSCSI adapter developed by VMware to provide optimal performance for virtualized business critical applications. The advantage of the PVSCSI adapter is that the added performance is delivered while minimizing the use of hypervisor CPU resources. This leads to lower hypervisor overhead required to run storage I/O-intensive applications.

As previously discussed, optimal storage I/O throughput and reduced latencies are important considerations in large and geographically-dispersed Skype for Business environments. The PVSCSI controller not only provides these capabilities, but it is also the most current SCSI controller supported in a vSphere infrastructure.

Virtual machines can be deployed with up to four virtual SCSI adapters. Each vSCSI adapter can accommodate up to 15 storage devices for a total of 60 storage devices per virtual machine. During allocation of storage devices, each device is assigned to a vSCSI adapter in sequential order. Not until a vSCSI adapter reaches its fifteenth device will a new vSCSI adapter be created. To provide better parallelism of storage I/O, VMware recommends that customers use as many of these controllers as necessary and distribute storage devices (VMDKs or RDMs) among the controllers in use. See the following diagram for a high-level illustration of this recommendation:

Figure 6. Storage Distribution with Multiple vSCSI Adapters

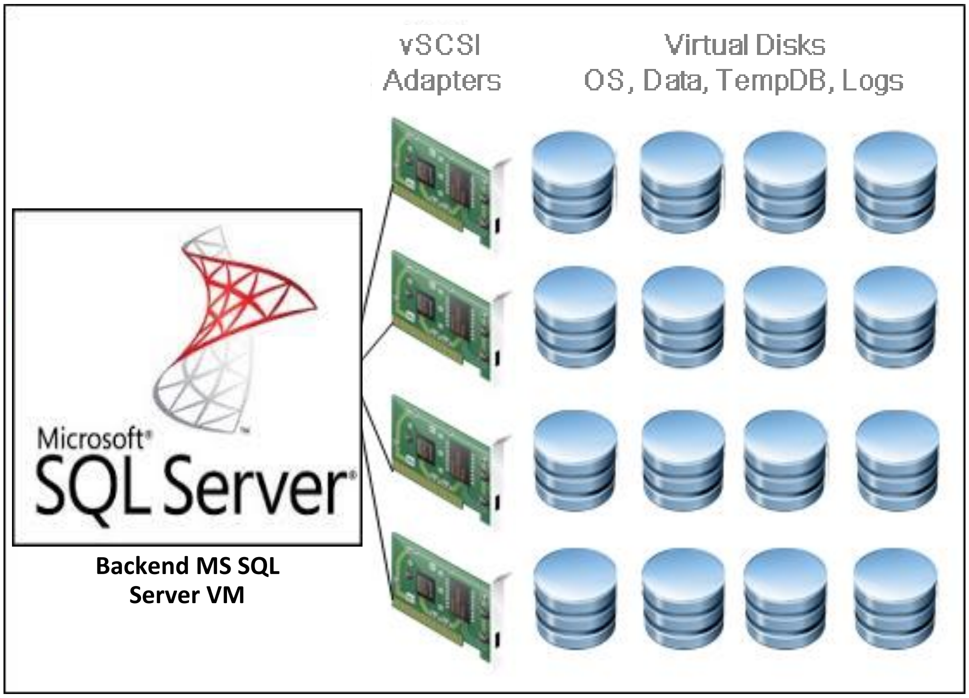

2.3.4 Virtual SCSI Queue Depth

Queue depth denotes the number of I/Os that can pass through a storage path at one time. All other I/Os beyond this number are queued until the path has more room to accommodate them. Because there are multiple points through which an I/O will traverse from within a guest operating system before it reaches the physical storage device, customers should pay special attention to their storage configuration to avoid unnecessary and undesirable I/O queuing for their Skype for Business Server 2015 virtual machines.

When presenting storage to a virtual machine, a virtual disk is connected to one virtual SCSI controller. Each SCSI controller has a finite (and configurable) queue depth, which varies among the virtual SCSI controller types available in vSphere.

The PVSCSI controller is the optimal SCSI controller for an I/O-intensive application on vSphere. This controller has a queue depth of 64 (per device) and 254 (per controller) by default (double the size of an LSI Logic SAS controller). The PVSCSI controller's per-device and per-controller queue depths can also

be increased to 254 and 1024 respectively, providing even more increased I/O bandwidth for the virtualized workload.

Because of these increased performance and throughput benefits, VMware recommends that customers choose PVSCSI as the virtual SCSI controllers, at least for their Skype for Business component servers' data volumes. See the preceding diagram for VMware recommendations for the backend Microsoft SQL Server's virtual SCSI and virtual disks.

Note The PVSCSI driver is not native to the Windows operating system. Customers using PVSCSI controllers must update the VMware Tools instances on their virtual machines on a regular basis. VMware Tools is the vehicle for delivering the PVSCSI drivers to the OS.

Figure 7. Common Points of Storage I/O Queues

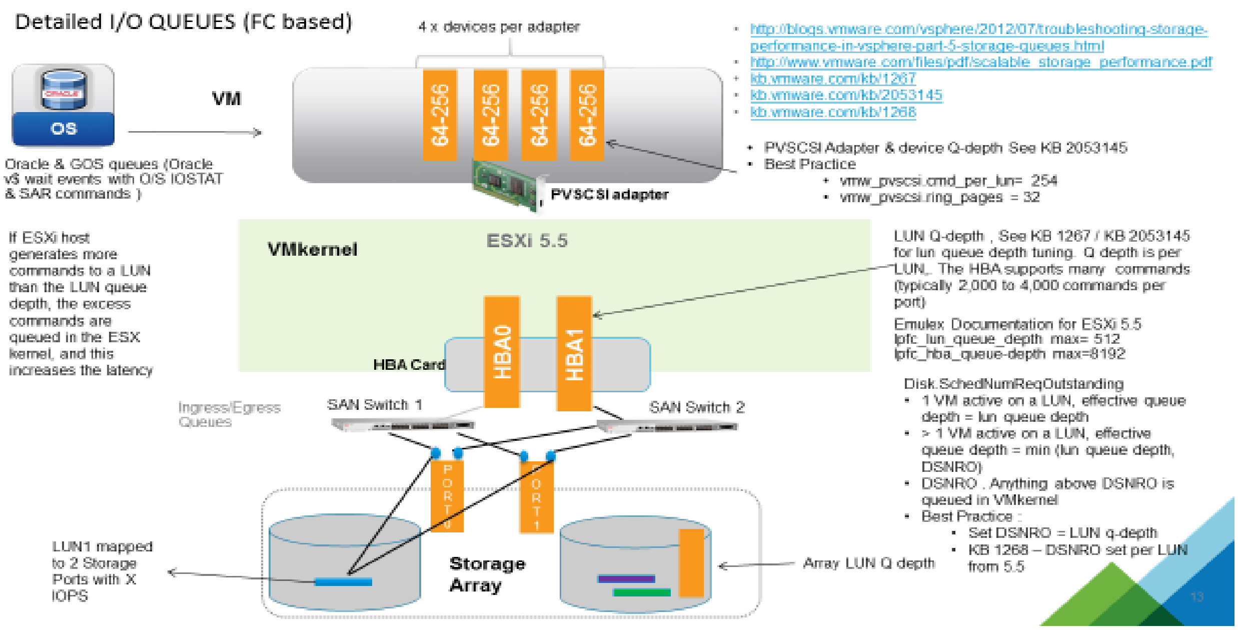

While increasing the default queue depth of a virtual SCSI controller can be beneficial to a Skype for Business Server 2015 virtual machine, the configuration can also introduce unintended adverse effects in overall performance if not done properly. VMware highly recommends that customers consult and work with their storage vendor's support personnel to evaluate the impact of such changes and obtain recommendations for other adjustments that might be required to support the increase in queue depth of a virtual SCSI controller. See the following references for further information:

- **Large-scale workloads with intensive I/O patterns might require queue depths significantly greater than Paravirtual SCSI default values**
- **Changing the queue depth for QLogic, Emulex, and Brocade HBAs**

2.3.5 Skype for Business Server 2015 on All-Flash Storage Array

Since all-flash storage is gaining popularity in corporate data centers, many people are examining the value of deploying Skype for Business Server 2015 on all-flash storage arrays. Skype for Business Server 2015 is generally perceived as a workload that requires lots of capacity, but is relatively less demanding on performance when compared to other Mission-critical application workloads. All-flash storage arrays are viewed as more expensive than traditional storage arrays when it comes to physical capacity cost per

GB. Why would customers be interested in an all-flash storage option for Skype for Business Server 2015?

While people are well aware of the performance benefits of all-flash storage, the latest generation of all-flash storage also offers:

- Built-in data services, such as always-on thin provisioning, inline data deduplication, and inline data compression that provide compelling data reduction ratios.
- Flash-optimized data protection that replaces traditional RAID methodologies and can simplify the Skype for Business Server 2015 sizing and capacity planning efforts while minimizing data protection overhead and performance penalty.
- Instant space-efficient copies through VSS integration that significantly increases efficiency and operational agility and can be used for local data protection, significantly the reducing recovery time objective and administrative efforts required in a failure or data corruption scenario.

Microsoft recommends and supports the use of flash-based storage arrays for Skype for Business Server 2015 workloads. VMware recommends that customers consult their storage array vendors for the most applicable tuning and optimization options available to them when virtualizing their Skype for Business workloads on vSphere.

2.4 Networking Configuration Guidelines

This section covers design guidelines for the virtual networking environment and provides configuration examples at the ESXi host level for Skype for Business Server 2015 installations.

Note The examples do not reflect design requirements and do not cover all possible Skype for Business Server 2015 network design scenarios.

2.4.1 Virtual Networking Concepts

The virtual networking layer consists of the virtual network devices through which virtual machines and the ESXi hosts interface with the rest of the network and users. In addition, ESXi hosts use the virtual networking layer to communicate with iSCSI and NAS storage devices.

The virtual networking layer includes virtual network adapters and virtual switches. Virtual switches are the key networking components in vSphere. The following figure provides an overview of virtual networking in vSphere.

Figure 8. vSphere Virtual Networking Overview

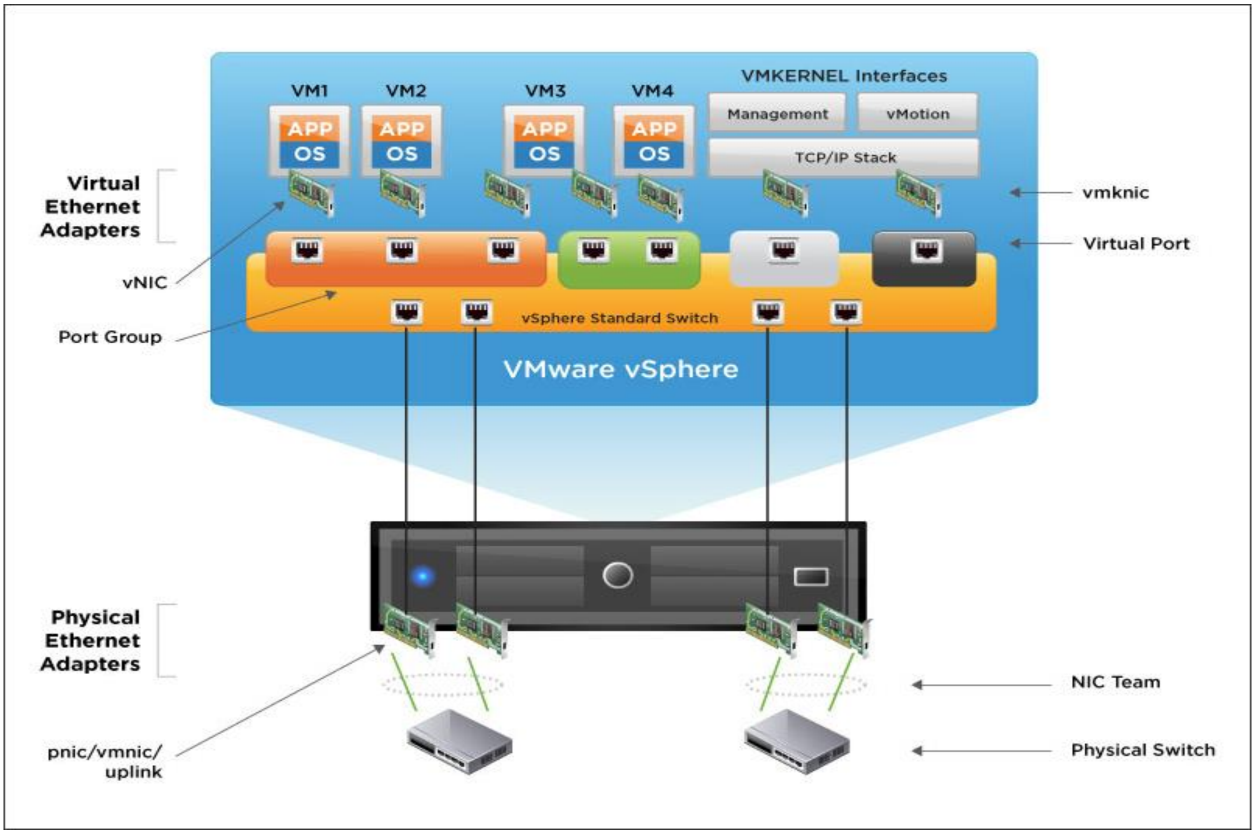

As shown in the preceding figure, the following components make up the virtual network:

