



MARS

VMware Assessment Report

Sample

Table of Contents

- 1 Executive Summary 3
- 2 Environment Overview 3
 - 2.1 Storage..... 3
 - 2.2 Host..... 4
 - 2.3 Virtual Machine 4
- 3 VMware Assessment Summary..... 5
- 4 Assessment Details 6
 - 4.1 Performance Issues..... 6
 - 4.1.1 Host..... 6
 - 4.1.2 Virtual Machine 7
 - 4.1.3 Storage..... 8
 - 4.1.4 Network 9
- 5 Appendix 10
 - 5.1 Appendix A: Slow Storage Device 11
 - 5.2 Appendix B: Resolving Slow Storage Issues 11

1 Executive Summary

This report presents the results of the assessment of the configuration and performance metrics of the VMware infrastructure at HISYS.

This analysis is based on the data collected between 2013-03-18 08:00 and 2013-03-18 17:00 by MARS.

2 Environment Overview

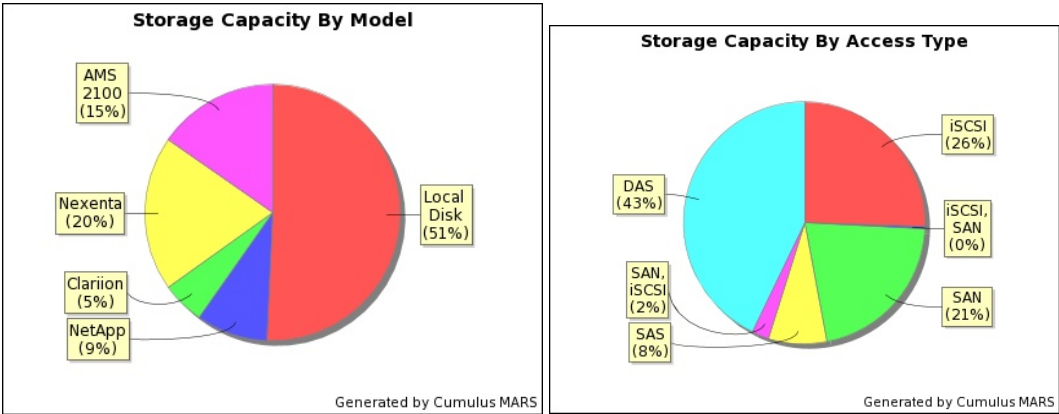
| | |
|--------------------------------|---------------|
| vCenter(s) Used | 192.168.20.13 |
| No. of ESX Hosts | 9 |
| No. of Virtual Machines | 262 |


2.1 Storage

Total Capacity : 16.95 TB

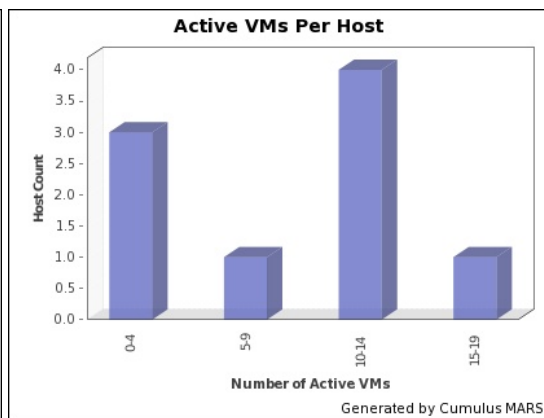

Used Capacity : 16.35 TB

Available Capacity : 607.86 GB


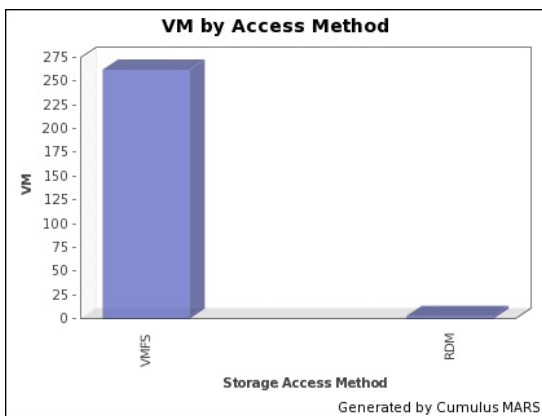








2.2 Host



2.3 Virtual Machine







3 VMware Assessment Summary

| Performance | | |
|------------------|---|---|
| Host |  | Did not pass 2 of 4 checks including 1 critical check |
| Virtual Machines |  | Did not pass 2 of 5 checks including 1 critical check |
| Storage |  | Did not pass 3 of 8 checks including 1 critical check |
| Network |  | Passed all 2 checks |






