

“Getting Started with HPE Morpheus VM Essentials”

The Ultimate Guide v1.0

By Bart Heungens

Table of Contents

TABLE OF CONTENTS	2
WHY THIS GUIDE	4
WHAT IS HPE MORPHEUS VM ESSENTIALS? WHAT HAS MORPHEUS TO DO WITH IT?	5
HPE MORPHEUS ENTERPRISE	5
HPE VM ESSENTIALS	8
TECHNICAL ARCHITECTURE	9
TERMINOLOGY	11
IS HPE VM ESSENTIALS ENTERPRISE READY?	13
WHERE TO FIND THE LATEST DOCUMENTATION AND FEATURES LIST?	14
SOFTWARE RELEASE NOTES.....	14
SOFTWARE DOCUMENTATION GUIDE	16
DEPLOYMENT GUIDE	17
QUALIFICATION MATRIX.....	18
HPE MORPHEUS ENTERPRISE	18
WHICH HARDWARE IS SUPPORTED? WHAT ABOUT VEEAM AND OTHER ISV SUPPORT?	19
WHAT ABOUT NETWORKING?	23
4 NICs – CONVERGED MANAGEMENT & VM TRAFFIC – ISCSI STORAGE – NATIVE VLANs.....	26
4 NICs – CONVERGED MANAGEMENT & VM TRAFFIC – ISCSI STORAGE – NO NATIVE VLAN	28
6 NICs – SEPARATE MANAGEMENT, VM & ISCSI STORAGE TRAFFIC – NO NATIVE VLAN	31
HOW TO GET STARTED? WHERE TO DOWNLOAD?	34
WHERE CAN I FIND FEATURES LIKE VMOTION? STORAGE-VMOTION? AFFINITY? HA & DRS? MAINTENANCE MODE?	42
INSTANCE VERSUS VIRTUAL MACHINE.....	42
EDIT A VIRTUAL MACHINE	43
VMOTION.....	45

STORAGE-VMOTION	45
HA & DRS	46
AFFINITY RULES	47
MAINTENANCE MODE.....	49
HOW TO SUCCESSFULLY DEPLOY A WINDOWS VM ON A KVM-BASED HYPERVISOR.....	50
MIGRATE VMWARE VM'S TO HPE VM ESSENTIALS	55
HPE BULK MIGRATION TOOL.....	56
HPE ZERTO AS A MIGRATION TOOL	58
RIVERMEADOW	59
OTHERS (VEEAM – STARWIND – ...)	61
HOW TO UPDATE HPE VM ESSENTIALS?	62
IS HPE VM ESSENTIALS SUPPORTED ON SIMPLIVITY?	66
USEFUL LINUX COMMANDS	70
VIRSH.....	70
GENERAL LINUX COMMANDS.....	79
USEFUL LINKS.....	87

Why this guide

I am Bart Heungens, founder of an IT consultancy company called Bitcon (nothing to do with bitcoins!) based in Belgium, and have built +30 years of experience in on premises IT infrastructure solutions. In my role I provide consultancy and training services to companies worldwide, mainly on HPE solutions and the extended HPE ecosystem, and deliver presales services and partner enablement for TD SYNnex.

This document is based on experiences and feedback that I got in this role in the past years when talking with virtualization users, ranging from 5 to 50.000 virtual machines. SMB to large enterprise... And all the challenges they see the past months and years with some major changes in the virtualization world with the acquisition of VMware by Broadcom.

In these discussions with C-level and technical teams I noticed a pattern of questions and concerns that kept coming back, independent of the size of the virtualized environment.

In this document I will give **my thoughts and answers on these questions and concerns**, and where all the available information can be found to proceed and succeed.

The steps and information in this guide are successfully used in the last year with several HPE partners and end customers worldwide, but cannot guarantee any success for your specific situation. Feel free to contact us for specific needs and support.

See it as a guideline with most of the best practices learned from the field, and will be kept updated when these practices change or new versions are released.

This version of the document is based on **HPE Morpheus VM Essentials v8.1.0** and will be updated accordingly over time. Be sure that U always use the latest version of this guide for the latest and greatest information.

Screen captures can change over time due to changes on the HPE website and in the product.

I hope this guide will help you in your journey with HPE Morpheus VM Essentials!



Bart Heungens

Blogger – Consultant – Cloud Architect –
Trainer – Storyteller

at www.hpepedia.com and www.bitcon.be

What is HPE Morpheus VM Essentials? What has Morpheus to do with it?

When I start my pitch on HPE Morpheus VM Essentials (which I will call from now on VM Essentials or even VME for the easiness), I always start with explaining what Morpheus has to do in the complete picture, because this is very important to know, for instance for my second bullet point in this guide (see further).

HPE acquired Morpheus Data, a US based software company founded in 2010, in the summer of 2024.

Morpheus developed through all those years a cloud automation and orchestration platform that helps customers and service providers create a customizable cloud service portal for their private and public clouds.

HPE Morpheus Enterprise

HPE Morpheus Enterprise is an application-centric abstraction framework with the ability to deliver rapid value and extreme flexibility through 90+ built-in codeless integrations and plugins, and the flexibility to integrate with +20 public and private clouds.

Provision Any Workload Into Any Cloud On-Demand

Orchestrate all the tools and dependencies for catalog deployment and lifecycle management of VMs (IaaS / VMaaS), containers (CaaS), app services (PaaS), multi-tier app stacks, and cloud services.

Simplify Runtime Deployment for Flexibility and Choice

Push button deployment and management of **Virtualization (VME) and Kubernetes (MKS) Clusters** plus cloud-managed EKS, AKS, GKE. Extend with brownfield of 3rd Party K8s like SUSE, OpenShift, etc.

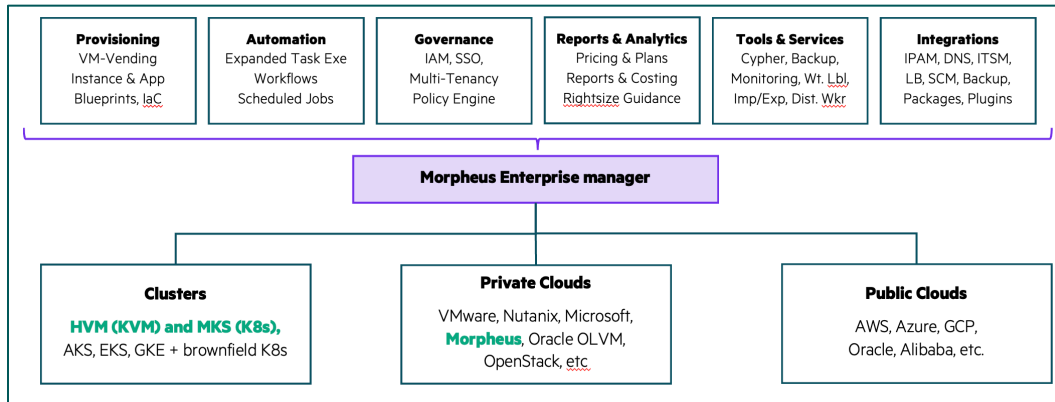
Lower Cloud Cost with Visibility and Optimization

Discover on-prem and public cloud instances, optimize resource use with analytics and rightsizing, sync cloud invoices, enforce budget policy, enable show back, and centralize reporting.



This platform is still available today under the HPE Morpheus Enterprise branding.

The acquisition was important for HPE since it was using the Morpheus control platform for some years in its HPE GreenLake cloud platform, just like many other cloud providers worldwide.



The core of the platform is the Morpheus Enterprise Manager that is deployed as a single VM at start but can scale to a multi-tier control plane.

Very similar to what the VMware VCF stack does, controlling one or more clusters of hypervisors (for VM's) and runtimes (for containers), and a mix of private and public clouds.

It provides a provisioning layer to deploy so called instances, which can be virtual machines, containers or bare metal deployments, and can be extended to App Blueprints representing multi-instance deployments for IaC environments. Provisioning can be triggered via GUI (Standard or Catalog), API/CLI, Terraform Provider, or ITSM plugin (ServiceNow). Service Catalogs can be used for 'easy button' provisioning of instances, apps, clusters, and workflows.

The automation part is quite extended in the Morpheus platform, since that is actually the root of why this platform was created at start.

Tasks are primarily created for use in Workflows, but a single Task can be executed on an existing instance. The automation includes scripting languages such as Ansible, Chef, HTTP/REST, Puppet, PowerShell, Bash, Python, vRealize Orchestrator, Terraform, etc.

Governance features like multi-tenancy and RBAC – Groups, Roles, Users, Policies.

Tenants are isolated environments with unique users and workloads. One master tenant per installation; all others are sub-tenants. Master and sub-tenants can connect to identity providers where you can setup roles and map groups of users.

Groups are a logical wrapper where resources meet RBAC. Clouds and cloud-related resources are added to groups. Virtually every feature, cloud, catalog item, etc is subject to RBAC enabling thousands of combinations to meet specific customer needs.

Reports and Analytics as well as 'FinOps' features for cloud cost management are by default part of the HPE Morpheus Enterprise solution.

Service plans and price sets are components of 't-shirt sizes' (cores, mem, storage, etc.) associated with Instances and Applications. Prices support different units (minute, hour, day, month, year, etc.) and Currency. Public cloud costs are synced from the provider based on the attached cloud accounts.