- Physical switch – vSphere host-facing edge of the physical local area network.
- NIC team – Group of physical NICs connected to the same physical/logical networks to provide redundancy.
- Physical network interface (pnic/vmnic/uplink) – Provides connectivity between the ESXi host and the local area network.
- vSphere switch (standard and distributed) – The virtual switch is created in software and provides connectivity between virtual machines. Virtual switches must uplink to a physical NIC (also known as vmnic) to provide virtual machines with connectivity to the LAN, otherwise virtual machine traffic is contained within the virtual switch.
- Port group – Used to create a logical boundary within a virtual switch. This boundary can provide VLAN segmentation when 802.1q trunking is passed from the physical switch, or it can create a boundary for policy settings.
- Virtual NIC (vNIC) – Provides connectivity between the virtual machine and the virtual switch.
- VMkernel (vmknic) – Interface for hypervisor functions, such as connectivity for NFS, iSCSI, vSphere vMotion, and VMware vSphere Fault Tolerance logging.
- Virtual port – Provides connectivity between a vmknic and a virtual switch.

2.4.2 Virtual Networking Best Practices

The following standard VMware networking best practices are applicable to Skype for Business Server 2015 workloads running on the vSphere platform:

- The choice between standard and distributed switches should be made outside of the Skype for Business Server 2015 design. Standard switches provide a straightforward configuration on a per-host level. For reduced management overhead, improved cross ESXi communication, and increased functionality, the distributed virtual switch should be used. Both virtual switch types provide the functionality needed by Skype for Business Server 2015.
- Traffic types should be separated to keep like traffic contained to designated networks. vSphere can use separate interfaces for management, vSphere vMotion, and network-based storage traffic. Additional interfaces can be used for virtual machine traffic. Within virtual machines, different interfaces can be used to keep certain traffic separated. Use 802.1q VLAN tagging and virtual switch port groups to logically separate traffic. Use separate physical interfaces and dedicated port groups or virtual switches to physically separate traffic. This is shown in 8 above.
- Leverage network interface teaming capabilities to provide redundant uplinks for virtual switches. To use this capability, assign at least two physical network interfaces per virtual switch.
- Use the VMXNET3 network adapter. This is a paravirtualized network device that provides better network I/O throughput and enhanced features and capabilities, combined with reduced hypervisor CPU utilization. It is imperative that customers regularly update VMware Tools on virtual machines using VMXNet3 virtual adapters.
- Where possible, VMware highly recommends that customers consider enabling jumbo frames on the virtual switch to which the vSphere vMotion VMkernel port group is connected. This is to improve vSphere vMotion operations and increase workload availability and recovery for the Skype for Business virtual machines.

NOTE: Before enabling jumbo frames within the vSphere infrastructure, confirm that jumbo frames are also enabled on your physical network infrastructure.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Figure 9. Avoid MTU Mismatches in Network Traffic Paths

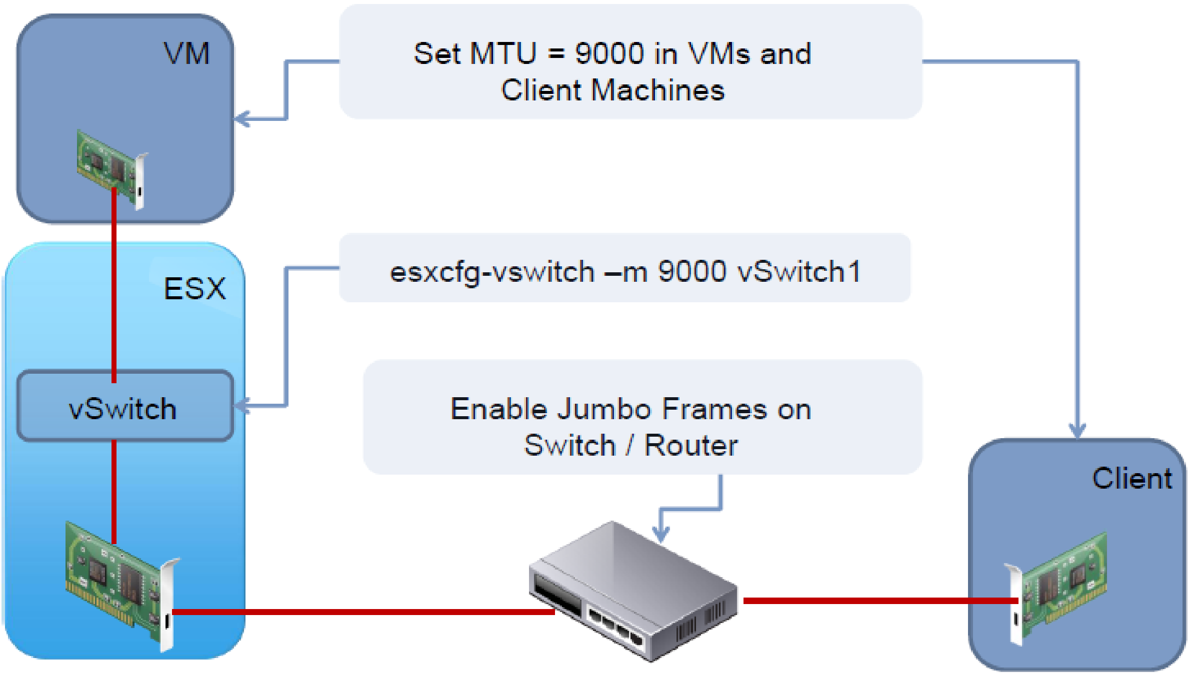

Follow the guidelines in the **Guest Operating System Networking** and **Hardware Networking Considerations** sections of the *Performance Best Practices for VMware vSphere 6.0* Guide.

2.4.3 Sample Skype for Business Server 2015 Virtual Network Configuration

Because of the flexibility of virtual networking, the topology can take many different forms. There is no single recommended practice because each provides its own sets of benefits. The following figure shows two examples of host-level configurations based on the most common configurations, which are single and multiple virtual switches.

Figure 10. Sample Virtual Network Configuration

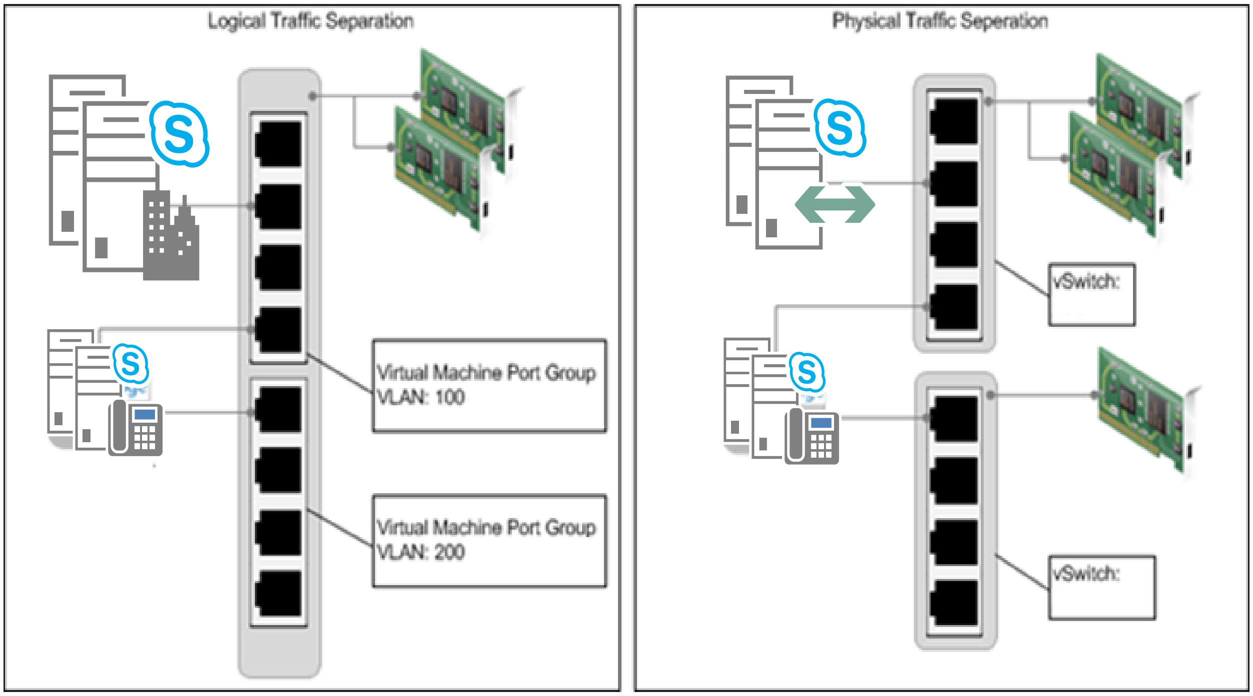

In the vSphere environment, traffic separation can be established using virtual or physical networks. The preceding figure provides examples of these two scenarios.

- The scenario on the left depicts an ESXi host with two network interfaces, teamed for redundancy and using virtual networks and port groups to provide the traffic separation for different Skype for Business component servers' traffic. This scenario can also utilize VMware vSphere Network I/O Control for dynamic traffic prioritization.
- The scenario on the right depicts an ESXi host with multiple network interfaces. Physical traffic separation is accomplished by allocating two vmnics on one network to a virtual switch. These vmnics are teamed and dedicated to client access network traffic.

In both scenarios, the Skype for Business Server virtual machine's network requirements can be accommodated without constraints. The flexibility allows vSphere administrators to determine the most applicable design and configuration options that best meet corporate, environmental, security, availability, or vendor-imposed requirements.

2.5 Power Management

Modern server hardware and operating systems are engineered to minimize power consumption by default. Windows and the vSphere ESXi hypervisor both favor reduced power consumption over performance. While previous versions of ESXi default to high performance power schemes, vSphere 6.0 defaults a balanced power scheme. For critical applications such as Skype for Business Server 2015, the default power scheme in vSphere 6.0 is not recommended.

The vSphere ESXi hypervisor provides a high performance and competitive platform that effectively runs many Tier 1 application workloads in virtual machines. By default, ESXi has been heavily tuned for driving high I/O throughput efficiently by utilizing fewer CPU cycles and conserving power, as required by a wide

range of workloads. However, many applications require I/O latency to be minimized, even at the expense of higher CPU utilization and greater power consumption.

VMware defines latency-sensitive applications as workloads that require optimizing for a few microseconds to a few tens of microseconds end-to-end latencies. Eliminating (or, at least, minimizing) delays and jitters are of paramount importance for these applications. Skype for Business Server 2015 is an example of such an application because enterprises consider the quality of the video, voice, and other services provided by this application to be of more importance than the need to save cost through a reduction in power consumption.

Configuring the vSphere hypervisor to achieve such performance objective is effortless, as shown in the following image.

