

HPE SimpliVity Sizing

Using data from Lanamark One

May 2019

Version 2.6

Acknowledgement:

All product and company names are trademarks™ or registered® trademarks of their respective holders. Use of them does not imply any affiliation with or endorsement by them.

Introduction

This document supports the Atrium course “HPE SimpliVity Sizing and Configuration Workshop” and is not a comprehensive or standalone reference document as it is intended for those who have attended the course.

This is intended to help pre-sales consultants to use the HPE SimpliVity Sizing tool more effectively, by extracting relevant data from the Lanamark One spreadsheet which is produced by running Lanamark One on an end-user’s infrastructure.

Using Lanamark One data for HPE SimpliVity Sizing

Lanamark One provides the detailed metrics required by the HPE SimpliVity Sizing tool. A useful inventory is created as a formatted document, which can be used to present information to the end-user. An example for a fictitious company, HLR Telco is shown in figure 1 below.

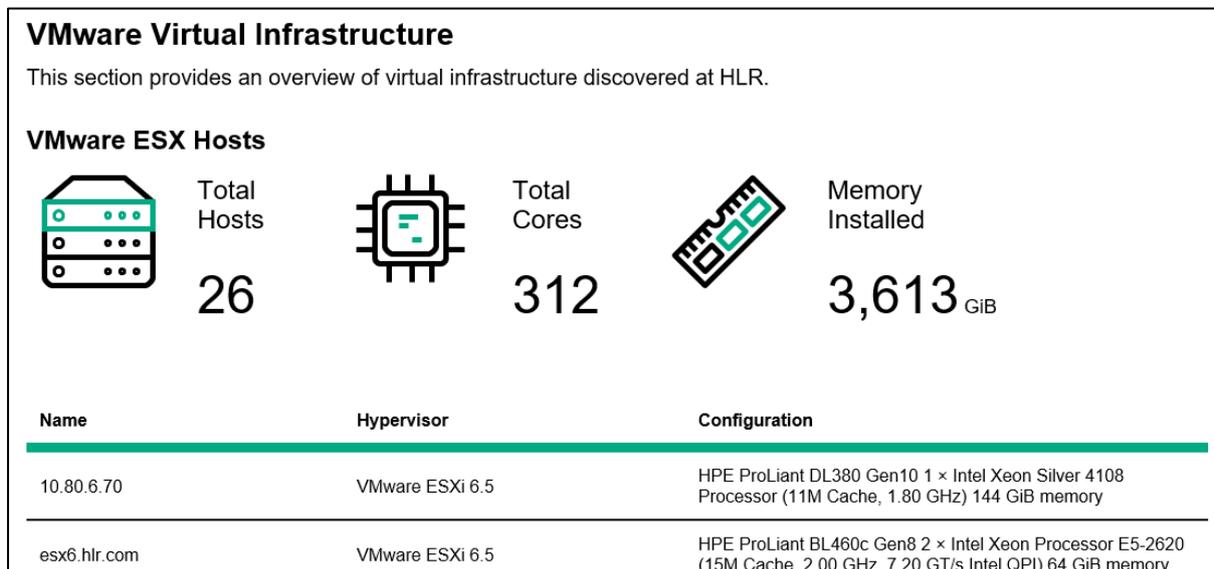


Figure 1

Lanamark One also provides a detailed spreadsheet containing the raw data, shown in Figure 2 below. The spreadsheet provides more detailed metrics than formatted document. The spreadsheet can be filtered if some of the virtual machines need to be excluded from the sizing exercise – for example if they are temporary test machines.