Miscellaneous tools and services are available to help when orchestrating hybrid cloud management.

Cypher is secure Key/Value store for encrypting credentials and other info used in instances, automation, etc.

Backup is a built-in basic data protection feature to provide for VM, container, host, files, DB, etc. It covers snapshots and synthetic full backups. It is not a replacement for external integration with Enterprise backup vendors like CommVault, Cohesity, Veeam, which is by default part of the solution.

Monitoring is built-in basic features for creating a 'check' for health of provisioned instances & apps including ServiceNow integration. It is not a replacement for external integration with Enterprise monitoring tools like OpsRamp, DataDog, etc.

Types of integrations exist for Automation, Backups, Clouds – Private and Public, Containers, Deployment (Git, GitHub), DNS, Identity Management, ITSM, Load Balancers, Logs, Monitoring, Networking, Storage and Trust. Let this be the strength of the platform.



Conclusion: HPE Morpheus Enterprise is the Swiss army knife that can help you in every situation! It can do it (almost) all!

And this is where VM Essentials came in the picture.

HPE VM Essentials

Within the HPE Morpheus Enterprise platform, Morpheus used their own hypervisor called **MVM**, based on Linux KVM hypervisor, to deploy virtual machines in their own IT environment.

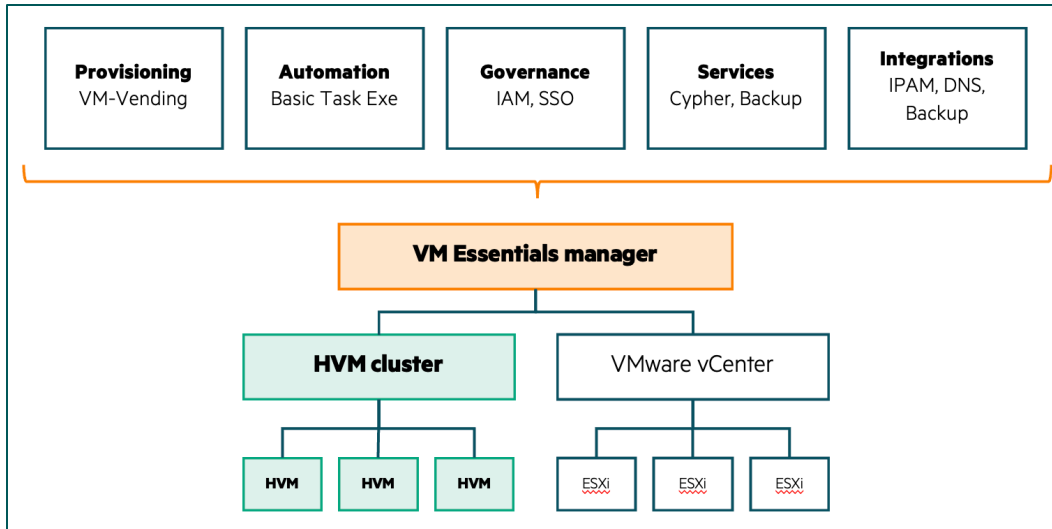
They never commercialized it, since they would compete with the other hypervisors on the market at that moment, being VMware, Nutanix, Hyper-V, and so on...

Enter Broadcom! Unless you lived under a rock the past 2 years you know what happened with the go to market strategy with the VMware software solution stack by Broadcom.

The entire world started looking for alternatives for the rising costs of the VMware hypervisor and vCenter stack.

Enter Morpheus MVM! Actually, HPE did not have to create their own hypervisor and management stack, since they owned it thanks to the Morpheus Data acquisition.

Morpheus MVM became the foundation of the HVM hypervisor that runs on each host in the cluster, and to replace the vCenter functionality and they took the Swiss army knife and cut off all the advanced features and functions of HPE Morpheus Enterprise that most of the customers do not require (like Terraform automation and all public cloud integrations) and kept only the essentials to run virtual machines...



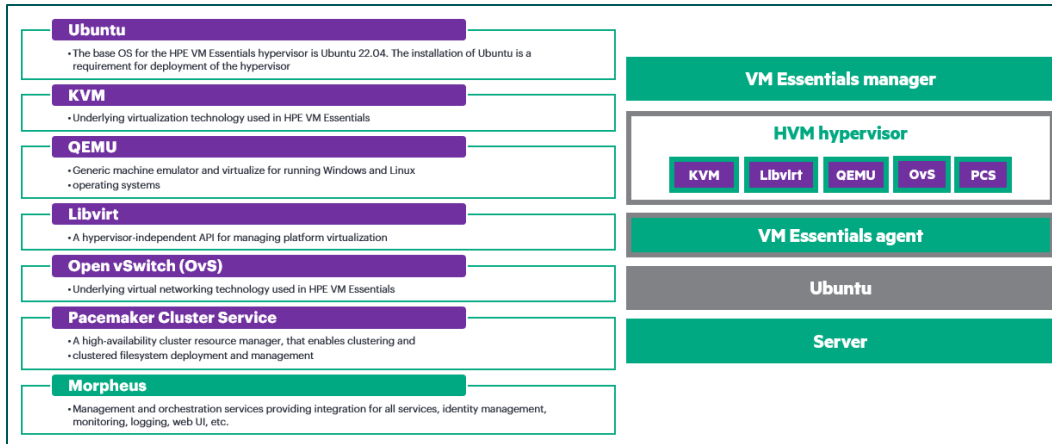
Enter HPE Morpheus VM Essentials!

And this is directly the answer to the next topic in this guide, being the enterprise readiness of VM Essentials... It exists already +10 years... Version 8.1.0 at the release of this guide.

Technical architecture

HPE VM Essentials is built upon established and proven opensource technologies like KVM, QEMU, LibVirt OvS and Pacemaker Cluster Service, and industry leading Morpheus Enterprise suite, with a feature set focused upon virtualization management and consumption.

The management software that powers HPE Morpheus VM Essentials Software is the same management software that powers Morpheus Enterprise. The license that is applied determines what features are exposed.

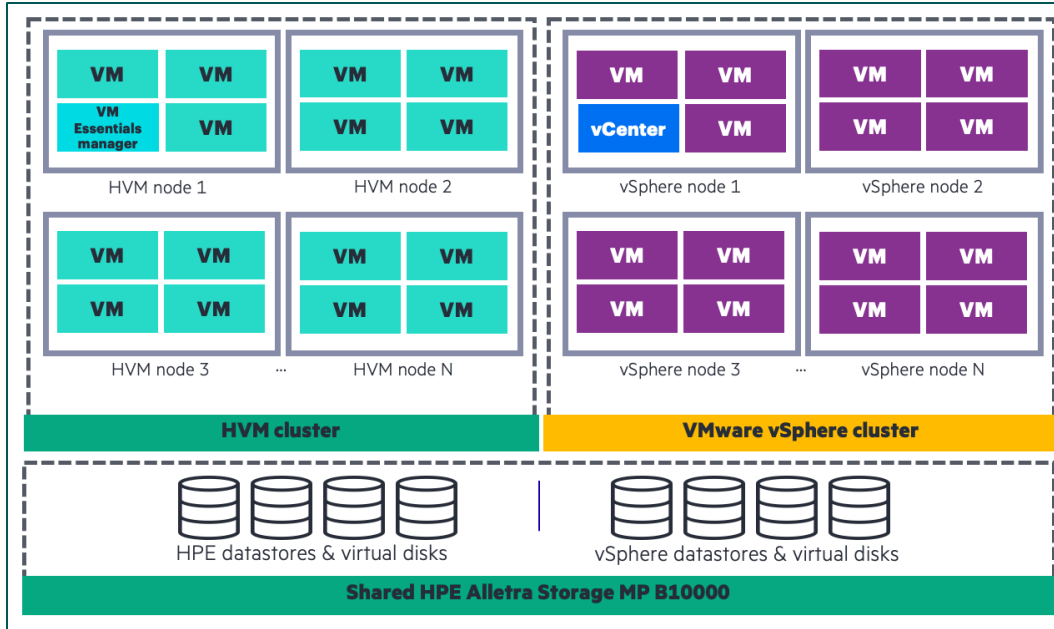


From a technical point of view, VM Essentials is very similar to a VMware deployment.

You deploy a set of hosts in a so-called cluster delivering features like high availability (see later for more details), which is managed by a manager function which runs as a virtual machine on that cluster, or outside the cluster.

The VM Essentials Manager is what we discussed above the HPE Morpheus Enterprise software, but with the features enabled just for virtual machines.

Nice thing to know about the VM Essentials Manager is that you can manage not only the HVM cluster but also can connect to the VMware vCenter server in your IT environment and can do the day2day operations from your VMware stack from within the same VM Essentials Manager console.



Terminology

HPE created this nice overview on the terminology they are using in the HPE Morpheus VM Essentials software stack, and how it compares with the VMware stack.

Familiarize yourself with it!