Figure 11. Default ESXi 6.0 Power Management Settings

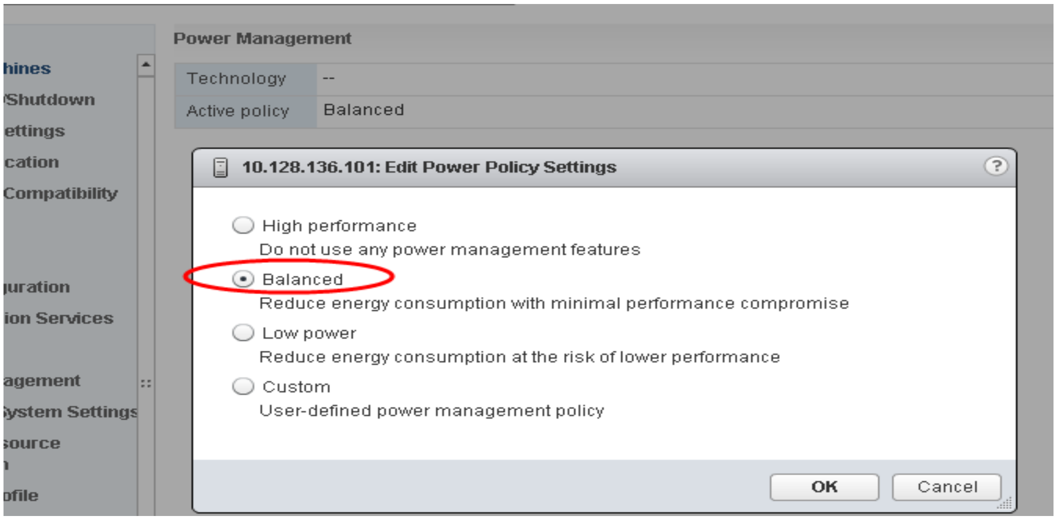

There are three distinct areas of power management in a vSphere hypervisor virtual environment:

- Server hardware
- Hypervisor
- Guest operating system

The following section provides power management and power settings recommendations that cover all of these areas.

2.5.1 Server Hardware BIOS Settings

Most servers with new Intel and AMD processors provide power saving features that use several techniques to dynamically detect the load on a system and put various components of the server, including the CPU, chipsets, and peripheral devices, into low power states when the system is mostly idle. Hardware-level power management adds latency to the path where an idle system (in one of several power savings modes) responds to an external event. Consequently, VMware recommends the following BIOS-level power management settings for Skype for Business Server 2015 on vSphere:

- Set **Power Management** (or its vendor-specific equivalent label) to OS controlled. This will enable the ESXi hypervisor to make the most judicious power management decisions based on actual utilization and the condition of the running virtual machines.

- Disable all processor C-states (including the C1E halt State). These enhanced power management schemes can introduce memory latency and sub-optimal CPU state changes (Halt-to-Full), resulting in reduced performance for the virtual machine.
- Enable Turbo Boost, if available.

2.5.2 ESXi Host Power Settings

ESXi can take advantage of several power management features that the host hardware provides to adjust the trade-off between performance and power use. You can control how ESXi uses these features by selecting a power management policy.

In general, selecting a high performance policy provides more absolute performance, but at lower efficiency (performance per watt). Lower-power policies provide lower absolute performance, but at higher efficiency. ESXi provides five power management policies. If the host does not support power management, or if the BIOS settings specify that the host operating system is not allowed to manage power, only the Not Supported policy is available.

Table 3. CPU Power Management Policies

Power Management Policy	Description
High Performance	The VMkernel detects certain power management features, but will not use them unless the BIOS requests them for power capping or thermal events. This is the recommended power policy for a Skype for Business Server 2015 running on ESXi.
Balanced (Default)	The VMkernel uses the available power management features conservatively to reduce host energy consumption with minimal compromise to performance.
Low Power	The VMkernel aggressively uses available power management features to reduce host energy consumption at the risk of lower performance.
Custom	The VMkernel bases its power management policy on the values of several advanced configuration parameters. You can set these parameters in the vSphere Web Client **Advanced Settings** dialog box.
Not Supported	The host does not support any power management features or power management is not enabled in the BIOS.

VMware recommends that customers choose the High Performance power policy option for all ESXi hosts in a vSphere cluster supporting a Skype for Business infrastructure. You select a policy for a host using the vSphere Web Client. If you do not select a policy, ESXi uses **Balanced** by default.

See **Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs** for a more detailed discussion of the benefits of this recommendation.

Figure 12. Recommended ESXi Host Power Management Settings

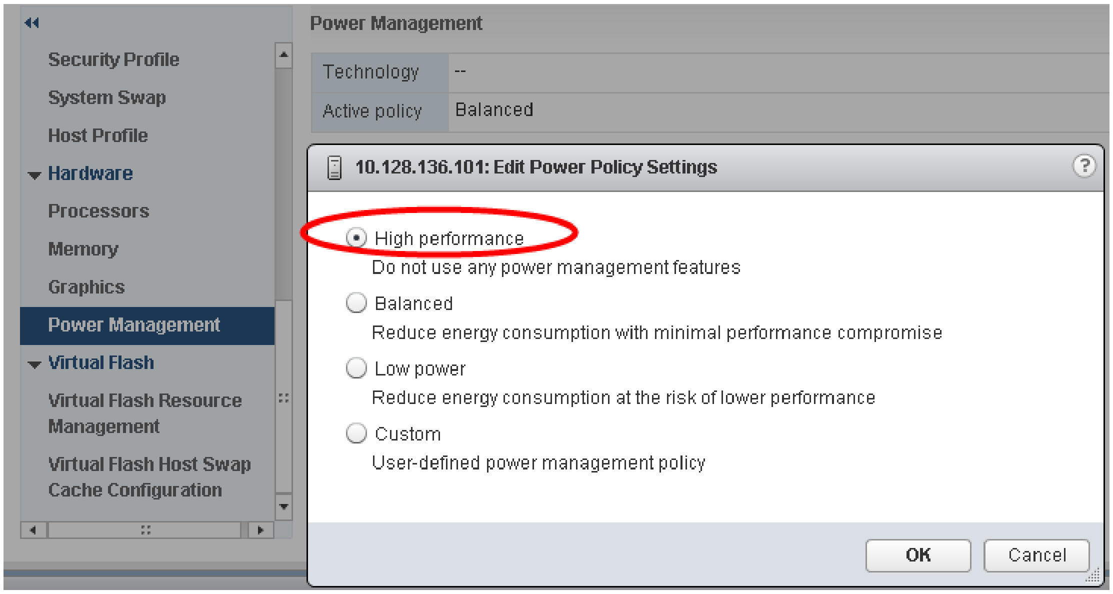

When a CPU runs at lower frequency, it can also run at lower voltage, which saves power. This type of power management is called dynamic voltage and frequency scaling (DVFS). ESXi attempts to adjust CPU frequencies so that virtual machine performance is not affected.

When a CPU is idle, ESXi can take advantage of deep halt states (known as C-states). The deeper the C-state, the less power the CPU uses, but the longer it takes for the CPU to resume running. When a CPU becomes idle, ESXi applies an algorithm to predict how long it will be in an idle state and chooses an appropriate C-state to enter. In power management policies that do not use deep C-states, ESXi uses only the shallowest halt state (C1) for idle CPUs.

VMware, therefore, recommends disabling all C-state settings at the hardware level for all ESXi servers hosting Skype for Business workloads.

2.5.3 Windows Guest Power Settings

The default power policy option in Windows Server 2012 is Balanced. This configuration allows the Windows OS to save power consumption by periodically throttling power to the CPU and turning off devices such as the network cards in the guest when Windows determines that they are idle or unused. This capability is inefficient for a Skype for Business Server 2015 workload due to the latency and disruption introduced by the act of powering-off and powering-on CPUs and devices. Allowing Windows to throttle power to CPUs results in what describes as *core parking.* See the following Microsoft blog posts for more discussion of core parking behavior:

- **Processor power policy on Windows**
- **Processor Power Management Options**
- **Troubleshooting High CPU utilization issues in Exchange 2013** ("Power Management" section)

Figure 13. Windows CPU Core Parking

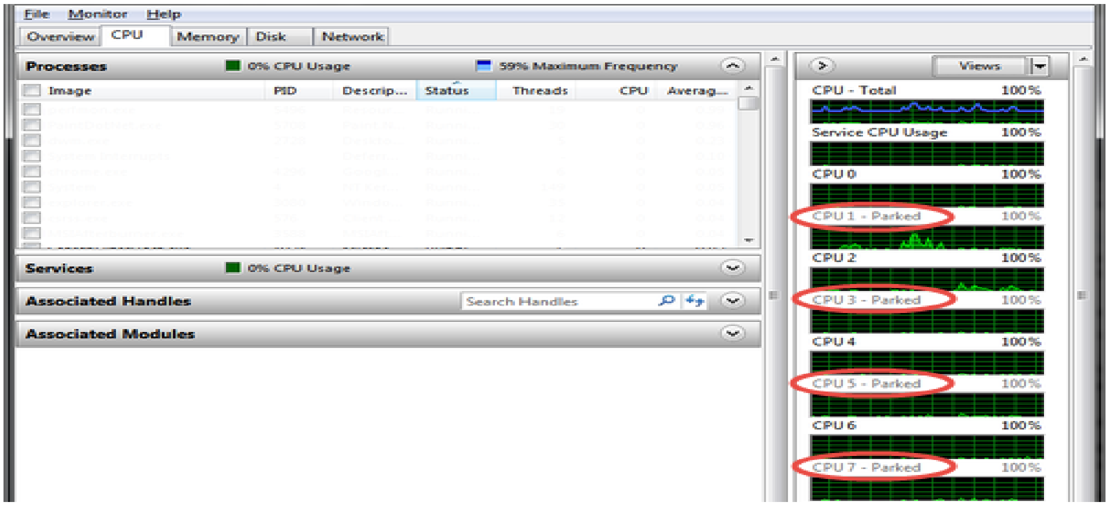

Microsoft recommends the High Performance power management policy for applications requiring stability and performance. VMware supports this recommendation and encourages customers to incorporate it into their server tuning and administration practices for virtualized Skype for Business Server 2015 virtual machines.

Note Although the recommended power plan can help avoid the performance issues associated with core parking, Microsoft documentation indicates that this might not be sufficient. See **Tweaking CPU Core Parking** for additional recommendations.

Although Microsoft documentation also indicates that core parking is disabled by default on Windows Server 2012 (https://support.microsoft.com/en-us/kb/2814791), VMware encourages customers to manually verify their Skype for Business servers for any indication of core parking behavior.

Figure 14. Recommended Windows Guest Power Scheme

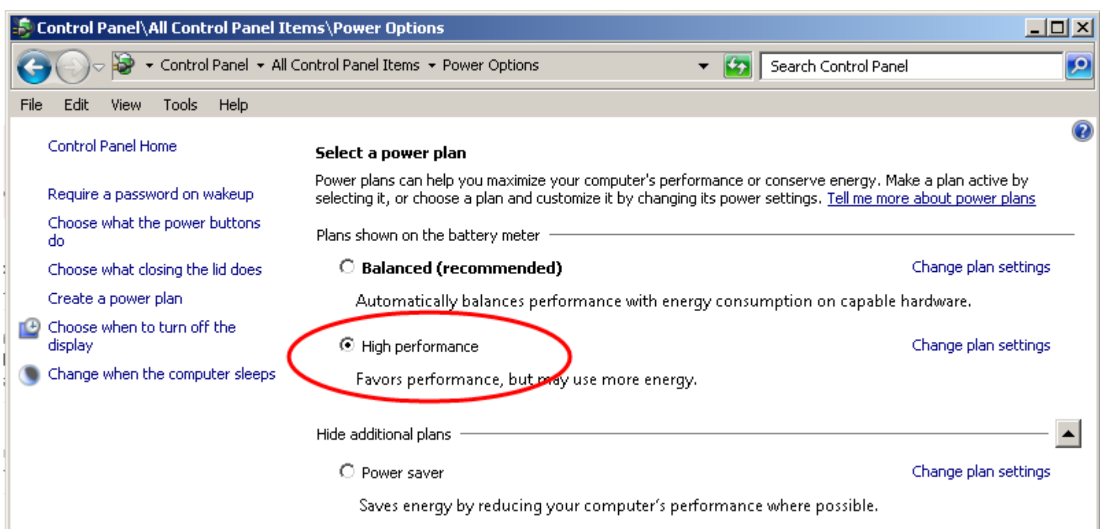

3. Using vSphere Technologies with Skype for Business Server 2015

This section explores the native features and configuration options available for building a highly available Skype for Business Server 2015 infrastructure in a vSphere cluster. This section will also discuss the recommended configuration practices and tuning options for using these features concurrently with the native high-availability and resilience features of Skype for Business Server 2015.