4 Assessment Details

4.1 Performance Issues







4.1.1 Host



| | |
|--|---|
| <p>1. CPU Saturation</p> <p><i>CPU Usage greater than 90% continuously for at least 5 minutes.</i></p> <p>Result: 1 CPUs violated this threshold during the assessment period</p> |  |
| <p>2. Swap-in Recurrence</p> <p><i>Host was swapping in at least 1 MB/s of data or more continuously for 5 minutes.</i></p> <p>Result: 0 hosts violated this threshold during the assessment period</p> |  |
| <p>3. Swap-out Recurrence</p> <p><i>Host was swapping out at least 1 MB/s of data or more continuously for 5 minutes</i></p> <p>Result: 1 hosts violated this threshold during the assessment period</p> |  |
| <p>4. Memory Ballooning</p> <p><i>Balloon driver has claimed at least 1GB of memory continuously for 5 minutes</i></p> <p>Result: 4 hosts violated this threshold during the assessment period.</p> <p>Analysis:</p> <p>Ballooning is part of normal operation when memory is over-committed. The fact that ballooning is occurring is not necessarily an indication of a performance problem. The use of the balloon driver enables the guests to give up memory pages that they are not currently using. However, if ballooning causes the guest to give up memory that it actually needs, performance problems can occur due to guest OS paging.</p> <p>Recommendation:</p> <p>Determine whether memory is over-committed on a Host by using vSphere client..</p> |  |

4.1.2 Virtual Machine



| | |
|--|---|
| <p>5. CPU Saturation</p> <p><i>CPU usage greater than 90% continuously for 5 minutes</i></p> <p>Result: 2 CPUs violated this threshold during the assessment period.</p> |  |
| <p>6. High ready times</p> <p><i>Ready time > 2000ms continuously for 5 minutes</i></p> <p>Result: 17 VMs violated this threshold during the assessment period</p> |  |
| <p>7. Swap-in Recurrence</p> <p><i>Host was swapping in at least 1 MB/s of data or more continuously for 5 minutes</i></p> <p>Result: 0 VMs violated this threshold during the assessment period.</p> |  |
| <p>8. Swap-out Recurrence</p> <p>Host was swapping out at least 1 MB/s of data or more continuously for five minutes.</p> <p>Result: 0 Virtual Machines violated this threshold during the assessment period.</p> |  |
| <p>9. Memory Ballooning</p> <p><i>Balloon driver has claimed at least 1GB of memory continuously for 5 minutes</i></p> <p>Result: 4 hosts violated this threshold during the assessment period.</p> <p><u>Analysis</u></p> <p>Ballooning is part of normal operation when memory is over-committed. The fact that ballooning is occurring is not necessarily an indication of a performance problem. The use of the balloon driver enables the guests to give up memory pages that they are not currently using. However, if ballooning causes the guest to give up memory that it actually needs, performance problems can occur due to guest OS paging.</p> <p><u>Recommendation</u></p> <p>Determine whether memory is over-committed on a Host by using vSphere client..</p> |  |

4.1.3 Storage

| | |
|---|---|
| <p>10. SCSI Command Aborts</p> <p><i>Number of times a SCSI command was aborted because the storage system was taking too long to respond.</i></p> <p>Result: 0 disks violated this threshold during the assessment period.</p> |  |
| <p>11. SCSI Bus Resets</p> <p><i>Number of times SCSI I/O request has failed</i></p> <p>Result: 0 disks violated this threshold during the assessment period.</p> |  |
| <p>12. Disk Read Latency</p> <p><i>Device Read Latency > 20 ms for at least 5 continuous minute.</i></p> <p>Result: 4 disks violated this threshold during the assessment period.</p> |  |
| <p>13. Disk Write Latency</p> <p><i>Device Write Latency > 20 ms for at least 5 continuous minutes</i></p> <p>Result: 8 disks violated this threshold during the assessment period.</p> <p><u>Analysis</u></p> <p><u>Recommendations</u></p> <ul style="list-style-type: none">o Size storage for performance as well as capacityo Moving to a virtualized environment changes storage workloads.o Understand the load being placed on storage devices.o Simple benchmarks can help isolate storage performance problems.o Consider the trade-off between memory capacity and storage demand. <p>For details, refer to Appendix B ‘Slow Storage Devices’ below..</p> |  |
| <p>14. Kernel Read Latency</p> <p><i>Kernel Read Latency > 20 ms for a minute</i></p> <p>The kernels read and write latencies are defined by VMware as the time spent in the VMKernel processing the request. Very busy kernel (too many VMs) results in this value being high.</p> <p>Result: 2 disk’s violated this threshold during the assessment period.</p> |  |
| <p>15. Kernel Write Latency</p> |  |