Description		HPE	VMware
Orchestration & automation solution which includes virtualization	The full solution that adds orchestration, governance, FinOps, in addition to virtualization	HPE Morpheus Enterprise Software	VMware Cloud Foundation / Aria Automation
Virtualization solution	The solution that provides a platform for running virtual machine workloads	HPE Morpheus VM Essentials Software	VMware vSphere
Virtualization management server	The management server that enables advanced functionality	VM Essentials manager	vCenter
Virtualization cluster	A group of hypervisors that are combined to provide an aggregated pool of resources	HVM cluster	vSphere/ESXi cluster
Virtualization hypervisor server	The physical host used to host virtual machines	HVM host	ESXi host
Hypervisor itself	The hypervisor	HVM hypervisor	ESXi

Most important to remember, and what will be referenced in the guide are the hypervisor (HVM versus ESXi) and the control plane (VME Manager versus vCenter).

Is HPE VM Essentials enterprise ready?

Easy answer: YES !

I got this question quite often from customers, asking me about deploying their mission critical virtual machines on a generation 1 product...

Well, HPE is using KVM as the hypervisor, which exist +20 years, and is also being by other leading vendors like Nutanix, so I don't see the real issue there. It I just the hypervisor.

The big challenge in the virtualization technology is the control plane. For this HPE uses HPE Morpheus Enterprise, which is version 8.1.0 at the moment of writing of this guide, which is definitely not version 1.x...

Knowing that multiple large enterprises and service providers worldwide, including HPE GreenLake, is using HPE Morpheus Enterprise code for the automation and orchestration of their large public and private cloud deployments, I cannot imagine that HPE would take that risk to put VM Essentials on the market if it was not enterprise ready.

Other side note: HPE transitioned all its private cloud offerings Private Cloud Business Edition (**PCBE** – for virtual machine workloads), Private Cloud Enterprise (**PCE** – for virtual machines, containerized and bare metal workloads) and Private Cloud AI (**PCAI** – for AI workloads) to use VM Essentials as hypervisor and not mandatory VMware anymore...

HPE would not do this if it had no confidence in the entire software stack.

Further from the 8.1.0 release notes:

- There is now a single binary installer for both HPE Morpheus Enterprise Software and HPE Morpheus VM Essentials Software. Feature access is controlled solely by licensing. Previously, HPE Morpheus Enterprise and HPE Morpheus VM Essentials had separate binary packages

My point: HPE VM Essentials is HPE Morpheus Enterprise, but just with some features turned off... Cool!

Another proof that this is not a generation 1 product: continue to the next topic on documentation...

A generation 1 product would not have a that detailed and in-depth documentation set available like it is now for the HPE Morpheus Enterprise stack. Love it!

Where to find the latest documentation and features list?

The one and only place to get started is this URL, and that you should first in your browser favorites:

<https://www.hpe.com/support/vme-docs>

Support Center HPE Morpheus VM Essentials Software

Access product support documents and manuals, software, download drivers by operating environment, and view product support videos.

Manuals Knowledge Articles Drivers and Software Alerts Videos Feedback

Find in Manuals

Sign up for Product Alerts

Change Product

All Management/Operations Parts/Specifications Reference Information

You are viewing the most relevant and current results for this product.

Showing 1-10 of 31

Subtype	Title	Date
Management/Operations Guide	HPE Morpheus VM Essentials Software Documentation v8.1.0	Mar 16, 2026
Management/Operations Guide	HPE Morpheus VM Essentials Software Release Notes v8.1.0	Mar 15, 2026
Reference Guide	Qualification Matrix for HPE Morpheus VM Essentials Software	Mar 07, 2026
Release Notes	HPE Storage Integration Pack for VM Essentials 8.1.0 Release Notes	Mar 04, 2026
Management/Operations Guide	HPE Morpheus VM Essentials Deployment Guide	Feb 15, 2026
Management/Operations Guide	HPE Morpheus VM Essentials Software Release Notes v8.0.13	Feb 09, 2026
Management/Operations Guide	HPE Morpheus VM Essentials Software Documentation v8.0.13	Feb 08, 2026
Management/Operations Guide	HPE Storage Integration Pack for VM Essentials 8.0.13 Release Notes	Feb 03, 2026
QuickSpecs	HPE Morpheus VM Essentials Software QuickSpecs	Jan 26, 2026
Management/Operations Guide	HPE Morpheus VM Essentials Software Documentation v8.0.12	Jan 10, 2026

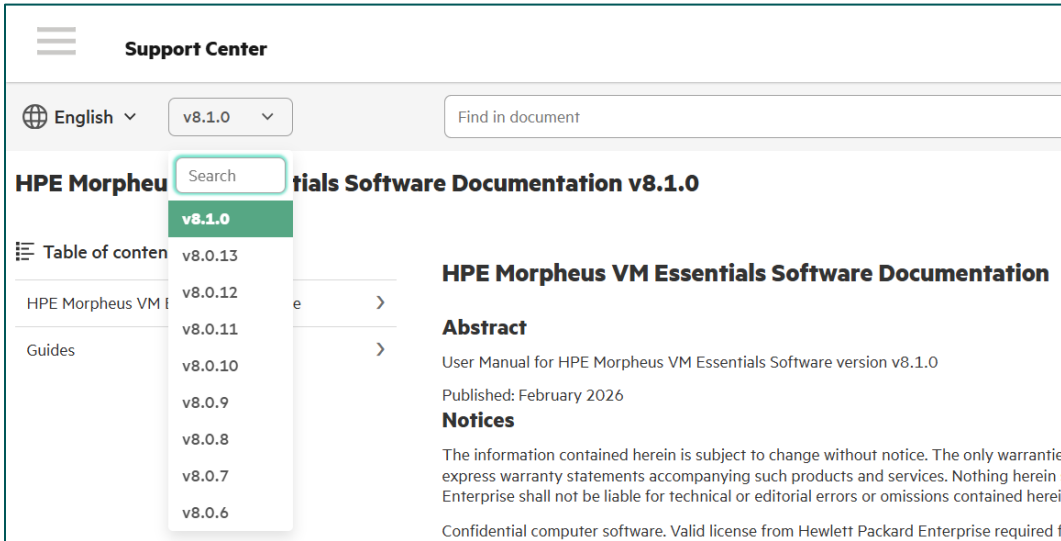
On this webpage you can find all the latest information and documentation guides on VM Essentials.

4 important documents here:

Software Release Notes

The first document to look out for are the Release Notes. Here you will find the latest version being released (8.1.0 at time of writing this guide), with all the updates that are available in that specific release.

A dropdown box allows you easily to switch between releases and look for all the updates over time.



2025 was a busy year for the HPE Morpheus team, where every month a new software update was released. Naming is quite easy; the last digit is/was the month of release. The important update 8.0.8 (with the Migration Tool and the Integrated Installer) was released in August 2025.

From 2026 on the update/upgrade plan will be a bit different, going to quarterly updates.

HPE Morpheus Software Release Schedule

Release Window	7.x.x	8.x.x	9.x.x	10.x.x
January 2026	Maintenance	8.0.13		
February 2026		8.1.0		
March 2026		8.1.1		
April 2026		8.1.2		
May 2026	EOL	Maintenance	9.0.0	
June 2026			9.0.1	
July 2026			9.0.2	
August 2026			9.1.0	
September 2026			9.1.1	
Oct '26 – Apr '27			9.1.2 – 9.3.2	
May 2027		EOL	Maintenance	10.0.0
June 2027				10.0.1
July 2027				10.0.2
August 2027				10.1.0

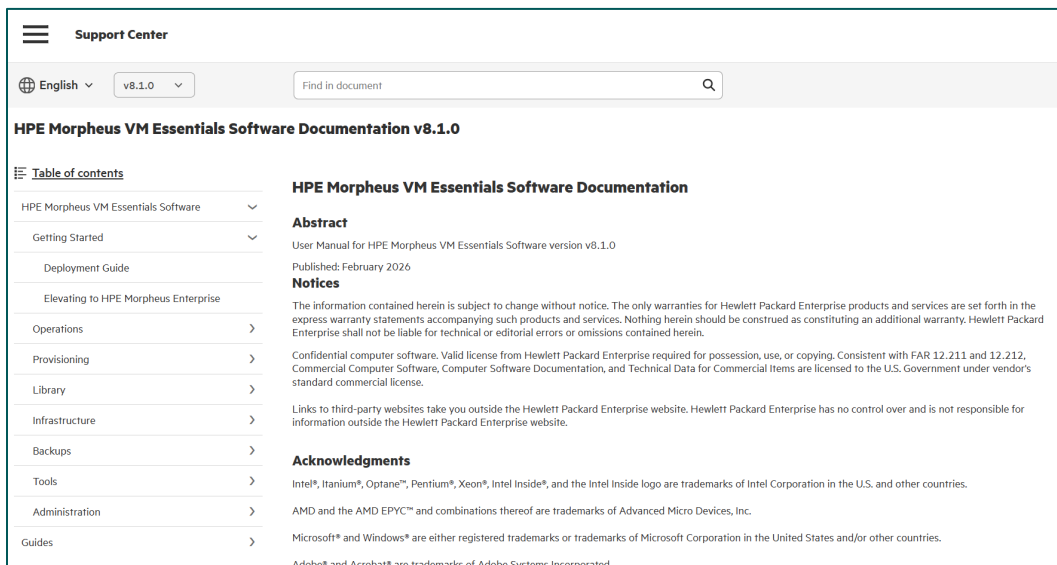
HPE Morpheus Software will support each major version for a period of two years. This support window is split between a twelve-month period of active development followed by another twelve-month period of maintenance support.