3.1 Overview of vSphere Technologies

The native high-availability and automated resource management features of vSphere consist of the following three major components:

- vSphere HA - vSphere HA provides high availability for virtual machines by pooling them and the hosts on which they reside into a cluster. Each ESXi host in a vSphere cluster is continually monitored and, in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts in the cluster.
- Distributed Resource Scheduler (DRS) – DRS dynamically allocates and balances computing capacity and virtual machine placement with resources pooled from multiple ESXI server hosts. DRS attempts to provide even and optimal resource utilization among all the ESXi hosts in a given vSphere cluster, avoiding the possibility that a host might be under resource utilization stress while other hosts in the cluster remain unused, or underutilized.
- vSphere vMotion (virtual machine and storage) - vSphere vMotion enables an administrator (or any authorized operator) to move a virtual machine and/or its associated disks and data from one ESXi host to another ESXi host (or one backend storage to another) while the virtual machine is running and providing services, without the need to power off the virtual machine or interrupt the services provided. The relocation of the virtual machine is transparent to the guest operating system, the applications and processes running inside the virtual machine, and especially the end users or other processes accessing the services provided by the virtual machine.

3.1.1 vSphere HA

With vSphere HA, Skype for Business Server 2015 virtual machines on a failed ESXi host can be restarted on another ESXi host. In addition to minimizing the downtime of a given virtual machine and its services, this feature provides a cost-effective proactive failover alternative to third-party clustering and replication solutions.

If you use vSphere HA, be aware of the following:

- vSphere HA handles ESXi host hardware failure and does not monitor the health status of the Skype for Business Server 2015 services. Although vSphere HA can monitor for heartbeat loss from a virtual machine, it is possible for a virtual machine to continue to provide heartbeat signals to vSphere while the applications and processes hosted on the virtual machine are unavailable. Customers should not consider vSphere HA as a Skype for Business monitoring solution.
- A vSphere HA heartbeat is sent using the vSphere VMkernel network, so optimal uplink bandwidth and redundancy in this network are strongly recommended.
- Allowing more than one Skype for Business servers hosting similar components and services to run on the same ESXi host for an extended period of time is not recommended. This condition will create a single-point-of-failure scenario if the host were to become suddenly unavailable. DRS anti-affinity or guest-to-host affinity rules should be used to avoid the possibility of this condition occurring.

- Microsoft has specific recommendations for using vSphere HA with Microsoft application servers such as Skype for Business. Microsoft requires that all host-level availability and recovery features must recover such application servers from a completely powered-off state. vSphere HA always recovers a virtual machine from a powered-off state.

3.1.2 vSphere vMotion

A well-designed and purpose-built vSphere infrastructure can provide seamless migration of running workloads with no special configuration. Even under ideal conditions, several factors can interfere with a vSphere vMotion operation, leading to undesirable outcomes for the virtualized workloads. It is, therefore, important for Skype for Business and vSphere architects, administrators, and operators to understand vSphere vMotion operations and the considerations required for providing its optimal performance.

The following is a very high-level description of a vSphere vMotion operation:

- Upon initiation, vSphere vMotion triggers the creation of a *shadow VM* (a representation of the original virtual machine) at the target ESXi host.
- A pre-copy operation begins copying the memory state of the virtual machine from its current host to its target host over the vSphere vMotion VMkernel network interface.
- The memory copy operation is performed iteratively to copy any memory pages that might have changed during the last copy.
- When the last changed pages become small enough to be copied in less than 500ms, the virtual machine is stunned (frozen) on the source host and resumed on the target host. A switch-over operation is performed.
    - Execution switch-over requires the virtual machine execution to be stopped on the source host. Then the virtual device state of the virtual machine is check-pointed on the source host and transmitted to destination host. After the received checkpoint is restored and all remaining changed pages are received, the virtual machine begins execution on the destination host.
- The old incarnation of the virtual machine is then transparently shut down and its process is deleted on the target host.

3.1.3 DRS

DRS takes vSphere vMotion a step further by adding an intelligent scheduler. DRS allows you to set resource assignment policies that reflect business needs. DRS does the calculations and automatically handles the details of physical resource assignments. It dynamically monitors the workloads of the running virtual machines and the resource utilization of the physical servers within a vSphere cluster.

vSphere vMotion and DRS perform best under the following conditions:

- The source and target ESXi hosts must be connected to the same network and the same shared storage.
- A dedicated gigabit (or higher) network for vSphere vMotion is recommended.
- The vSphere vMotion or DRS target host must have enough resources to accommodate the virtual machine that will be migrated.
- The virtual machine must not have physical devices such as a CD-ROM or floppy disk attached during the vSphere vMotion operation.
- The source and destination hosts must have compatible CPU models, otherwise migration with vSphere vMotion fails. If the vSphere cluster contains hosts with different CPU generations, then Enhance vMotion Compatibility (EVC) must be enabled in the cluster to allow vSphere vMotion

operations to succeed. See *EVC and CPU Compatibility FAQ* (http://kb.vmware.com/kb/1005764) for more information on EVC.

- Virtual machines with smaller memory sizes are better candidates for migration than larger ones.
- Persistent resource over-commitment in a vSphere cluster can impede the efficiency of vSphere vMotion operations

3.2 vSphere High Availability

By providing a higher level of availability than is possible out-of-the-box for most applications, vSphere HA has become the default HA solution for vSphere virtual machines. Regardless of operating system or application, vSphere HA can provide protection from ESXi host failures, guest operating system failures, and, with the help of third-party add-ons, application failures.

Like most enterprise applications, Microsoft Skype for Business Server 2015 includes native features and configuration options to enhance service and component availability. The default recommended High Availability option for Skype for Business component servers (with the exception of the backend Microsoft SQL Server) is to have multiple instances of those component (role) servers and configure them in a topology Microsoft describes as a "Pool"

The goal of High Availability is primarily the need to ensure uninterrupted and continuous availability of the services or features provided by an application, or (when a failure occurs) minimize the duration and impact of such outage. By deploying multiple copies (instances) of a given Skype for Business role, your Skype for Business infrastructure become more resilient to a single (or, when properly designed, multiple) Server outages. This is achievable simply due to the fact that there are multiple copies (pools) of the roles or components (e.g. Front End, Edge, Directors, etc) deployed in a single Skype for Business infrastructure, creating both server and role redundancy as a result.

VMware provides robust and very matured native redundancy and High Availability capability for workloads (virtual machines) running in a vSphere environment. vSphere HA is the mechanism for this capability.

Within a vSphere infrastructure, a hardware failure results in virtual machines being powered back on by vSphere HA, restoring availability levels quickly and keeping utilization balanced. With increased availability, customers are able to reduce the duration of a Skype for Business Server outage by ensuring that when a role server (VM) becomes unavailable (e.g. during a host failure), the VM is automatically recovered and powered on another ESXi host in the vSphere cluster. This rapid recovery compliments and enhances the native HA capability of Skype for Business Server. This concurrent use of Host and VM HA features is fully supported by Microsoft due to the fact that vSphere HA always recovers a VM from powered-off state.

Without the complimentary capabilities of vSphere HA, a failed Skype for Business role server will remain in such failed state (reducing the number of such roles in the pool), requiring manual administrative intervention.

The following sections provide recommendations for using vSphere HA with Skype for Business Server 2015.

3.2.1 Admission Control

The amount of failover capacity provided by hosts in a vSphere cluster determines how many host failures can be tolerated. For example, in a four-host cluster, each host needs to reserve at least 25 percent of its resources to accommodate the workloads of one host in the event of a failure. Reserving this capacity is a manual task without admission control.

vSphere HA's Admission Control allows administrators to define policies that are used by VMware vCenter® to manage failover capacity. How vCenter failover capacity is calculated and reserved is determined by the following policies:

- The number of host failures a cluster can tolerate.
- A percentage of the cluster resources that are reserved as failover spare capacity.
- Dedicated failover hosts.
- Sizes of the virtual machines in the cluster.

Surviving a physical host failure without compromising performance or availability is a major driver for virtualizing business critical applications. By configuring admission control, vCenter Server monitors utilization and manages failover capacity. To provide sufficient resources in the case of a hardware failure, or even during maintenance, enable and configure admission control.

For more information on admission control and the policies see "vSphere HA Admission Control" in the *vSphere Availability Guide* (http://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-601-availability-guide.pdf).

3.2.2 Virtual Machine Monitoring

Along with ESXi host monitoring, vSphere HA can also provide monitoring at the virtual machine level. Virtual machine monitoring can detect guest operating system failures and, with the help of third-party software, application failures. vSphere HA establishes a heartbeat with VMware Tools, installed within the guest operating system, and monitors this heartbeat. If heartbeat communication fails between VMware Tools and vSphere HA, a secondary check is made against network and storage I/O activity. If both network and storage I/O activity has halted, vSphere HA triggers a virtual machine restart.

Enable virtual machine monitoring if there is a desire to reduce downtime due to guest operating system failures that would otherwise require manual intervention. This monitoring feature cannot detect Skype for Business Server 2015 processes, health, or failures. It is, therefore, not a replacement for Microsoft's recommended Skype for Business Server 2015 HA configuration options.

3.2.3 Using vSphere HA with Skype for Business Pools

In a physical environment, it is not uncommon to deploy three or more Skype for Business role servers in a pool to protect from planned or unplanned outages. This helps ensure that the environment continues to provide the required services in the event of a failure of one or more role servers in the pool. The main rationale behind this large pool design is the realization that, in a 2-member pool configuration, the failure of a single role server in the pool exposes the entire pool to possible complete outage, for example, if the surviving role server were to encounter a failure before the first failed node has been repaired, recovered and introduced to the pool.

Large pool configuration increases availability, but at the cost of additional licensing cost, administrative efforts and possible sprawl.

Skype for Business Server 2015 environments built on vSphere are typically designed with two copies of the same role servers for availability purposes. vSphere HA is then added to protect these servers from planned and unplanned outages. vSphere HA restarts a pool member if the host where it was running experiences a hardware failure or becomes otherwise unavailable.

When enabling a vSphere cluster for HA with the intention of protecting Skype for Business Pool members, consider the following:

- Members of the same Pool should not reside on the same vSphere host for an extended period of time, and should be symmetrically distributed between among as many hosts in the vSphere cluster as is required to avoid a co-location scenario for VMs performing similar roles..
- Allowing two pool members to run on the same host for a short period of time (for instance, after a vSphere HA event), is acceptable, even if doing so might violate resource availability constraints and DRS rules. This allows the Skype for Business Server 2015 virtual machine to become operational and for the pool to become more resilient, even at reduced performance. Pool members should be separated as soon as operationally feasible (for example, as soon as the ESXi host becomes available or additional capacity has been added to the vSphere cluster).
- To adequately protect from an extended server outage, vSphere clusters should be designed in an N+1 configuration, where N is the largest number of Skype for Business role members in a given pool. If a hardware failure occurs causing vSphere HA to power on a failed pool member, Skype for Business Server 2015 maintains the same levels of performance and protection as during normal runtime because the “+1” host’s resources become immediately available for use in the environment.
- Use anti-affinity rules to keep Pool members separated. vSphere HA might violate this rule during a power-on operation (one caused by a host failure), but DRS fixes the violation during the next interval. To eliminate the possibility of Pool members running on the same host (even for a short period), **“must not run on host”** anti-affinity rules must be used.

3.3 vSphere vMotion and Skype for Business

Note The information provided in this section applies solely to virtual machines running the following Skype for Business Server components/roles:

- Frontend Server
- Mediation Server
- Edge Server
- Persistent Chat Server

Please see section "8.3 vMotion Considerations" in the **Microsoft SQL Server on VMware vSphere Availability and Recovery Options** guide for a discussion of vSphere vMotion impact on Microsoft SQL Server.