| | |
|--|---|
| <p><i>Kernel Write Latency > 20 ms for a minute</i></p> <p>The kernels read and write latencies are defined by VMware as the time spent in the VMKernel processing the request. Very busy kernel (too many VMs) results in this value being high.</p> <p>Result: 4 disks violated this threshold during the assessment period.</p> | |
| <p>16. Queue Read Latency</p> <p>Queue Read Latency > 20 ms for a minute.</p> <p>The amount of time spent once a command has been queued up for execution at the storage. Some people also occasionally call this 'HBA queue'.</p> <p>Result: 2 disks violated this threshold during the assessment period.</p> |  |
| <p>17. Queue Write Latency</p> <p><i>Queue Write Latency > 20 ms for a minute</i></p> <p>Result: 4 disks violated this threshold during the assessment period.</p> |  |

4.1.4 Network

| | |
|--|---|
| <p>18. Received Packets Dropped</p> <p><i>One or more received packets dropped for at least 5 continuous minutes</i></p> <p>Result: 0 Received packets</p> |  |
| <p>19. Transmitted Packets Dropped</p> <p><i>One or more transmitted packets dropped for at least 5 continuous minutes.</i></p> <p>Result: 0 Transmitted packets dropped</p> |  |

5 References

- CPU Saturation Check: VMware document, “vsphere41-performance-troubleshooting-performance community document Page 29”
- Memory Issues: Swap In VMware document “vpshere41-performance-troubleshooting-performance community document Page 36”
- Memory Issues Swap Out VMware document “vpshere41-performance-troubleshooting-performance community document Page 36”
- Memory Issues Balloon VMware document “vpshere41-performance-troubleshooting-performance community document Page 36”
- Ready time > 2000 ms VMware document “vpshere41-performance-troubleshooting-performance community document Page 13”
- Device Read Latency > 20 ms for 5 or more minutes VMware document “vpshere41-performance-troubleshooting- performance community document Page 20”

6 Appendix

6.1 Appendix A: Slow Storage Device

In order to understand whether a high disk read or write latency value represents an actual problem, it is necessary to understand the storage workload. There are three main workload factors that affect the response-time of a storage subsystem:

- o I/O arrival-rate
A given configuration of a storage device will have a maximum rate at which it can handle specific mixes of I/O requests. When bursts of requests exceed this rate, they may need to be queued in buffers along the I/O path. This queuing can add to the overall response-time.
- o I/O Size
The transmission rate of storage interconnects, and the speed at which an I/O can be read from or written to disk, are fixed quantities. As a result, large I/O operations will naturally take longer to complete. A response-time that is slow for small transfers may be expected for larger operations.
- o I/O Locality
Successive I/O requests to data that is stored sequentially on disk can be completed more quickly than those that are spread randomly.

In addition, read requests to sequential data are more likely to be completed out of high-speed caches in the disks or arrays.

Storage devices typically provide monitoring tools that allow data to be collected in order to characterize storage workloads according to these, and other, factors. In addition, storage vendors typically provide data on recommended configurations and expected performance of their devices based on these workload characteristics. If the problem check for Slow Storage indicated that storage may be causing performance problems, the monitoring tools should be used to collect workload data, and the storage response-times should be compared to expectations for that workload. If investigation determines that storage response-times are unexpectedly high, corrective action should be taken.