During active maintenance, an upgrade containing new features is released once per quarter (**version x.1.x, x.2.x**, etc). Each month between the quarterly feature upgrades, HPE will deliver a maintenance upgrade containing improvements and bug fixes (**version x.x.1 and x.x.2**). After this twelve-month period of active development ends, a new major upgrade is delivered (**version 9.x.x, 10.x.x**, etc) and a twelve-month period of maintenance support begins. During maintenance support, there are no scheduled monthly upgrades. However, support services are still available and critical security upgrades will be delivered should the need arise. After this twelve-month period of maintenance support ends, all versions within that major release umbrella will reach end of life.

Software Documentation Guide

The most important documentation is the software documentation guide.

The tree structure at the left gives you a good guideline of how to get started with VME, and then all subsections from the management UI on Provisioning, Library, Infrastructure, Backups, Tools and Administration.



Support Center

English v8.1.0 Find in document

HPE Morpheus VM Essentials Software Documentation v8.1.0

Table of contents

- HPE Morpheus VM Essentials Software
- Getting Started
- Deployment Guide
- Elevating to HPE Morpheus Enterprise
- Operations
- Provisioning
- Library
- Infrastructure
- Backups
- Tools
- Administration
- Guides

HPE Morpheus VM Essentials Software Documentation

Abstract

User Manual for HPE Morpheus VM Essentials Software version v8.1.0
Published: February 2026

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

AMD and the AMD EPYC™ and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

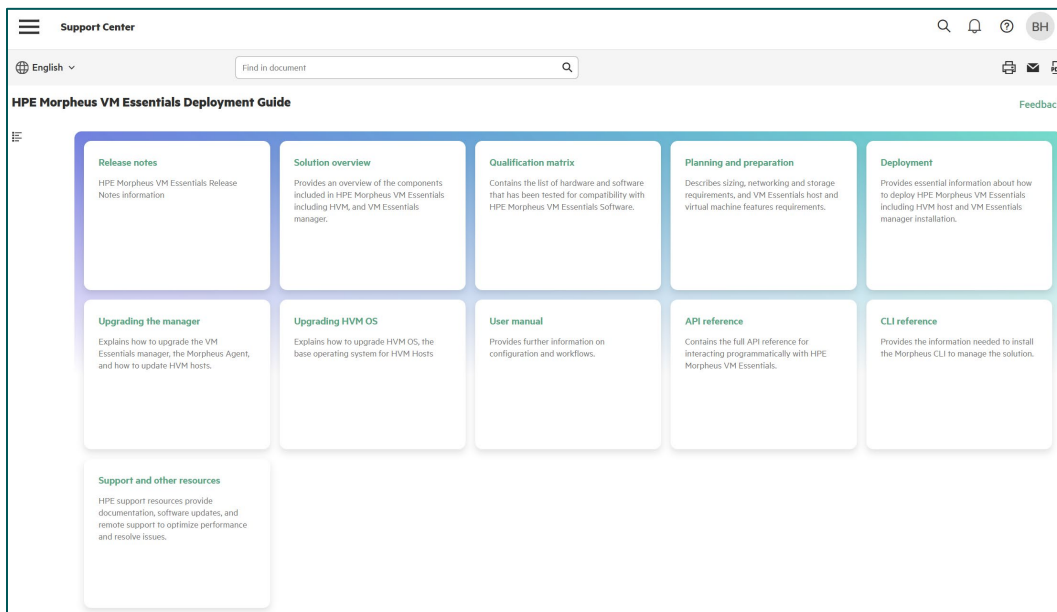
Like I mentioned in the introduction of this guide, I will not rewrite the entire software documentation guide here (since everything is well documented there) but will cover all the most important topics and steps for you to be successful in sizing, installing, configuring, demo-ing and troubleshooting VM Essentials.

But the latest and most up-to-date information will be found there!

Notice that in previous versions of the Software Documentation Guide the installation and configuration section was part of this guide. HPE has created now a new dedicated section specific for the deployment and initial configuration of HPE VM Essentials, which can be found in the HPE Morpheus VM Essentials Deployment Guide.

Deployment Guide

The place to get started with VM Essentials is the Deployment Guide.



Multiple tiles guide, from a general overview, the qualification matrix (see below), planning and preparation all the way to how to install and configure VM Essentials, and shows as well the steps of how to upgrade VM Essentials to a newer version.

Useful addition in the documentation set for sure...

Qualification Matrix

The Qualification Matrix (originally called the Compatibility Matrix) gives the list of supported hardware and software for the most recent version of released software.

Check the next topic in this guide for more details.

HPE Morpheus Enterprise

Know that there is also a vanity URL for the HPE Morpheus Enterprise documentation, where you can find more in-depth information and documentation, specifically on the integrations part of the software stack...

<https://docs.morpheusdata.com>

If you decide one day to do the upgrade to Morpheus Enterprise, this URL will tell you all the latest and greatest.

Also if you are looking for more details on API and CLI options, this is the place to be.

Which hardware is supported? What about Veeam and other ISV support?

HPE maintains a Software Compatibility Matrix document, which is recently changed to the Qualification Matrix, that is publicly available on the HPE Support website:

<https://www.hpe.com/support/VME-Compatibility-Matrix>

At launch only HPE ProLiant Gen11 servers were supported.

At the time of writing this guide Gen9, 10, 10+, 11 and 12 servers are supported, and also Synergy compute modules and even **Dell PowerEdge servers**.

Hardware Family	Platform Type	Hardware SKU	Comments
HPE ProLiant	AMD 2 RU server	HPE ProLiant DL145 Gen 11	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL320 Gen 12	
HPE ProLiant	AMD 1 RU server	HPE ProLiant DL325 Gen 12	
HPE ProLiant	AMD 2 RU server	HPE ProLiant DL345 Gen 12	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL360 Gen 9	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL360 Gen 10	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL360 Gen 10+	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL360 Gen 11	
HPE ProLiant	Intel 1 RU server	HPE ProLiant DL360 Gen 12	
HPE ProLiant	AMD 1 RU server	HPE ProLiant DL325 Gen 11	
HPE ProLiant	AMD 1 RU server	HPE ProLiant DL365 Gen 11	
HPE ProLiant	AMD 2 RU server	HPE ProLiant DL385 Gen 11	
HPE ProLiant	AMD 2 RU server	HPE ProLiant DL345 Gen 11	
HPE ProLiant	Intel 2 RU server	HPE ProLiant DL340 Gen 12	
HPE ProLiant	Intel 2 RU server	HPE ProLiant DL380 Gen 9	
HPE ProLiant	Intel 2 RU server	HPE ProLiant DL380 Gen 10	
HPE ProLiant	Intel 2 RU server	HPE ProLiant DL380 Gen 10+	
HPE ProLiant	Intel 2 RU server	HPE ProLiant DL380 Gen 11	

On the storage side, a complete series of HPE Storage Alletra 4000, 5000, 6000, 9000 and MP B10000 arrays are supported, and even the MSA Gen6 and Gen7 through native iSCSI, FC and even SAS connectivity.

Besides that, also some NetApp, Pure Storage and EMC storage arrays are supported as well. All details and requirements are mentioned in the Qualification Matrix.

Table 2. Storage Hardware Support			
Vendor	Hardware Family	Platform Type	Comments
HPE	HPE Alletra Storage	HPE Alletra 4000 (4100, 4120) (iSCSI)	Update storage hardware to the latest available firmware
HPE	HPE Alletra Storage	HPE Alletra 5000 (5010, 5030, 5050) (iSCSI)	Update storage hardware to the latest available firmware
HPE	HPE Alletra Storage	HPE Alletra 6000 (6010, 6030, 6050, 6070, 6090) (iSCSI)	Update storage hardware to the latest available firmware
HPE	HPE Alletra Storage	HPE Alletra 9000 (Fibre Channel)	Update storage hardware to the latest available firmware
HPE	HPE Alletra Storage	HPE Alletra Storage MP B10000 (iSCSI, Fibre Channel)	
HPE	HPE MSA Storage	HPE MSA Gen7 (2070, 2072) (iSCSI, Fibre Channel)	Update storage hardware to the latest available firmware
HPE	HPE MSA Storage	HPE MSA Gen6 (1060, 2060, 2062) (iSCSI, Fibre Channel, SAS)	Update storage hardware to the latest available firmware
HPE	HPE XP Storage	HPE XP8 (Fibre Channel, iSCSI)	Update storage hardware to the latest available firmware
NetApp	All-Flash FAS	AFF A-Series (all models)	iSCSI, FC, NFS; Minimum ONTAP version: 9.14
NetApp	All-Flash FAS	AFF C-Series (all models)	iSCSI, FC, NFS; Minimum ONTAP version: 9.14
Dell	Dell	PowerStore 4.0.0.2	iSCSI, FC, NFS
Pure Storage	FlashArray	FlashArray//X (all models)	iSCSI, FC, NFS; Minimum Purity version: FA 6.6.1
Pure Storage	FlashArray	FlashArray//XL (all models)	iSCSI, FC, NFS; Minimum Purity version: FA 6.6.1
Pure Storage	FlashArray	FlashArray//C (all models)	iSCSI, FC, NFS; Minimum Purity version: FA 6.7.1

The most asked question however is support from Veeam for the backup of the running VM's.