VMware Virtual Machines						
						State
Name	vCenter	Host	Cluster	Power	Connection	
HPEOneView-DCS_4.00-CP-synergy	vc6.hlr.com	esx2.hlr.com	infrastructure_a	Powered Off	Connected	
AD demopool.local	vc6.hlr.com	esx6.hlr.com	infrastructure_a	Powered Off	Connected	
HPEOneView_3.10.04_299553_GHLWr	vc6.hlr.com	esx1.hlr.com	infrastructure_a	Powered Off	Connected	
ITA_001	vc6.hlr.com	esx1.hlr.com	Cluster NY	Powered Off	Connected	
PLDC	vc6.hlr.com	esx2.hlr.com	infrastructure_a	Powered Off	Connected	
VCSIM 3	vc6.hlr.com	esx1.hlr.com	infrastructure_a	Powered Off	Connected	
Veeam9 Demo	vc6.hlr.com	esx1.hlr.com	infrastructure_a	Powered On	Connected	
OneView410	vc6.hlr.com	esx1.hlr.com	Cluster NY	Powered On	Connected	
Cluster Management Utility 8.1	vc6.hlr.com	esx3.hlr.com	infrastructure_a	Powered Off	Connected	
CTXWS2012r2-016	vc6.hlr.com	esx4.hlr.com	infrastructure_a	Powered Off	Connected	
Nagios	vc6.hlr.com	esx4.hlr.com	infrastructure_a	Powered Off	Connected	

Figure 2

The online HPE SimpliVity Sizing tool requires detailed input to size a solution. If one of the fields is left blank then that metric won't be used in the sizing, and therefore the end results would be skewed. This is very likely to result in the solution being over specified. Figure 3 below shows the sizing tool in basic mode, which is the default option when using the tool.

Sizing: HLR

Compression Ratio: Dedupe Ratio:

Cluster: SVT1

Name	VM Count	Total vCPU Cores	V:P CPU Ratio	Peak pCPU (GHz)	Allocated Memory (GiB)	Peak Capacity (GiB)	95% Peak Storage IOPS	Latency Tolerance (ms)
workload1	519	1798	6 :1	103	4339	82861	8656	30

Add VM Group

Cluster Growth: 0% 0% 0%

External Compute: 0GHz 0GiB

Compute HA: N+1

Product Family:

Subtotal	519	1798	6.0	103	4339	82861	8656	30
-----------------	------------	-------------	------------	------------	-------------	--------------	-------------	-----------

Figure 3

The table below summarises how to find each value from the Lanamark One spreadsheet which the HPE SimpliVity Sizing tool requires (Figure 1).

Sizing Tool Metric	Data Source
Compression Ratio	The value we select here depends on how compressible the data is. It is possible the current infrastructure has statistics which can inform us here. Otherwise, the recommendation is to leave at the default of 1.5:1
Dedupe Ratio	As most SimpliVity customers see a 1.5:1 dedupe rate, we can change the setting on 1.5:1
VM Count	<p><i>Guest OS Instances</i> worksheet -> the total of entries in the <i>Name</i> column.</p> <p><i>Example formula</i> =SUBTOTAL(103,'Guest OS Instances'!A4:A999)</p>
Total vCPU Cores	<p><i>Guest OS Instances</i> worksheet -> the total of entries in the <i>CPUs</i> column.</p> <p><i>Example formula</i> =SUBTOTAL(109,'Guest OS Instances'!E4:E999)</p>
V:P CPU Ratio	<p><i>VMware ESXi Hosts</i> worksheet -> A total derived from: <i>CPUs</i> multiplied by <i>Cores/CPU</i>, divided by the Total vCPU Cores value recorded above.</p> <p>This would give the ratio from the current environment. When sizing a SimpliVity cluster, a ratio of 4:1 or 5:1 is typical. For critical VMs a lower ratio may be required.</p> <p><i>Example formula for physical cores</i> =SUMPRODUCT('VMware ESXi Hosts'!O4:O29,'VMware ESXi Hosts'!N4:N29)</p>
95% Peak pCPU (GHz)	<p><i>VMware ESXi Hosts</i> worksheet -> A total derived from: the <i>CPUs</i> column multiplied by <i>Cores/CPU column</i>, then multiplied by <i>Clock Ghz</i>, and then multiplied by the percentage value in the <i>Processor (95th)</i> column.</p> <p><i>Example formula</i> =SUMPRODUCT('VMware ESXi Hosts'!N4:N30,'VMware ESXi Hosts'!O4:O30,'VMware ESXi Hosts'!P4:P30,'VMware ESXi Hosts'!U4:U30)</p>