6.2 Appendix B: Resolving Slow Storage Issues

Due to the complexity and variety of storage infrastructures, it is impossible to specify specific solutions for slow or overloaded storage in this document. vSphere is capable of high storage performance using any supported storage technology, and, beyond rebalancing load, there is typically little that can be done from within vSphere to solve problems related to a slow or overloaded storage device. Documentation from your storage vendor should be followed to monitor the demand being placed on the storage device, and vendor-specific configuration recommendations should be followed to configure the device for the demand. If the device is not capable of satisfying the I/O demand with good performance, the load should be distributed among multiple devices, or faster storage should be obtained. To help guide the investigation of performance issues in the storage infrastructure, we offer some general notes on storage

performance, and the differences between physical and virtual environments, that should be considered when investigating storage performance problems in a vSphere environment.

- Size storage for performance as well as capacity.

As physical disks continue to grow in storage capacity, it becomes possible to satisfy the storage capacity requirements of applications with a very small number of disks. However, the I/O rate that each individual disk can handle is limited. As a result, it is possible to fulfill the storage space requirements of an application with a disk device that cannot physically handle the I/O request rates the application generates. The ability of a storage device to handle the bandwidth and IOPS demands of an application is as important a selection criterion as is the capacity of the device. Storage vendors typically provide data on the IOPS that their devices can handle for different types of I/O loads and configuration options. Sizing storage for performance is particularly important in a virtualized environment. When applications are running on a server in a non-virtualized environment, the storage for that application is often placed on a dedicated LUN with specific performance characteristics. When moved to a vSphere environment, the storage for an application may be taken from a pre-existing VMFS volume, which may be shared by multiple VMs running different applications. If the LUN on which the VMs are placed was not sized and configured to meet the needs of all VMs, the storage performance available to the applications running in the VMs will be sub-optimal. As a result, it is critical to consider the performance capabilities of VMFS volumes and their backing LUNs, as well as available capacity, when placing VMs.

- Moving to a virtualized environment changes storage workloads.

In a non-virtualized environment, the configuration of LUNs is often based on the I/O characteristics of individual applications. Characteristics such as IOPS rate, I/O size, and disk access-pattern all affect the configuration of storage and its ability to satisfy the performance demands of an application. When multiple applications are consolidated onto a single VMFS volume, the workload presented to that volume will be different than the storage workload of the individual applications. A particular example is that the interleaving of I/Os from multiple applications may cause sequential request streams to be presented to the storage device as random requests. This will increase the number of physical disk devices needed to meet the storage performance needs of the applications.

- Understand the load being placed on storage devices.

In order to properly troubleshoot storage performance problems, it is important to understand the load being placed on the storage device. Many storage arrays have tools that allow workload statistics to be captured. In addition to providing information on IOPS and I/O sizes, these tools may provide information on queuing in the array, and statistics regarding the effectiveness of the caches in the array. Refer to the vendor's documentation for details. VMware ESX also provides tools for understanding storage workloads. In addition to the high-level data provided by the vSphere Client and esxtop, the vscsiStats tool can provide detailed information on the I/O workload generated by a VM's virtual SCSI device. See <http://communities.VMware.com/docs/DOC-10095> for more details.

- Simple benchmarks can help isolate storage performance problems. Once slow storage has been identified, it can be difficult to troubleshoot the problem using application-level loads. When using live applications, problems may occur intermittently or only during peak periods. Driving applications with synthetic loads can be time-consuming and requires complex tools. Disk benchmarking tools, such as IOmeter, can be used to generate loads similar to application loads in a controllable manner. vscsiStats tool can help to capture the I/O profile of an application running in a VM. Details such as outstanding I/O requests, I/O request size, read/write mix, and disk access-pattern can be obtained using vscsiStats. These details can be used to create a workload in IOmeter. This I/O load should simulate the I/O behavior of the actual application. See <http://communities.VMware.com/docs/DOC-3961> for more information.
- Consider the tradeoff between memory capacity and storage demand. Some applications can take advantage of additional memory capacity to cache frequently used data, thus reducing storage loads. This typically has the added benefit of improving application performance. Because running applications in VMware ESX makes it easy to add additional memory capacity to a VM, this tradeoff should be considered when storage is unable to meet existing demands. Refer to the application-specific documentation to determine whether an application can take advantage of additional memory. Some applications require changes to configuration parameters in order to make use of additional memory.