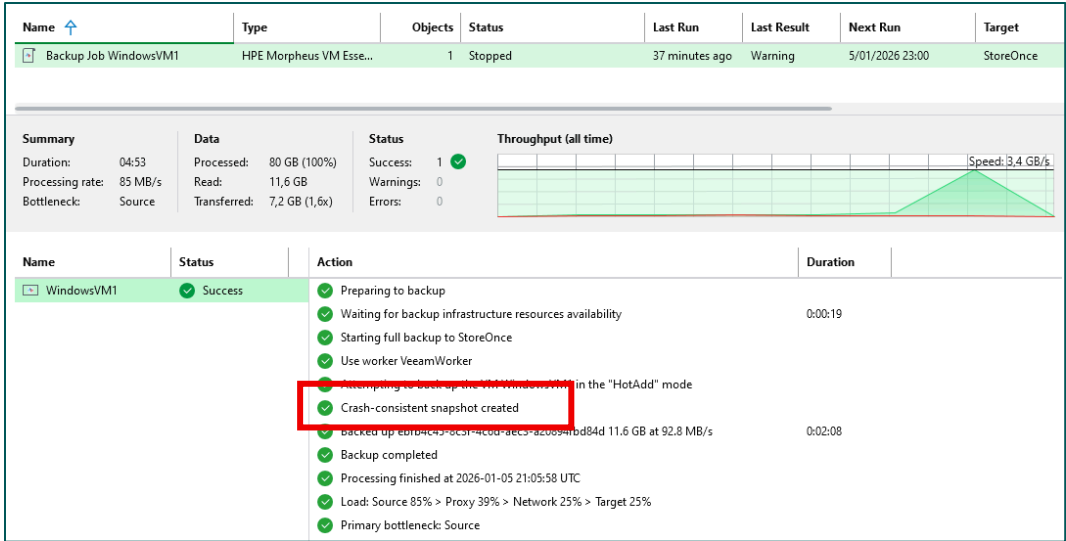
Short answer is that Veeam supports HPE VM Essentials already since several months by using the agent-based backups. I discovered a lot of objections of this 'limitation' that no image-based backups are supported, although it took several years for VMware ESXi to get this feature many years ago.

Veeam released a public beta version from Veeam 13.0.1 that supports VM Essentials with image-based backups with Changed Block Tracking (CBT) and since early March 2026 it is finally GA. Tests in my environment shows successful backups and restores from the VM's that I am running with snapshot integration, so successful agentless backups.

Further testing meanwhile shows that I am able to restore backups from VMware-based VM's on my VM Essentials cluster as well. So also the cross-hypervisor integration works as promised!

Other backup vendors like Commvault are on my to-do list to be tested and evaluated. Stay tuned.

The final objection to not get started with VM Essentials removed!



In the same Qualification Matrix you will find many other popular enterprise software solutions that support meanwhile HPE VM Essentials as hypervisor stack for their software. In the list you can find for instance Cohesity, Commvault, Microsoft SQL Server, Oracle, MongoDB, Elastic, Citrix, Omnissa (ex VMware Horizon), HP Anyware and Splunk, besides several others.

Always check the Qualification Matrix for the latest updates, since I see regular updates every few weeks.

Table 3. Independent Software Vendor (ISV) Support					
Partner	Product Name	Product Version	Deployment	Validation Type	Resources
Aerospike	Aerospike Enterprise Edition Database	8.1.0.2		HPE	Blog
Apache	Cassandra DB	5.0.6		HPE	Blog
Veeam	Backup and Replication	12.3	Agent-based	Partner	Technical Brief , Blog (part 1, part 2, part 3) , Video
CelerData	CelerData Enterprise (StarRocks)	3.4.4			
Cohesity	DataProtect	7.1.2 and later	Agent-based	Partner	Technical Brief , TekTalk-on-Point , Blog
Cohesity	NetBackup	11	Agent-based	Partner	Release Notes
Commvault ²	Commvault Cloud Backup and Recovery	11.40	Agent-Based, Image-based	Partner	Technical Brief , TekTalk-on-Point , Blog
HP	HP Anyware	25.03.1		HPE	Technical Paper
OpenText	OpenText Analytics Database (Vertica)	26 or 25.2		HPE	
Oracle ¹	Database	19c	Single instance only; Oracle RAC support TBD	HPE	Technical Brief , Blog , TekTalk-on-Point
Oracle	MySQL Community Edition	8.4.6			
Medical Informatics Corp. (MIC)	Sickbay Clinical Platform	3.45.4.0		HPE	
Microsoft ¹	SQL Server	SQL Server 2016, 2017, 2019, 2022	Single instance with Availability Groups	HPE	Technical Brief , Blog , TekTalk-on-Point
MongoDB ¹	Enterprise Advanced	8.0.0		HPE	Technical Brief , Blog , Video
Elastic ¹	Elastic Stack	9.0.0-1		HPE	Technical Paper
Citrix ^{1,3}	Citrix Virtual Apps and Desktops	7.2402 LTSR CU1		HPE	
Omnissa ^{1,3}	Horizon	8.13.1 (Build 11490723527)		HPE	Blog
Milestone	XProtect	Professional+, Expert, Corporate	2025 R2		
PostgreSQL Global Development Group	PostgreSQL	17.6			
SMInformation	ClusterPlex	5.0.5.5		Partner	ClusterPlex Support Matrix
Splunk	Splunk Enterprise	9.4.3	Distributed clustered deployment, single site	HPE	
Splunk	Splunk Enterprise SmartStore	9.4.3	Distributed clustered deployment, single site	HPE	
Thales Group	CipherTrust Manager	2.22.0 and 2.11.1		Partner	
Vali Cyber	ZeroLock	3.9.8			
Veeam	Veeam Backup & Replication	13.0.1	Image-based with Changed Block Tracking (CBT)	Partner	Blog , Video
Virtual Cable	UDS Enterprise	4.0 and above			

What about Networking?

I put on purpose the networking section before the actual installation section, since the networking part has a big impact on the **success or failure** of the installation of HPE VM Essentials... Don't blame HPE, but Ubuntu and KVM networking... Keep on reading.

I also want to highlight that I am not going to rewrite the Networking section from the Deployment Guide, which got a recent update/rewrite with better pictures and scenario's (thanks Sean, great job!), which I will use as a reference for this section in this guide. I am planning anyway another blog post specifically on VM Essentials networking. It will be based on my (still very popular) blog post on HPE SimpliVity networking in combination with VMware, but then updated to HPE VM Essentials networking. Stay tuned!

The short URL to the HPE VM Essentials documentation is already mentioned elsewhere in this document, but specifically for networking I want to give the direct link since it is too important in the installation and configuration process. Use this as a reference to get started and avoid to get into problems, like I did when starting to test with VM Essentials and refused to read the documentation... Would have saved a lot of time then and fewer re-deployments of Ubuntu OS...

https://support.hpe.com/hpesc/public/docDisplay?docId=sd00007332en_us&page=GUID-BDF3EFBF-EA90-4E6F-88D3-87ABF18C4170.html

The screenshot shows the HPE Support Center interface. At the top, there is a 'Support Center' header with a search bar and a language dropdown set to 'English'. Below this, the document title 'HPE Morpheus VM Essentials Deployment Guide' is displayed. A table of contents on the left lists sections like 'Solution overview', 'Qualification matrix', 'Planning and preparation', 'Host Requirements', 'VME Manager sizing', and 'Example Network Configurations'. The 'Example Network Configurations' section is expanded, showing a breadcrumb trail: 'Planning and preparation > Example Network Configurations'. The main content area is titled 'Example Network Configurations' and contains a paragraph of text: 'If you have the capability to do so, it's recommended you set up networking with full redundancy. Such a setup could include two network switches and hosts with at least four network interfaces spread across two network cards. This would allow for failover in the event of the loss of a switch and/or one of the host network cards in addition to separating management and compute network traffic to their own interfaces. Hosts with only four NICs each can still be configured for full redundancy but would have to converge management and compute traffic across the same interfaces. Keep in mind also that these examples use MPIO (multi-path input/output) for storage. It would also be possible to use bonding for storage depending on capabilities of the environment. MPIO is recommended whenever fibre channel or iSCSI LUNs are being used for GFS2 datastores.' Below this text, there are two subtopics: 'Scenario 1: Recommended Converged Networking Setup' and 'Scenario 2: Recommended Decoupled Networking Setup'. Navigation arrows are visible at the bottom right of the content area.

In order to run VM Essentials effectively in production, network redundancy and throughput must be considered. **Network bonding** is an important component to building redundancy into the environment so the networking section in the documentation guide is showing some example network configurations to explain the options and decisions that you will need to take and configuration changes that will need to be made.

Ultimately the environment is your own but this discussion and the example network configurations that follow will help in planning out an effective operating environment for VM Essentials.