3.3.1 vSphere vMotion Failure Scenarios

From our high-level description of a vSphere vMotion operation workflow in Section 3.1.2, vSphere vMotion, the success of a vSphere vMotion operation depends on several factors, including the following:

- Available network bandwidth – Since the virtual machine memory pages are transferred over the network, a vSphere vMotion operation completes more quickly on faster network connections.
- Optimal storage throughput – vSphere vMotion needs to be able to allocate storage blocks for the virtual machine. Factors such as I/O congestion or prolonged disk/device queue depth saturation (over 20 seconds) can impair this storage allocation process.
- The size of the virtual machine's compute resources (particularly memory) – Larger virtual machines take longer to copy over.
- Availability of sufficient compute resources – The target ESXi host must have enough resources to completely accommodate the virtual machine that is being moved.
- The rate of changes happening on the virtual machine at the time of the vSphere vMotion operation – If pages are being dirtied at a rate faster than they could be copied, vSphere will trigger a Stun During Page Send (SDPS) process, which begins to inject incremental delays inside the virtual machine so that the memory copy transmit rate exceeds the page dirty rate. SDPS helps achieve faster memory pre-copy convergence, which then allows the vSphere vMotion operation to succeed.

vMotion Operation Timeout

The following possible failure scenarios can also cause a vSphere vMotion operation to timeout, halting the migration attempt and leaving the virtual machine at its original location:

- The shadow VM is unable to power on at its destination within 90 seconds.
- The shadow VM powers on, but is unable to receive data from source for up to 120 seconds after power-on.
- The target host becomes unavailable before the switch-over of the virtual machine completes.
- Unreliable network (including intermittent network timeout) causing vMotion to be unable to transmit or receive packets for 20 seconds or longer.
- Slow storage throughput (including I/O congestion) – If the disk (or any storage component's) queue is full and vSphere vMotion is unable to allocate disk blocks for longer than 20 seconds or longer.

The net result of the preceding scenarios and issues described is an overly long vSphere vMotion operation process that might do one of the following:

- Cause the vSphere vMotion operation to completely fail, leaving the virtual machine at its original (source) host.
- Induce failures within the guest's operating system, applications, or services hosted on the virtual machine.

3.3.2 Optimizing vSphere vMotion Operations for Skype for Business

As mentioned previously, some Skype for Business components do not tolerate external interference with their operations gracefully. Such interference might result in degraded performance in the services provided by these components. For example, network latencies, or the effects of blocked vSphere vMotion operations performed on a frontend role server, can result in call drops, distorted voice and video output, or even a complete session disconnection for users connected to the impacted frontend server. This results in sub-optimal user experiences and overall user dissatisfaction.

To avoid such outcomes, VMware recommends that Skype for Business operators adhere to the following recommended practices when configuring vSphere vMotion in a virtualized Skype for Business infrastructure hosted on vSphere:

3.3.2.1. 10 Gb Uplink or Jumbo Frames for vSphere vMotion Uplink

The efficiency of a vSphere vMotion operation depends to a large extent on the overall speed and available bandwidth of the underlying network infrastructure. VMware recommends that the virtual switch on which the vMotion port group resides be connected to a 10 Gb uplink.

When a 10 Gb uplink is not available, customers should enable jumbo frames on both the physical network uplink and the virtual switch to which the vMotion port group is connected

Standard Ethernet frames are limited to a length of approximately 1500 bytes. Jumbo frames can contain a payload of up to 9000 bytes. Support for jumbo frames on VMkernel ports was added to vSphere 4.0 for both ESX and ESXi. This added feature means that large frames can be used for all VMkernel traffic, including vSphere vMotion. Using jumbo frames reduces the processing overhead to provide the best possible performance by reducing the number of frames that must be generated and transmitted by the system.

The use of jumbo frames requires that all network hops between the vSphere hosts support the larger frame size. This includes the systems and all network equipment in between. Switches that do not support (or are not configured to accept) large frames will drop them. Routers and Layer 3 switches might fragment the large frames into smaller frames that must then be reassembled, which can cause performance degradation, latencies, or unnecessary delays during a vSphere vMotion operation. Do not enable jumbo frames within a vSphere infrastructure unless the underlying physical network devices are configured to support this setting.

3.3.2.2. Multiple vSphere vMotion Interfaces

vSphere vMotion operations can be substantially improved by using multiple dedicated vSphere vMotion network interfaces. In most cases, the interfaces that are used for vSphere vMotion are also used for management traffic. Because management traffic is relatively light, this does not add significant overhead.

vSphere provides the ability to use multiple vmnic interfaces for vSphere vMotion traffic to effectively balance the vSphere vMotion traffic load. Testing has shown up to a 25 percent increase in throughput achieved when multiple vSphere vMotion interfaces are used.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Enabling multiple interfaces for vSphere vMotion requires configuring multiple VMkernel ports on different port groups. Each port group can be assigned multiple vmnic interfaces for redundancy. See **Multiple-NIC vMotion in vSphere** for detailed configuration procedures.

Before deciding on using this configuration option, see **vMotion migrations fail when using multiple VMkernel ports for vMotion in different IP subnets**.

3.4 vSphere Distributed Resource Scheduler

Distributed resource scheduling provides active load balancing of virtual machine workloads within a vSphere cluster. In addition to the active monitoring and load balancing functions, DRS provides the following features:

- Virtual machine placement during power on, based on resource requirements and availability.
- Virtual machine evacuation during ESXi host maintenance mode.
- Virtual machine and host groups for grouping like objects.
- Rules to keep virtual machines together or apart and on or off of a set of hosts.

DRS helps make a virtualized Skype for Business Server 2015 environment more agile. The following sections provide recommendations for using DRS with Skype for Business Server 2015

3.4.1 Enable DRS in Partially Automated Mode

DRS provides the following levels of automation:

- Manual – Migration recommendations are provided by DRS. No migrations are performed by DRS.
- Partially automated – Virtual machines are automatically placed on hosts during power-on, migration recommendations are provided by DRS, and no automatic migrations are performed by DRS.
- Fully automated – Virtual machines are automatically placed on hosts during power on and are automatically migrated between hosts to optimize resource usage.

When designed according to VMware recommendations, vSphere clusters that have been purpose built for Skype for Business Server 2015 possess sufficient resources and do not experience many DRS migrations. However, when an ESXi host is placed in maintenance mode, DRS makes recommendations on placement of virtual machines running on that host. DRS also periodically monitors resource consumption within vSphere and, where resource contention or consumption imbalance is detected, can migrate virtual machines within the cluster to alleviate the pressure or imbalance.

Given the importance of stability in a Skype for Business infrastructure, a critical design objective for Skype for Business on vSphere is the minimization of virtual machine migration operations in the environment. vSphere allows for this objective to be achieved while affording customers the opportunity to simultaneously leverage the superior high availability features of vSphere.

Setting the DRS automation level to partially automated helps customers achieve both objectives with very little administrative effort. In partially automated mode, DRS will not migrate a running virtual machine. Rather, DRS will provide visual alerts to the administrator, informing them of the resource imbalance and recommending a manual migration of the impacted virtual machines to a more optimal ESXi host in the cluster. Administrators can then choose whether to follow the recommendation immediately, or defer it for the next round of scheduled maintenance.

The net benefit of this configuration option is the assurance that vSphere will not induce any unintended migration of Skype for Business Server workloads in production, minimizing the possibility of introducing any interference with Skype for Business services from vSphere. The configuration option also makes it

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

so that administrators are proactively made aware of the health state of their virtual infrastructure, enabling them to plan accordingly.

Figure 15. Recommended DRS Automation Level

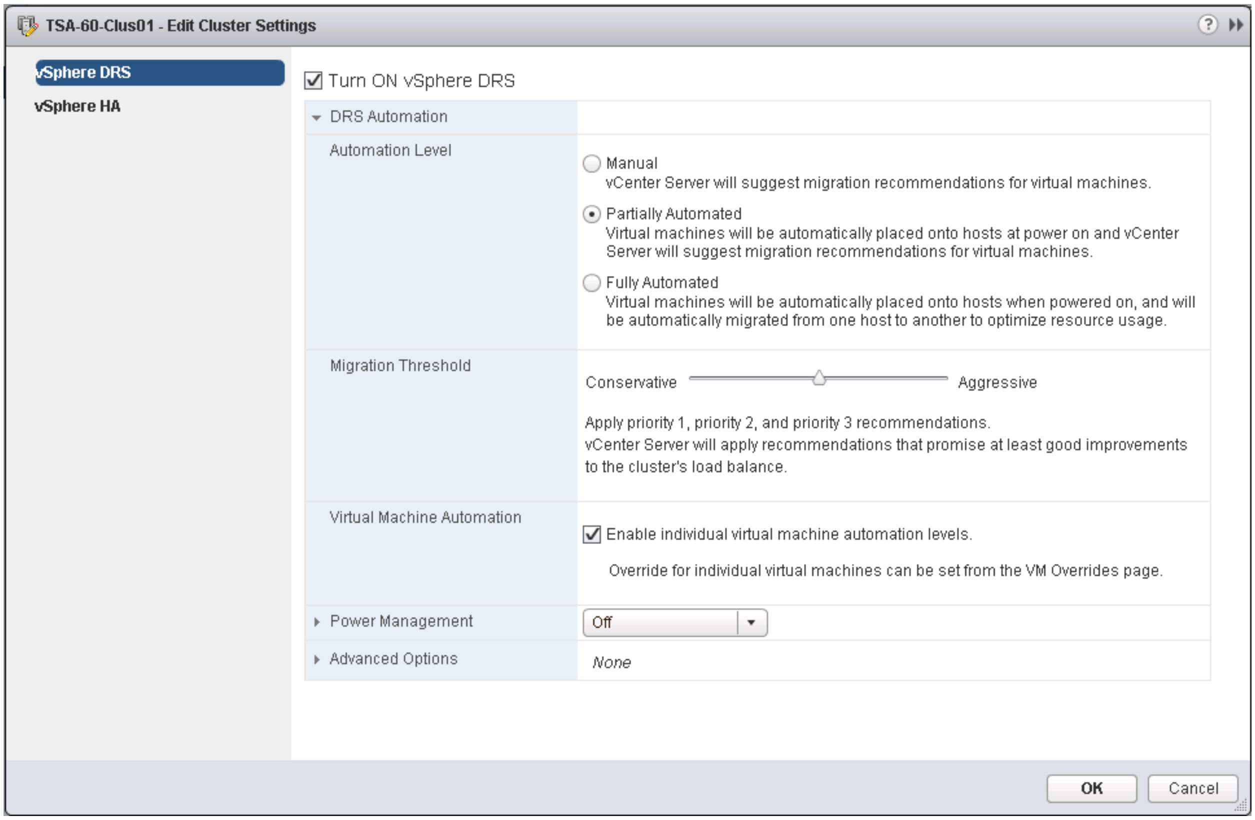

In a shared vSphere infrastructure (where Skype for Business workloads are co-located with other workload types that might require a more aggressive level of automation), enabling DRS at the cluster in fully automatic mode is still not problematic for Skype for Business. This is because vSphere allows administrators to override the cluster-level automation settings at the individual virtual machine level. For example, each Skype for Business Server virtual machine can be configured with a different automation level suitable for customers' needs and requirements, as shown in the following figure:

Figure 16. VM-Level DRS Automation Override

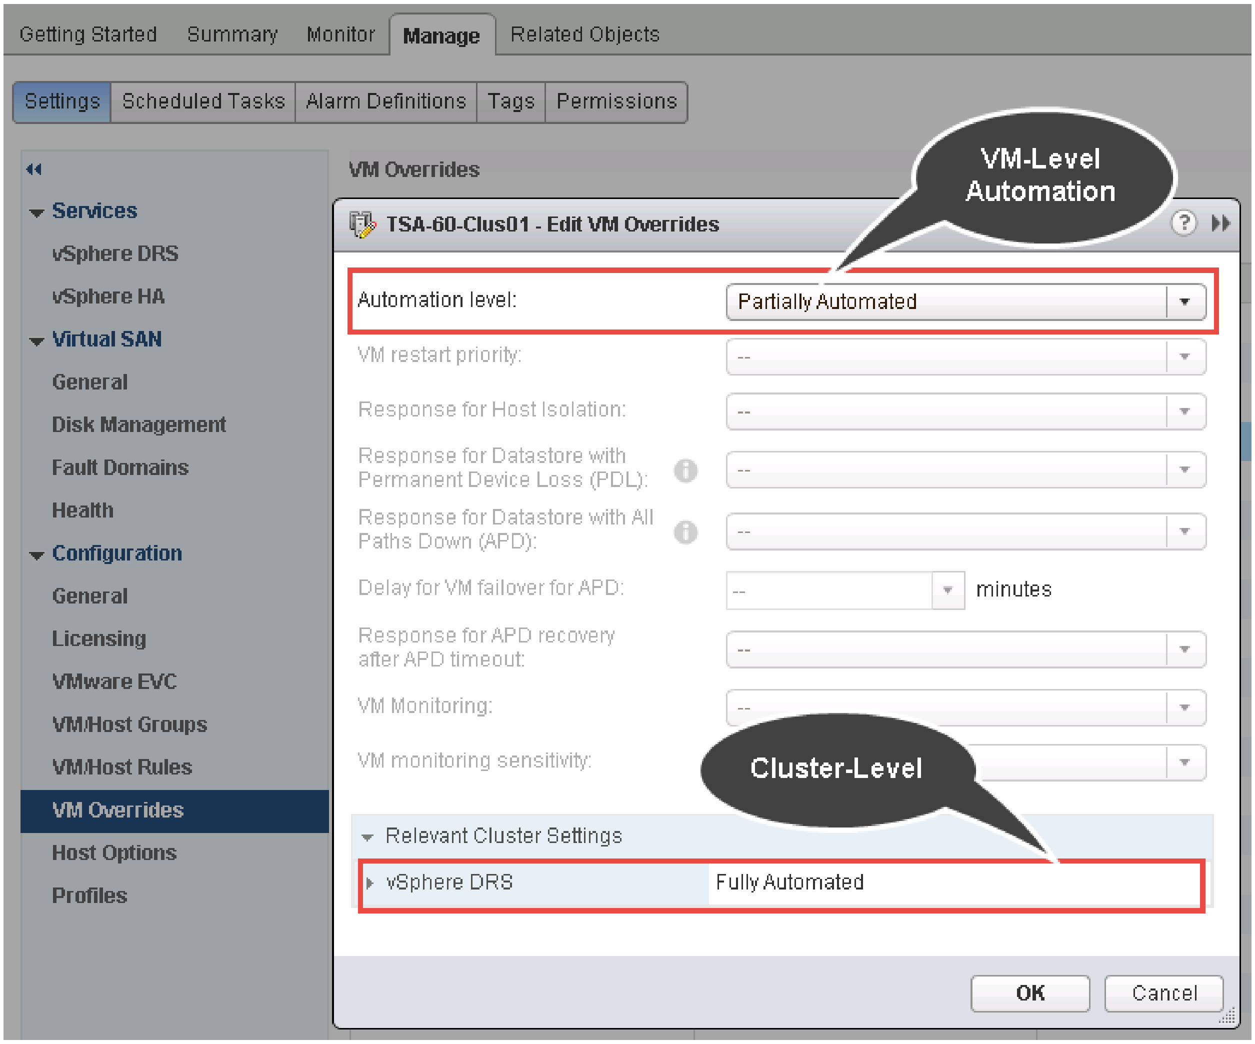

3.4.2 Use Anti-Affinity Rules for Skype for Business Server 2015 Virtual Machines

DRS provides rules for keeping virtual machines apart or together on the same ESXi host or group of hosts. In a Skype for Business Server 2015 environment, the common use case for anti-affinity rules is to keep Skype for Business Server 2015 virtual machines with the same roles installed separated from each other.

The most prevalent HA configuration option for Skype for Business workloads is a scale-out approach, wherein instances of the same roles are installed on multiple servers in the same Skype for Business Server farm. This approach confirms that, in the event of failure (or non-availability) of one of these servers, the roles hosted on it continue to be available in the environment. This approach reduces the impact of a single role server in such environments. Although certain administrative or automated interventions might be required to resuscitate services for impacted users during this failure scenario,

such efforts are significantly less than would otherwise be required if the affected roles were available only on the impacted server.

When making placement decisions (initial placement or during a vSphere vMotion operation), DRS uses resource availability as its primary deciding factor. This means that a vSphere vMotion operation can place Skype for Business virtual machines that host identical roles and services on the same ESXi host, effectively defeating the administrative objectives of avoiding a single-point-of-failure scenario. If the ESXi host on which these similar virtual machines have been placed become suddenly unavailable, all the virtual machines will become simultaneously unavailable.

Due to this consideration, VMware strongly recommends that customers should use the native vSphere affinity and/or anti-affinity features to actively prevent virtual machines that host similar functionalities from being co-located on the same ESXi host for an extended period of time.

3.4.3 DRS Groups and Group-Based Rules

Defining DRS groups helps enforce virtual machine placement in a way that is not possible with only affinity and anti-affinity rules. Host groups can contain hosts licensed to run an application or hosts in a blade chassis or rack. Virtual machine groups usually contain virtual machines performing similar roles and functionalities in an application farm environment. With groups defined, *virtual machines-to-hosts rules* are available for use. Virtual machines-to-hosts rules can be created with four variations:

- Virtual machine group must run on hosts in group.
- Virtual machine group should run on hosts in group.
- Virtual machine group must not run on hosts in group.
- Virtual machine group should not run on hosts in group.

“Must run on” rules provide hard enforcement of virtual machine placement. If a rule stipulates that a group of virtual machines (or specific virtual machine) must run on a group of ESXi hosts (or a particular ESXi host), both DRS and vSphere HA obey these rules. If all hosts in the group are down, the virtual machines are unable to run on any other host in the vSphere cluster.

Note An ESXi host can belong to more than one host group

A virtual machine can be a member of more than one virtual machine group.

In the following figure, two virtual machine groups and two host groups are defined. Two “must run on” rules, shown in the solid black ovals, keep the virtual machines in each group running on their respective host groups. The virtual machine in the middle is not tied to a group or a rule and might roam among all available hosts in the cluster. In the case of a failure of all hosts in the group, all virtual machines bound to those hosts by a “must run on” rule remain offline until a host from that group is brought back online.

One common use for “must run on” rules is to achieve the objective of pinning a specific virtual machine to a specific host, enabling administrators and operators to know the exact location of such virtual machines in the entire vSphere infrastructure.

Figure 17. "Must Run on" Rule Example

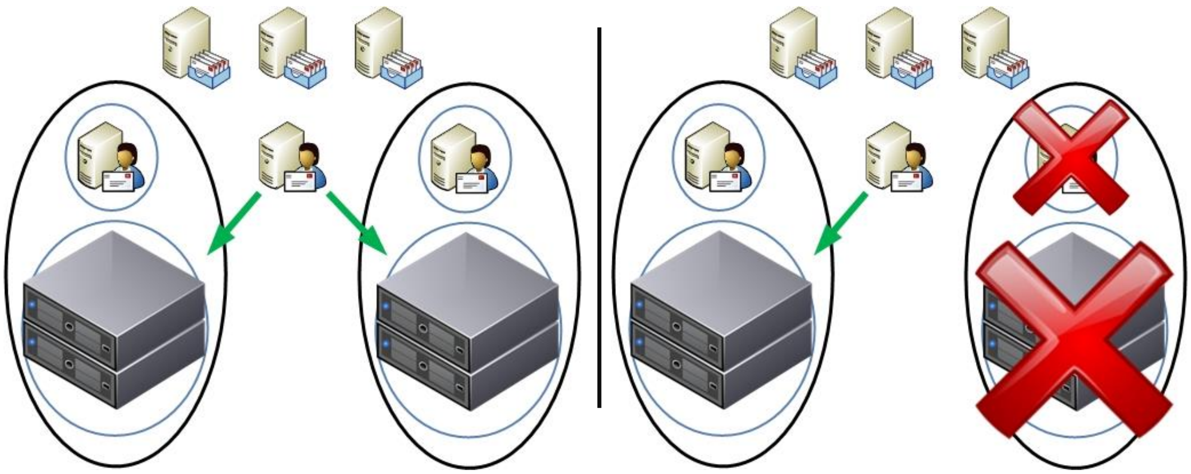

"Should run on" rules provide soft enforcement of virtual machine placement. If a rule stipulates that a group of virtual machines should run on a group of ESXi hosts, those virtual machines will always be preferentially placed on hosts in the host group. They can still run on other hosts in the vSphere cluster outside of the host group, if needed (for example, if all the hosts in the host group are unavailable or otherwise unsuitable for the virtual machine).

In the following figure, two virtual machine groups and two host groups are defined. Two "should run on" rules, shown in the broken green ovals, keep the virtual machines in each group running on their respective host group. The virtual machine in the middle is not tied to a group or a rule and might roam among all available hosts in the cluster. In the case of a failure of all hosts in the group, virtual machines bound to those hosts by a "should run on" rule can be brought back online by vSphere HA.

Figure 18. "Should Run on" Rule Example

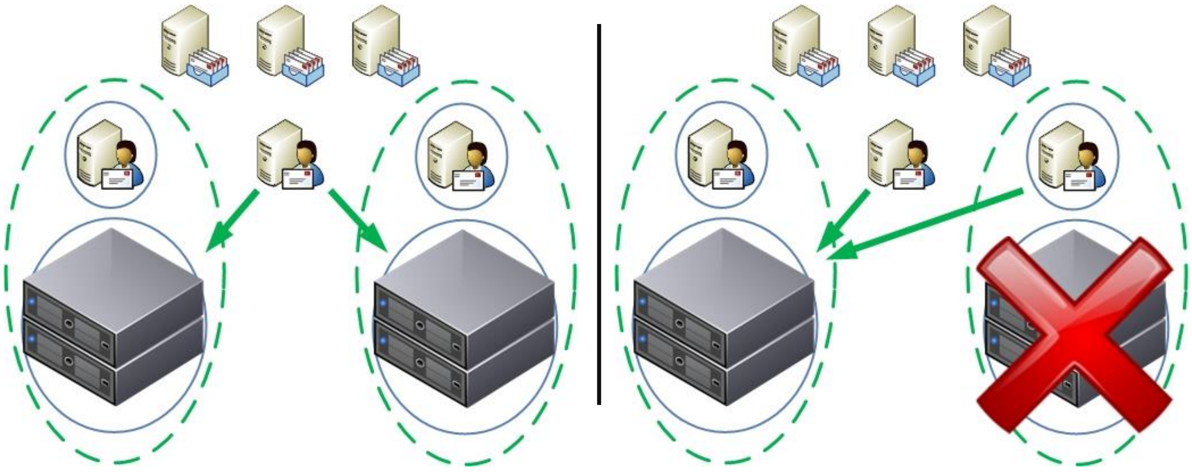

In a Skype for Business Server 2015 environment, virtual machine-to-host rules can be used to provide soft or hard enforcement of virtual machine placement. As an example, consider creating groups of ESXi hosts based on a failure domain, such as a blade chassis or server rack. Create two virtual machine groups with each containing half of the Skype for Business Server 2015 virtual machines, and create rules to link each virtual machine group to a host group. In the case of a complete chassis or rack failure,

any virtual machines that have failed can be powered back on by vSphere HA. Because DRS evaluates its rules at a 5-minute interval, DRS will detect that the virtual machines impacted by the pervious failure have been restarted on non-preferred hosts. If the failed hosts are available, DRS will relocate the virtual machines to enforce and satisfy the configured placement rules.

Rather than use the previously described approach, a better option for Skype for Business workloads on the vSphere infrastructure is the use of the DRS machine-to-machine anti-affinity rules. Anti-affinity rules are easier to configure, update, interpret, and enforce.

3.4.4 DRS Anti-Affinity Rules

Anti-affinity rules enforce virtual machine separation during power on operations and vSphere vMotion migrations due to a DRS recommendation, including a host entering maintenance mode. VMware recommends that customers configure at least one anti-affinity for each group of virtual machines that perform similar Skype for Business roles to prevent them from being co-located on the same ESXi host.