Allocated Memory (GiB)	<p><i>Guest OS Instances</i> worksheet -> the total of the <i>Memory Allocated (GiB)</i> column.</p> <p><i>Example formula</i> =SUBTOTAL(109,'Guest OS Instances'!L4:L999)</p>
95% Peak Capacity (GiB)	<p>VMware Virtual Machines worksheet -> the total of the <i>Storage Capacity (GiB) Vdisk</i> column.</p> <p>This is a more useful than the <i>VMware datastores worksheet</i>. The <i>VMware datastores</i> worksheet would exclude raw device mapped LUNs and would include storage efficiencies from the current infrastructure.</p> <p><i>Example formula</i> =SUBTOTAL(109,'VMware Virtual Machines'!Y4:Y999)</p>
95% Peak Storage IOPS	<p><i>VMWare Virtual Machines</i> worksheet -> the total of both <i>I/OPS (95th) Read</i> and <i>I/OPS (95th) Write</i> added together.</p> <p>When using the SimpliVity Sizing tool and OCA you are required to choose 4000 series or 6000 series SSDs, depending on whether the workload is read or write heavy. Using the IOPs figures we have just gathered we can calculate the percentage of reads and writes:</p> <p>Total of <i>I/OPS (95th) Read</i> divided by Total of all IOPs calculated above = percentage of read operations.</p> <p>Total of <i>I/OPS (95th) Write</i> divided by Total of all IOPs calculated above = percentage of write operations.</p> <p><i>Example formula</i> =SUM('VMware Virtual Machines'!AD4:AD522,'VMware Virtual Machines'!AE4:AE522)</p>
Latency Tolerance	<p>This metric is not calculated by Lanamark One. The Sizing tool is effectively asking you what latency tolerance is acceptable to the customer. The default value of 30ms seems rather slow for most application workloads, so you may want to consider 10ms or less.</p>

Figure 4

VM Groups

The HPE SimpliVity Sizing tool allows the user to create additional VM groups. Each VM group is essentially a workload, for example VM group might be e-mail servers or file servers. By dividing the virtual machines in this way, the sizing requirements for each application service can be seen more clearly. A suitable backup policy can then be configured for those services using the advanced mode option.

Advanced Mode

Advanced mode is an option in the HPE SimpliVity Sizing tool which is used to add local backup sizing into our HPE SimpliVity solution. The first step in this process is to create a backup policy. The policy defines backup schedules and retention. Different backup policies can be applied to different VM groups. An example backup policy is show in Figure 5 below.

In the advanced mode main window, the backup policy can be selected as well as the rate of data change for the virtual machines (Figure 6).

Backup Policies		
☰ Standard Backup Policy		
Frequency	Retain	For
Hourly	24	Hour(s)
Daily	7	Day(s)
Weekly	4	Week(s)
Monthly	12	Month(s)

Figure 5

Hourly Change Rate (%)	Daily Change Rate (%)	Weekly Change Rate (%)	Monthly Change Rate (%)	Yearly Change Rate (%)	Local Backup Policy
0.5	1.5	8	20	0	Standard Backup Policy ▾

Figure 6

Determining accurate values for these fields requires an analysis of the current infrastructure. It is better to estimate a value rather than leave the field blank.

The sizing tool allows the user to create a second cluster and specify the backups which are to be replicated to that second cluster. Note however, that the tool does not automatically size the second cluster to include the backup values from the primary cluster. The additional capacity that the backups create need to be manually added to the second cluster's capacity value.

References

HPE SimpliVity Sizing Quick Start Guide

Understanding VMware Virtual Machine storage and Virtual Datastores (web article found at <http://help.lanamark.one/lanamark-one/understanding-vmware-virtual-machine-storage-and-virtual-datastores>)

HPE SimpliVity Online Sizing Tool Presentation by Tim Antonowicz