During the installation of VM Essentials with the Unified Installer (or the DIY way by installing Ubuntu yourself and the VM Essentials Agent software afterwards) the first screen that you will get, and which is the most important one, is about networking.

```
Network configuration
Configure at least one interface this server can use to talk to other machines
access for updates.

NAME      TYPE  NOTES
[ bond0   bond  -           ▶ ]
  bond master for ens33, ens34

[ bond0.3 vlan  -           ▶ ]
  static   192.168.131.244/23
  VLAN 3 on interface bond0

[ ens33   eth   enslaved to bond0 ▶ ]
  00:50:56:93:2b:d6 / VMware / VMXNET3 Ethernet Controller

[ ens34   eth   enslaved to bond0 ▶ ]
  00:50:56:93:f7:06 / VMware / VMXNET3 Ethernet Controller

[ ens35   eth   -           ▶ ]
  DHCPv4   192.168.130.54/23
  00:50:56:93:f7:ec / VMware / VMXNET3 Ethernet Controller

[ ens36   eth   -           ▶ ]
  DHCPv4   192.168.130.53/23
  00:50:56:93:7b:1a / VMware / VMXNET3 Ethernet Controller

[ Create bond ▶ ]
```

You will see the 2, 4, 6 or even more interfaces from the hardware that you are using to install Ubuntu (in this picture ens18, ens19, ens20 and ens21) on. Know that these names can be different depending on the hardware that U use. So these might be different even

when using ProLiant hardware but Gen9, Gen10, Gen11 and so on... Always pay attention fir this already.

Typically you will see 2 networks in an HPE VM Essentials environment being the Management and Compute network. There might be a third Storage network when using iSCSI-based storage. Of course can you create more networks for your specific needs but these are the 3 2+1 typical ones...

Network bonding is the combining of multiple network adapters into a single logical interface, often called Network Teaming but in the Linux world called a bond. This is done to build in redundancy and to increase throughput. Network bonds are configured at the operating system level and there are multiple types of network bonds depending on hardware support and other factors.

First advice here is to create a bond **HERE**, so as early as possible in the installation & configuration process. You **CAN** do it afterwards, if you are good with VI editor and Netplan configuring, but if you are not the Linux expert for several years you can save yourself a lot of stress by setting up the bonding here.

Second advice: if you are testing or evaluating and want to keep it easy, then of course you don't need to use bonding, however you will see later in the installation process when deploying the VM Essentials Manager for instance, that you need to mention the interface to have it connected with. Or you know all the interface names of all your servers, and I hope they are the same across your hardware, or you just create a bond called Bond0 with just 1 interface connected to it, but it will save you a lot of potential issues later... Just saying.

So we create a Bond0 with 1 or more interfaces and need to set the type of Bond. HPE advices to use Active-Backup (when not using LACP on the switch side) or 802.3ad (LACP aka switch assisted load balancing) when using LACP on the switch.

For the ease of comparing how it works in VMware, this little comparison table of features and functions between both platforms:

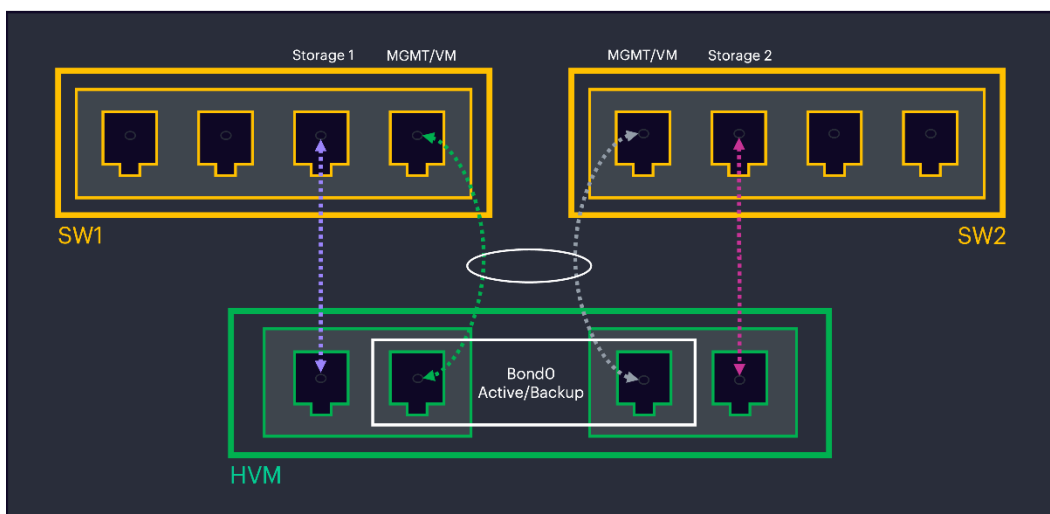
Feature	VMware LBT	Active-backup	802.3ad (LACP)
Switch config required	✗	✗	✓ Yes (LACP Support)
Outbound load balancing	✓ (NIC load based)	✗	✓ (Outbound + Inbound)
Inbound Failover	✗ (Depends on switch)	✓ (Instant failover)	✓ (Switch handles failover)
NIC Utilization	Dynamic	One Active NIC	Full duplex load balancing
Failover Speed	Fast (Outbound only)	Very Fast (link state)	Fast (Switch + Driver)
Complexity	Low	Very low	High (Switch + Host Config)
Best for	Balanced outbound traffic	Simplicity + Resilience	High performance switch-aware setups

In the Networking section of the Deployment Guide are three scenario's explained in detail of 'typical' networking setups that I will see in the field...

Let's start with the most common one:

4 NICs – converged Management & VM traffic – iSCSI storage – native VLANs

I see a lot of deployments already that will follow this schema: 4 NICs in the server, where 2 interfaces will be used for Management and VM traffic, and the 2 other interfaces for the iSCSI traffic.

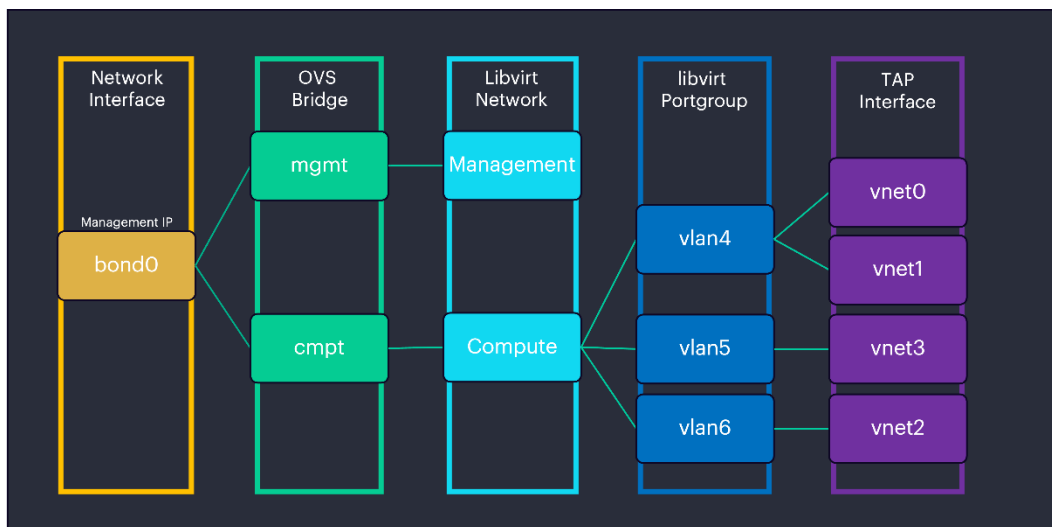


This is a logical visualization, or course it is better to have dedicated storage switches, but this does not change anything to the networking configuration inside VM Essentials.

Inside VM Essentials you will see all this appear. I like this picture a lot since it shows all components used in the VM Essentials architecture (see above) like the use of OVS Bridges, the LibVirt networks and portgroups that are used and the TAP interfaces. Welcome in Linux and KVM networking! Get used to it.

Below is a chart showing the connection from the bonded network interfaces created in the example network configuration above to OVS bridges (a virtual switch that connects VMs to the physical network), LibVirt Networks (logical virtual networks managed by LibVirt), LibVirt portgroups (subgroups within LibVirt Networks that can expose configurations, such as VLAN tagging), and tap interfaces (the virtual connection that connects the VM to the OVS bridge).

All these components you will see later appear when using the VM Essentials Manager.



So in this picture we are using default/native VLANs on the trunk port for the management network. The tagged packets will go through the Compute Network and will be untagged by the LibVirt portgroups to have them connected to the correct VM's.

And this is how it will look like during the Ubuntu deployment:

```
Network configuration
Configure at least one interface this server can use to talk to other machines
access for updates.

NAME      TYPE  NOTES
[ bond0   bond  -           ▶ ]
static    192.168.131.244/23
bond master for ens33, ens34

[ ens33   eth   enslaved to bond0 ▶ ]
00:50:56:93:2b:d6 / VMware / VMXNET3 Ethernet Controller

[ ens34   eth   enslaved to bond0 ▶ ]
00:50:56:93:f7:06 / VMware / VMXNET3 Ethernet Controller

[ ens35   eth   -           ▶ ]
DHCPv4    -
00:50:56:93:f7:ec / VMware / VMXNET3 Ethernet Controller

[ ens36   eth   -           ▶ ]
DHCPv4    192.168.130.54/23
00:50:56:93:7b:1a / VMware / VMXNET3 Ethernet Controller

[ Create bond ▶ ]
```

To keep in mind: the more you go to the right in this picture the more you can control in the VME Manager GUI. The more you go to the left, the more challenging it will be to get configured in the GUI with the current version. Cannot talk about the future, but this is a nice to have that I told already to VM Essentials product management that VMware did a great job in all these years to configure networking end to end.