Figure 19. vSphere Distributed Resource Scheduler Anti-Affinity Rule

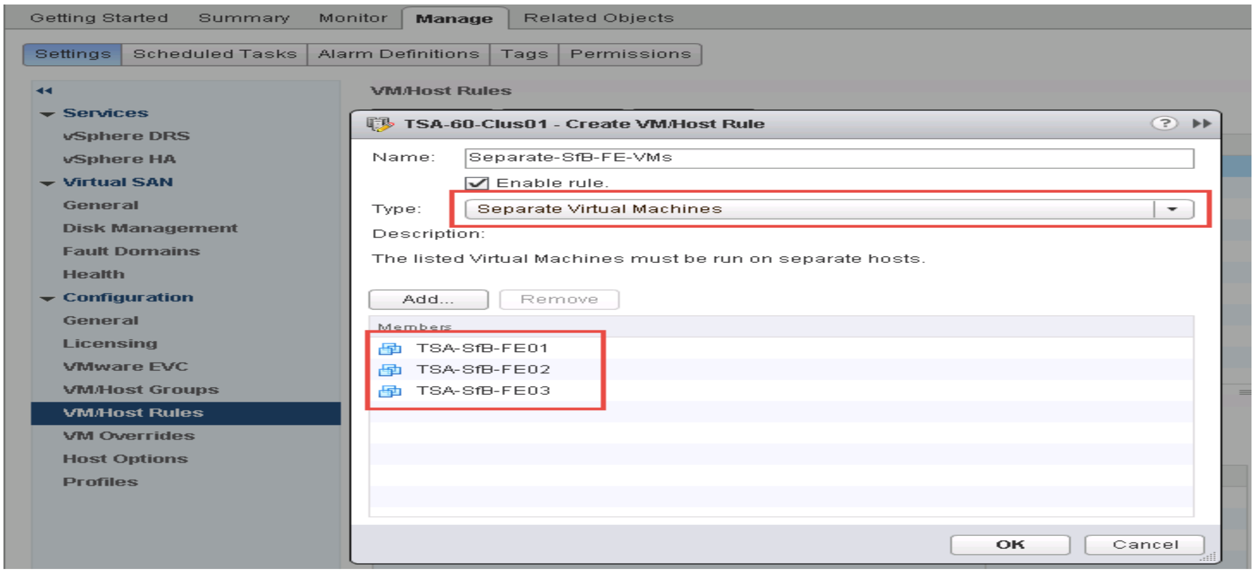

If a virtual machine is enabled for vSphere HA and a host experiences a failure, vSphere HA might power on the virtual machine and violate a DRS anti-affinity rule. This is because the main objective of vSphere HA is the rapid recovery and re-introduction of a failed virtual machine into service, substantially reducing downtime and improving service availability. vSphere HA does not inspect DRS rules during a recovery task, rather it seeks to place the virtual machine on the most optimal ESXi host available at the time of the recovery. DRS will rectify this incorrect virtual machine placement during the next DRS evaluation cycle (every 5 minutes). The virtual machine is migrated to fix the violation.

To avoid this condition when utilizing DRS with vSphere 5.5, VMware encourages customers to apply the following vSphere HA advanced configuration option to their vSphere clusters:

```
das.respectVmVmAntiAffinityRules = TRUE
```

This setting instructs vSphere HA to inspect and respect VM-to-VM anti-affinity rules when restarting virtual machines after a host failure. The Skype for Business Server 2015 virtual machines separated by an anti-affinity will not be co-located on the same ESXi host.

Figure 20. HA Advanced Configuration Option for DRS Anti-Affinity Rules

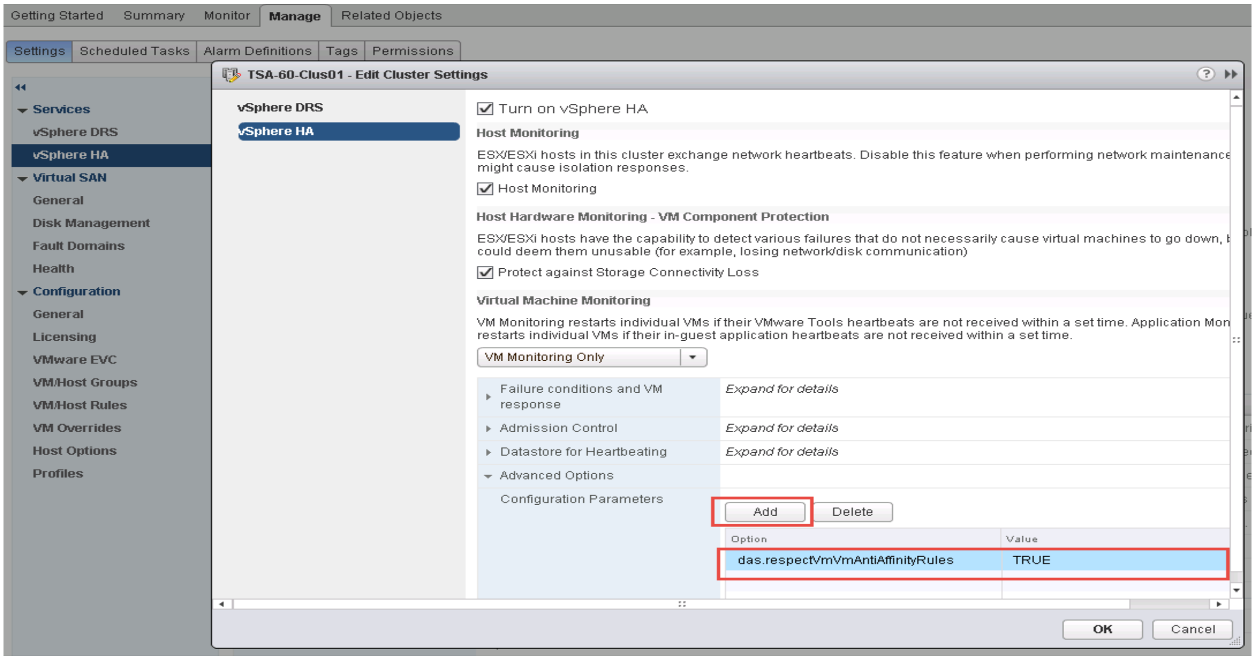

As shown in the following figure, vSphere 6.0 includes an improved, GUI-based configuration option to control how vSphere HA responds to all DRS rules in a cluster. It is no longer required to configure the advanced configuration parameter manually.

Figure 21. Improved vSphere HA and DRS Interoperability in vSphere 6.0

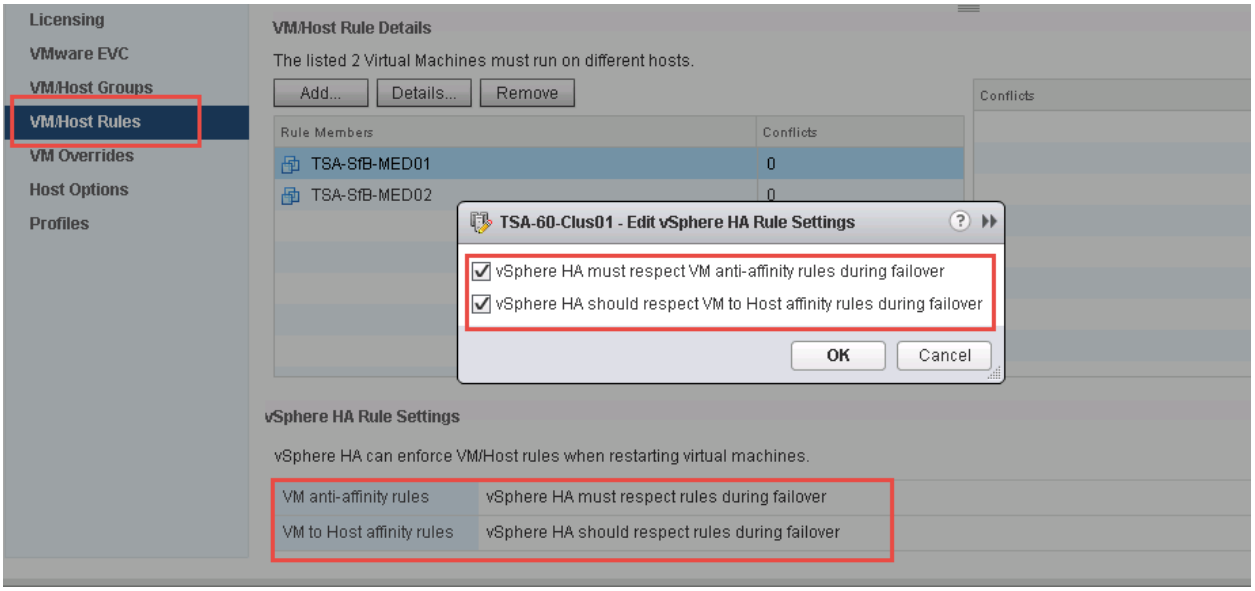

4. Skype for Business Performance on vSphere

VMware and its partners have been testing, validating, and virtualizing Microsoft Lync Server (and now Skype for Business) since 2010, both in controlled and production environments. The outcome of these exercises has been the availability of practical knowledge and mastery of the technical and operational configuration practices required to enable organizations to successfully virtualize and operate their Skype for Business Server workloads on the vSphere platform, and to maximize their hardware resources and investments in the process. Many organizations looking to virtualize resort to physical boxes when necessary due to the needs of the application or support stance of the manufacturer.

Microsoft has supported all of the server roles since Microsoft Lync Server 2010, including the voice roles (mediation and frontend).

Skype for Business' server requirements make the application very attractive to virtualize to companies. With most modern servers coming with 40+ cores and terabytes of RAM, dedicating a single server to a single role would be wasting a lot of these compute resources. Microsoft calls for 12 CPUs and 32 GB of RAM for the frontend server role that consumes the most compute resources. Dedicating any modern server hardware (the only possible option without virtualization) is neither economically advantageous nor practical.

This makes virtualization the most optimal and cost-effective configuration option for Skype for Business. Virtualization helps customers overcome these limitations by efficiently partitioning and maximizing the compute resources and allocating them to the Skype for Business virtual machines as required, without inducing a performance bottleneck.

Many enterprises have adopted virtualization as their default design choice for new workloads. This choice is logical for a Skype for Business infrastructure for the following reasons:

- Advances in server hardware, such as multicore processors, higher memory density, and advances in storage technology, are far outpacing the performance requirements for applications, including Skype for Business. Virtualization becomes an effective way to leverage the full power of these systems.
- The advances in Skype for Business and server hardware technology have coincided with advances in vSphere. On vSphere 6.0, virtual machines support up to 4 TB RAM and 128 vCPUs, both of which are significantly higher than the maximum recommended compute resources for a single Skype for Business Server. The question of whether or not virtualization can support the performance requirements of Skype for Business is, therefore, moot.

4.1 Key Performance Considerations

A variety of factors can affect Skype for Business performance on vSphere, including processor and memory allocation to the guest virtual machine, storage layout and design, virtual machine placement, and high availability methods. The following are tips for achieving the best possible performance:

- Fully understand your organization's business and technical requirements for implementing Skype for Business.
- Fully understand the Skype for Business workload requirements. Microsoft gives recommended resources for each machine type (https://technet.microsoft.com/en-us/library/dn951388.aspx). These recommendations will support thousands of users. It is not recommended to scale down if you have fewer users. Doing so could potentially introduce issues or call quality problems.
- Although I/O is reduced in Skype for Business compared with prior versions of Lync, there is still a requirement to provide adequate throughput and low latency. Dedicate physical storage for Skype for Business to avoid compromising I/O by having other workloads running on the same physical disks.

- Use the correct sizing tools to obtain the appropriate configuration guidelines for your Skype for Business virtual machines. The primary tools for sizing a Skype for Business infrastructure are as follows:
  - Skype for Business Capacity Calculator (https://www.microsoft.com/en-us/download/details.aspx?id=51196) – Obtain the correct version of these tools for your sizing exercise. For example, using the Lync Server 2010 version of the tool to size a Skype for Business implementation will result in sizing recommendations that will result in noticeable performance bottleneck.
  - Skype for Business Bandwidth Calculator (https://www.microsoft.com/en-us/download/details.aspx?id=19011) – This tool is used to determine how much network bandwidth is needed by Skype for Business at each location. Having the appropriate amount of network bandwidth will have an impact on voice quality and should be considered carefully.

  VMware strongly encourages customers to utilize these tools when sizing their Skype for Business workloads.
- Follow the best practices in Section 2, ESXi Host Best Practices for Skype for Business Server 2015, to optimize the ESXi host environment for enterprise applications such as Skype for Business.

4.2 Key Health Indicators

Skype for Business includes many performance counters (literally thousands) which enable administrators to dive deep for troubleshooting. Microsoft has provided a reduced set of these performance counters that are crucial to the health of the Skype for Business environment. Microsoft termed these performance counters as Key Health Indicators (KHI).