Don't understand me wrong, you can do all this as well in the Linux – KVM world, but there is today quite some CLI involved...

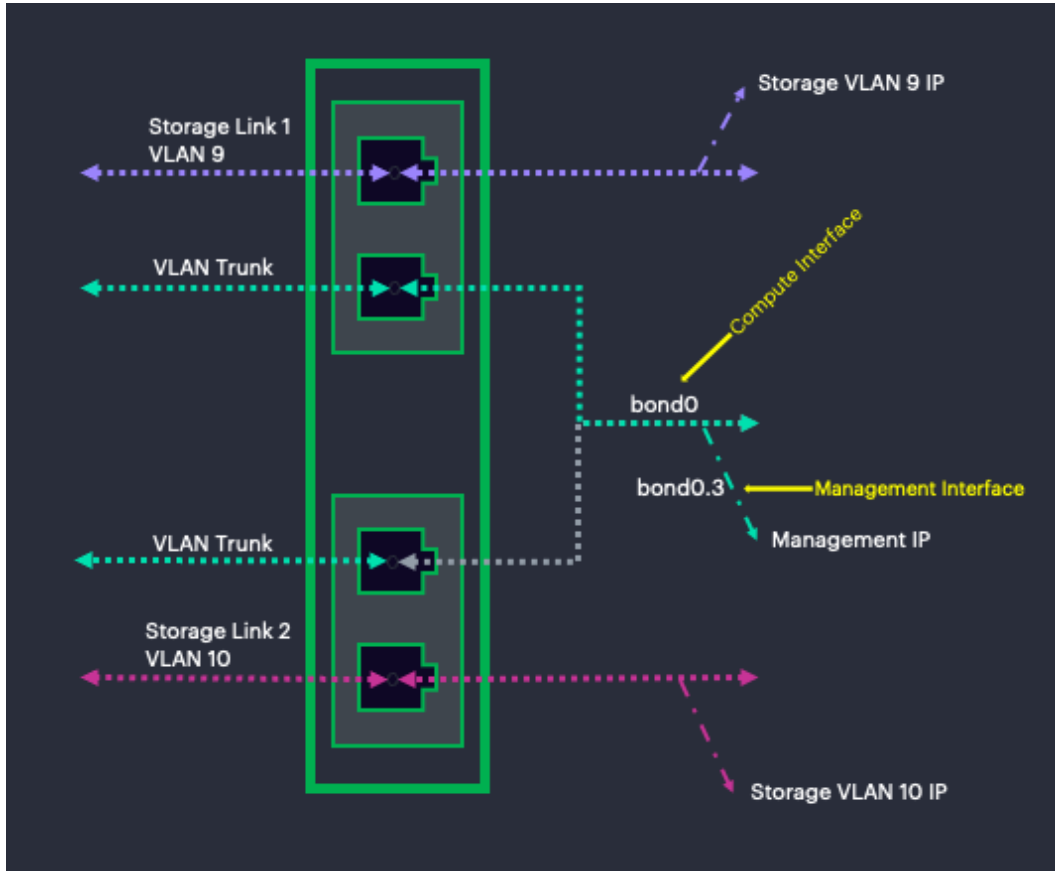
That is why I say that all the networking that you can configure during the deployment of the server will not bother you anymore afterwards... So my advice here is to create that bond, with or without VLAN's, or multiple bonds during the installation, and not afterwards... You can, but I hope you have some good Linux knowledge. Or the VM Essentials developers listen to our feedback and enhance this experience... Stay tuned!

So we arrive at scenario 2!

4 NICs – converged Management & VM traffic – iSCSI storage – no Native VLAN

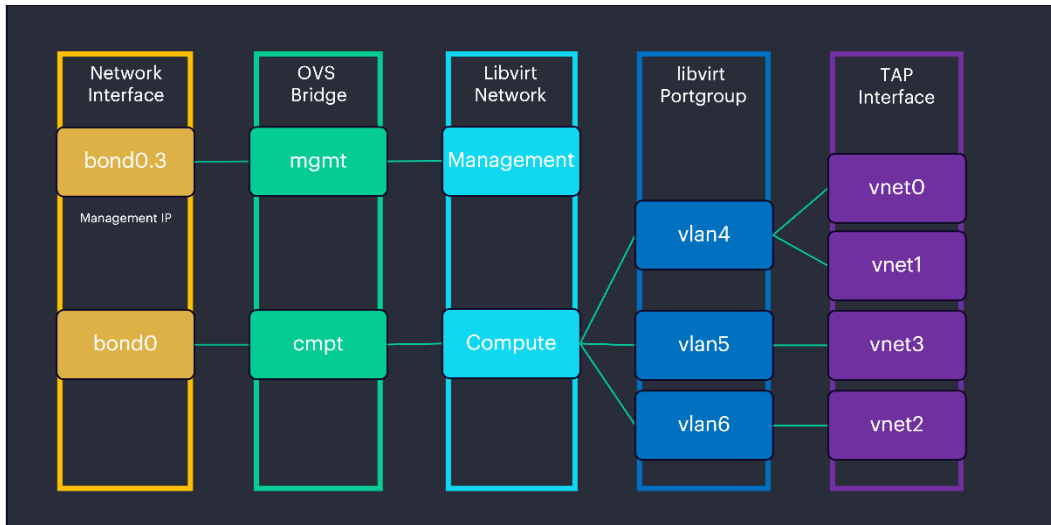
When not using native VLANs you will need to create VLAN tags on the interfaces because the switch is no longer handling the tagging for you. Your IP addresses then go on those

VLAN tags. In this example bond0 becomes the compute interface and bond0.x becomes the management interface when configuring the cluster in a later step.



Logically, and easier to compare this would give you this configuration

And yes this needs to be done ideally during the installation of Ubuntu, so this is why the Networking section is covered before the actual installation bits and pieces...



So during the bond creation you have the option to add a VLAN tag to this bond

```

Network configuration
Configure at least one interface this server can use to talk
access for updates.

NAME      TYPE  NOTES
[ bond0   bond  -
static    192.168.131.244/23
bond master for ens33, ens34

[ ens33   eth   enslaved to bond0
00:50:56:93:2b:d6 / VMware / VMXNET3

[ ens34   eth   enslaved to bond0
00:50:56:93:f7:06 / VMware / VMXNET3 Ethernet Controller

◀ (close)
Info
Edit IPv4
Edit IPv6
Edit bond
Add a VLAN tag
Delete
  
```

By doing this you will see a second bond appear in the list called bond 0.x, where x is the VLAN ID that you assigned. Bond 0.3 in my case. Use that bond 0.3 when deploying your VME Manager afterwards (see below).

```
Network configuration

Configure at least one interface this server can use to talk to other machines
access for updates.

NAME      TYPE  NOTES
[ bond0   bond  -           ▶ ]
bond master for ens33, ens34

[ bond0.3 vlan  -           ▶ ]
static    192.168.131.244/23
VLAN 3 on interface bond0

[ ens33   eth  enslaved to bond0 ▶ ]
00:50:56:93:2b:d6 / VMware / VMXNET3 Ethernet Controller

[ ens34   eth  enslaved to bond0 ▶ ]
00:50:56:93:f7:06 / VMware / VMXNET3 Ethernet Controller

[ ens35   eth  -           ▶ ]
DHCPv4    192.168.130.54/23
00:50:56:93:f7:ec / VMware / VMXNET3 Ethernet Controller

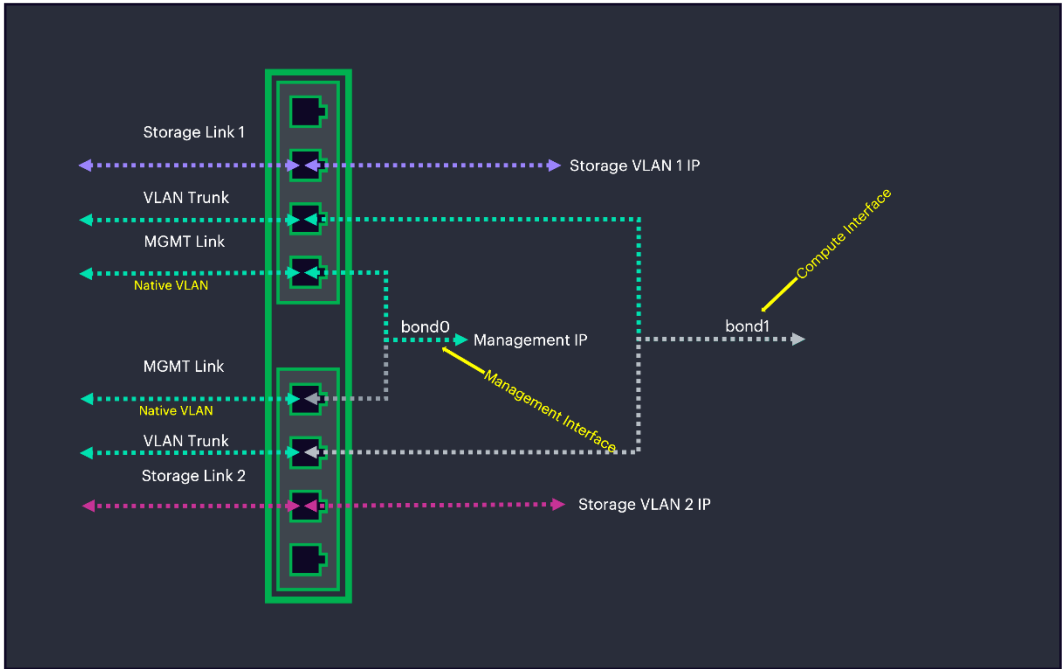
[ ens36   eth  -           ▶ ]
DHCPv4    192.168.130.53/23
00:50:56:93:7b:1a / VMware / VMXNET3 Ethernet Controller

[ Create bond ▶ ]
```