The KHIs are broken down into three rings. The rings build upon each other, that is, if Ring 0 has problems, then you will see issues in Rings 1 and 2. Solve the issues in the lowest number ring before moving on to the next.

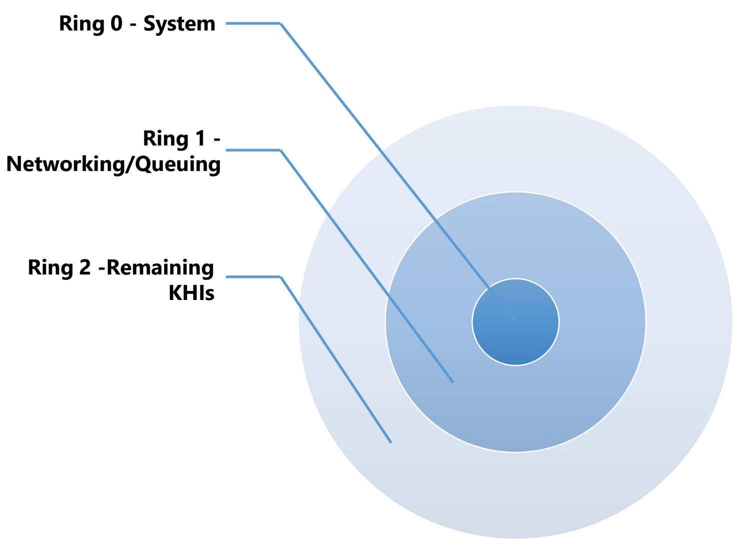

- Ring 0 is for all the system KHI's (CPU, RAM, and disk). Issues in this ring will be based on the infrastructure that is running Skype for Business (host, virtual machine, and storage).

- Ring 1 consists of queue latencies and network. These KHIs are heavily impacted by the KHIs in Ring 0.
- Ring 2 is all the remaining KHIs.

To setup and collect the KHIs, see https://www.microsoft.com/en-us/download/details.aspx?id=46895.

When collecting the KHI data, you can set the time interval for how often the performance counters are recorded. When collecting data for the first time or in a reactive model, set the interval for 15 seconds. If collecting in an on-going model, set the collection interval for every 60 seconds.

For this document, we will highlight the KHIs in Ring 0 as they are impacted by the VMware design.

4.2.1 System KHIs

Performance Counter	Description	Value
% Processor Time	Processor usage across all vCPUs.	< 80%
Physical Disk Read/Write	Represents the average time (in seconds) of either disk read or write latency.	< 25ms (0.025)
Available Megabytes	Displays the memory available to the server.	> 2.5 GB
Network Interface – Output Queue Length	Measures the packet queue length of the output packet queue.	0
Network Interface – Inbound/Outbound Packets Discarded	Incoming/outgoing packets not in error that have been discarded.	< 100

4.3 Performance Testing

Every Skype for Business environment is different, with varying business and technical requirements, many server and storage options. To definitively test the underlying infrastructure, VMware recommends that customers utilize the **"Stress Test"** utility from Microsoft (https://www.microsoft.com/en-us/download/details.aspx?id=41935). This test should only be run in a non-production, test environment. The Stress Test utility will create user accounts and perform activities to put load on the servers. It will report back on the number of calls and any errors it encountered. It will not report back any call quality statistics.

By utilizing the Stress Test utility alongside watching the KHIs, it can be determined whether or not the physical hardware and the virtual machines can withstand the anticipated load.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

5. Summary of VMware Recommended Settings for a Virtualized Skype for Business Infrastructure

The following table is a comprehensive listing of the configuration options and choices recommended by VMware. These options have been tested and validated in both controlled and production environments and found to be the most optimal for satisfying the performance, availability, recoverability, and ease-of-administration requirements of an enterprise-level Skype for Business infrastructure.

VMware encourages customers to adhere to these recommendations as closely as possible to modify them where required, only after a robust testing and validation effort.

Table 4. Summary of VMware's Recommended Settings

Server Hardware (BIOS) Settings **Please consult vendor references for instructions on configuring recommended settings**	
Option	**Setting**
Virtualization Technology Support	Enabled
CPU Power Management	Maximum Performance
Memory Frequency	Maximum Performance
NUMA	Enabled (Disable node-interleaving, if available as an option)
Hyper-Threading	Enabled
All C-States and P-States	Disabled
Turbo Boost	Enabled
Energy Efficient Turbo	Disabled
Collaborative CPU Performance Control	Disabled
Number of CPU Cores Enabled	All
MWAIT	Enabled
Energy Policy	Performance
X2Aic Mode	Enabled
Processor Idling	Disabled
QPI Speed	Maximum Data Rate

Alternate RTID	Disabled
Adjacent Cache Line Prefetch	Disabled (Disable all prefetching” options)
Address Translation Services (ATS)	Enabled
Patrol Scrub	Disabled
Demand Scrub	Disabled
Memory RAS	Maximum Performance
Low Voltage DDR Mode	Performance-mode
Channel or Rank Interleaving	Automatic
Intel Speedstep	Enabled
Execute Disabled Bit	Disabled
ESXi Host Settings **VMware vSphere Settings**	
Cluster Settings (including DRS/HA)	
Power Management	**Setting**
Power Plan	High Performance
Networking	**Setting**
Virtual Switches (where vMotion PG is connected)	Enable Jumbo Frames – If supported by physical network infrastructure
Use NTP	Enabled
NTP Client	Enabled
NTP Service Startup Policy	Start and Stop with Host
NTP Servers	pool.ntp.org (or anything reliable)

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Configuration	Setting
vSphere HA	Enabled
vSphere DRS	Enabled
DRS Automation Level	Partially Automatic
Host monitoring	Enabled
Host Hardware Monitoring	Protect Against Storage Connectivity Loss
Host Isolation Response	Disabled
Power Management	Off (Disables DPM)
Datastore Heart-beating	Automatic
PDL and APD Response	Disabled
Recovery After APD Timeout	Disabled
Admission Control	Do Not Reserve Failover Capacity (Need to bring all VMs back up ASAP, even if it results in temporary degraded performance)
VMware EVC Mode	Disabled
Migration Threshold	Conservative
Virtual Machine Automation	Enabled
VM Monitoring	VM Monitoring Only
VM Monitoring Sensitivity	High (Default)
VM restart priority	Medium (This is global. Will use per-VM configuration to prioritize VM Restart)
Fault Tolerance	Disabled
Hyper-threading	Enabled (BIOS settings, verify in vCenter)
Swap File Location	Default
Default VM Compatibility	ESXi 6.0 and Later
vSphere HA and DRS Affinity Rules	HA must respect VM anti-affinity rules during failover
	HA should respect VM to Host affinity rules during failover

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Virtual Machine Settings	
Latency-Sensitivity	**Setting**
FE, Mediation and Edge Role VMs	High
SQL Server and other VMs	None
Networking	**Setting**
Virtual Network Card Type	VMXNet3
Time Synchronization	**Setting**
Synchronize Time with Host	Disabled (See **Error! Reference source not found.** section for the advanced configuration options to add to VM templates that will be used for deployment)
Check and upgrade VMware Tools before each power-on	Enabled (Might cause slower VM startup, but only if there is a new VMware Tools version)
CPU Allocation	**Setting**
Backend Microsoft SQL Server (Full version)	If allocating more vCPUs than maximum sockets supported by SQL version, see section 2.1.5.1 for notes and recommendation Otherwise, allocate by Sockets (Leave “Cores Per Socket” at 1)
Embedded SQL Express version	See section 2.1.5.1 for notes and recommendation
CPU Hot-Add	Disabled
Limit	Unlimited
Reservation	0 (Minimum)
All other settings	Default

Memory Allocation	**Setting**		
Reserve all guest memory	Enabled		
Memory Hot-Plug	Disabled		
Limit	Unlimited		
All other settings	Default		
Power Management	**Setting**		
Standby Response	Put the guest OS into standby mode and leave the virtual machine powered on		
Storage Allocation	**Setting**		
Datastore Size	Determined by size of customer's infrastructure and needs		
SCSI Controller	Virtual SCSI Controller Type	VMware Paravirtual (inject OS-specific drivers during OS install, or deploy from template where vSCSI has been converted to PVSCSI)	
	Virtual SCSI Controller Count	Minimum 2	
	SCSI Bus Sharing	None (if using FCI for MS SQL Server, then choose Physical)	
Virtual Hard Disk	Type	SQL Server, FE, Mediation VMs	Eager-zero Thick
		Others	Lazy-zero Thick
	Sharing	Unspecified (Do not set to Multi-Writer)	
	Shares	Normal	

	Limit IOPs	Unlimited
	Virtual Flash Read Cache	0
	Disk Mode	Dependent
Windows Settings		
Disk Management	**Setting**	
Layout	GPT	
File System	NTFS	
Perform a Quick Format	Disabled	
PVSCSI Ring Buffer Adjustment	See **Error! Reference source not found.** section for registry key options	
Network Card	**Setting**	
MTU Setting	To be determined by customers, based on physical network capabilities	
IPv4 Checksum Offload (and all other "Offloads")	Enabled	
Maximum number of RSS Processors	Not Present (Default)	
Receive Side Scaling	Enabled Also enable it in Windows kernel (See **Error! Reference source not found.** section for instructions)	
Receive Throttle	Not Present (Default)	
RSS Base Processor Number	Not Present (Default)	
Receive Segment Coalescing (RSC)	Disabled (Unless persistently high CPU usage or packet drops are observed)	
Receive Window Auto-tuning Level	Auto	

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Rx Ring #1, #2, and so on (all Ring Size settings)	Not Present (Default)
Speed and Duplex	Auto
All other settings	Default
Power Management	**Setting**
Power Plan	High Performance Use powercfg -list to verify
Time Synchronization	**Setting**
All Domain-joined VMs Should Get Their Time from the Domain Hierarchy	Use w32tm /query /peers| find /i "peer:" Look for a domain controller's name in the returned result

Notes

The following table shows the VM advanced configuration options for disabling VMware Tools guest clock reset.

Configuration Option	Setting
`tools.syncTime`	0
`time.synchronize.continue`	0
`time.synchronize.restore`	0
`time.synchronize.resume.disk`	0
`time.synchronize.shrink`	0
`time.synchronize.tools.startup`	0
`time.synchronize.tools.enable`	0
`time.synchronize.resume.host`	0

Ensure that RSS is Enabled:

- At the NIC level, use PowerShell

  - Get current status of RSS for your NIC
    - Get-NetAdapterRss |fl name, enabled
  - Identify the name of the NIC from the result in previous command. Enable RSS, if not enabled
    - Enable-NetAdapterRss -name <NIC-displayname>
- At the Windows kernel, use netsh
  - Check the NIC Properties inside the Guest OS. Verify that "Receive Side Scaling" is enabled
  - Run "netsh int tcp show global" from a CMD prompt to see if Receive-Side Scaling (RSS) is enabled.
    - If it's not enabled, run "netsh int tcp set global rss=enabled" from an escalated CMD prompt to enable it

PVSCSI Ring Buffer (http://kb.vmware.com/kb/2053145):

- REGEDIT:

```
Windows Registry Editor Version 5.00
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\pvscsi\Parameters\Device]
"DriverParameter"="RequestRingPages=32,MaxQueueDepth=254"
```

- POWERSHELL:
  - New-Item -Path "HKLM:\SYSTEM\CurrentControlSet\services\pvscsi\Parameters\Device"
  - Set-ItemProperty -Path "HKLM:\SYSTEM\CurrentControlSet\services\pvscsi\Parameters\Device" -name DriverParameter -value "RequestRingPages=32,MaxQueueDepth=254" | Out-Null

Backend Microsoft SQL Server VMs Only

Windows Server Failover Clustering (AG or FCI) Timeout Settings:

- (get-cluster).SameSubnetThreshold = 10
- (get-cluster).CrossSubnetThreshold = 20
- (get-cluster).RouteHistoryLength = 40
- Get-Cluster| fl *history*

Enable Large Pages:

- PowerShell:
  - Set-Itemproperty -Path "HKLM:\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQLServer\Parameters" -name SQLArg3 -value "-T834"

    Note: “MSSQL13” is for Microsoft SQL Server 2016 version. Please adjust based on your Microsoft SQL Server version.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com