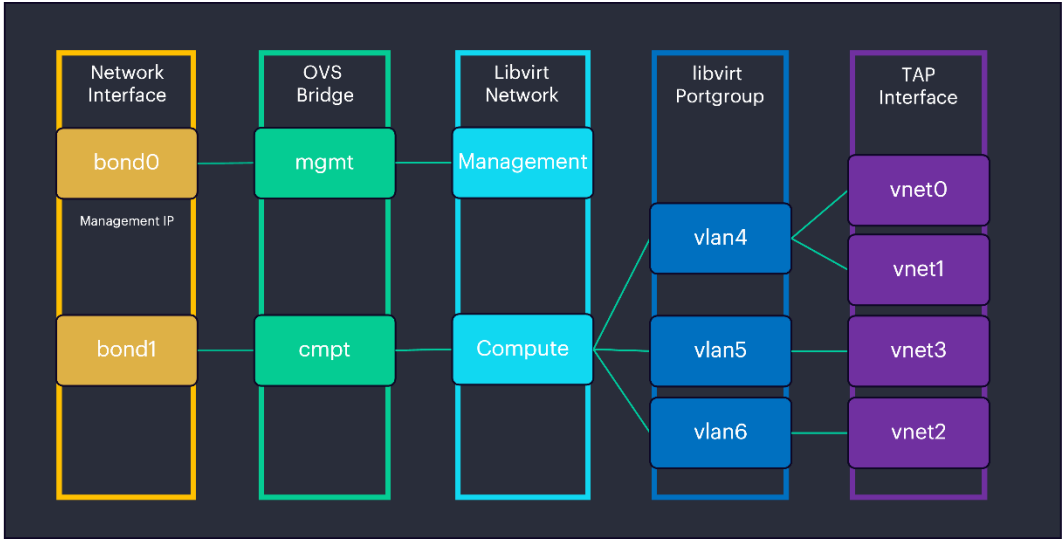
Which brings us to the next scenario:

6 NICs – separate Management, VM & iSCSI storage traffic – no Native VLAN

When hosts have six or more network interfaces spread across at least two cards, management and VM traffic can be separated for even greater network throughput. The overall configuration is very similar. The storage lanes are still separated as they were previously. In this case, create an additional bond across two of the interfaces for VM traffic. As previously, a bond is also created for management traffic. In this example, these bonds are created in an active/backup configuration.



Logically it will look like this:



And during the Ubuntu installation and configuration it will look like this:

```
Network configuration
Configure at least one interface this server can use to talk to other machines
access for updates.

NAME      TYPE  NOTES
[ bond0   bond  -           ▶ ]
disabled
bond master for ens33, ens34

[ bond1   bond  -           ▶ ]
disabled
bond master for ens35, ens36

[ ens33   eth   enslaved to bond0 ▶ ]
00:50:56:93:2b:d6 / VMware / VMXNET3 Ethernet Controller

[ ens34   eth   enslaved to bond0 ▶ ]
00:50:56:93:f7:06 / VMware / VMXNET3 Ethernet Controller

[ ens35   eth   enslaved to bond1 ▶ ]
00:50:56:93:f7:ec / VMware / VMXNET3 Ethernet Controller

[ ens36   eth   enslaved to bond1 ▶ ]
00:50:56:93:7b:1a / VMware / VMXNET3 Ethernet Controller

[ ens38   eth   -           ▶ ]
DHCPv4      192.168.130.63/23
00:50:56:93:ba:81 / VMware / VMXNET3 Ethernet Controller

[ ens39   eth   -           ▶ ]
DHCPv4      192.168.130.62/23
00:50:56:93:2c:dd / VMware / VMXNET3 Ethernet Controller

[ Create bond ▶ ]
```

Of course a lot of varieties exist, depending on the hardware being used and the physical network configuration. When using FC or SAS storage arrays, you can remove the 2 iSCSI connections of course...

My final advice: the more you can configure here during the initial Ubuntu installation, the less you will need to do afterwards in the hpe-vm tool or in Netplan... And trial and error is your friend with your first deployment. And also the great tool **PING**...

As soon as you configure the ethernet interface here with an IP address, or it shows up by using DHCP, just ping and check for a reply... No reply? Check network cabling and switch configuration... And make the necessary changes until you get a reply.

How to get started? Where to download?

I started myself with version 8.0.1 somewhere at the end of 2024 when the first version was pre-released.

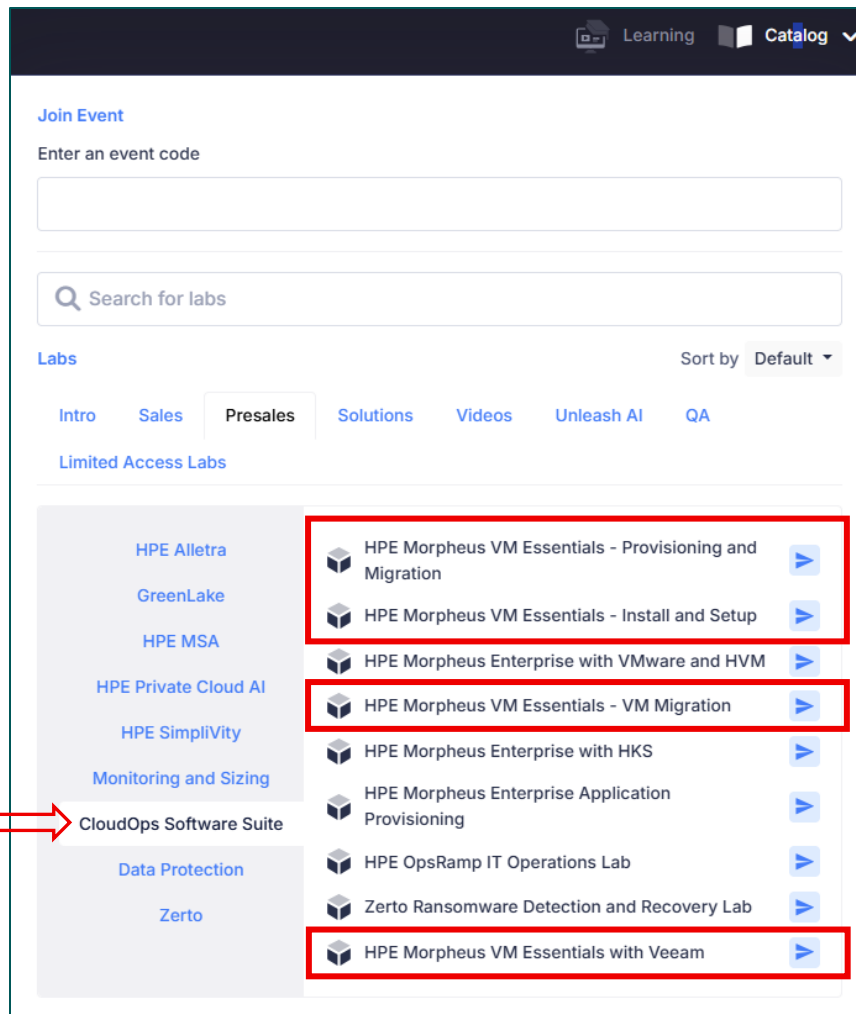
Since then, I tried and tried again, restarted, troubleshooted and tried again... Good for building experience in Ubuntu (networking) and KVM hypervisor troubleshooting, however this is not ideal when facing a lot of IT challenges that need to be taken care of outside that great VM reset.

This is why I recommend to get started by learning how to use the product in a **pre-staged environment provided by HPE**, and where you can try yourself setting up and configuring HPE VM Essentials, then deploy and maintain VM's on it, and try the migration from a VMware VM towards HPE VM Essentials.

All this will take just a few hours end to end, but will give you the full confidence of starting to deploy the software yourself on your own home lab or test environment.

<https://salesenablement.ext.hpe.com/>

To do this, navigate to the HPE Sales Enablement portal and launch one of the 4 available VME labs under the Presales – HPE VM Essentials section



Once that you have the HPE VM Essentials Manager functions and features in your fingers you are ready to deploy the HPE VM Essentials software on your own environment.

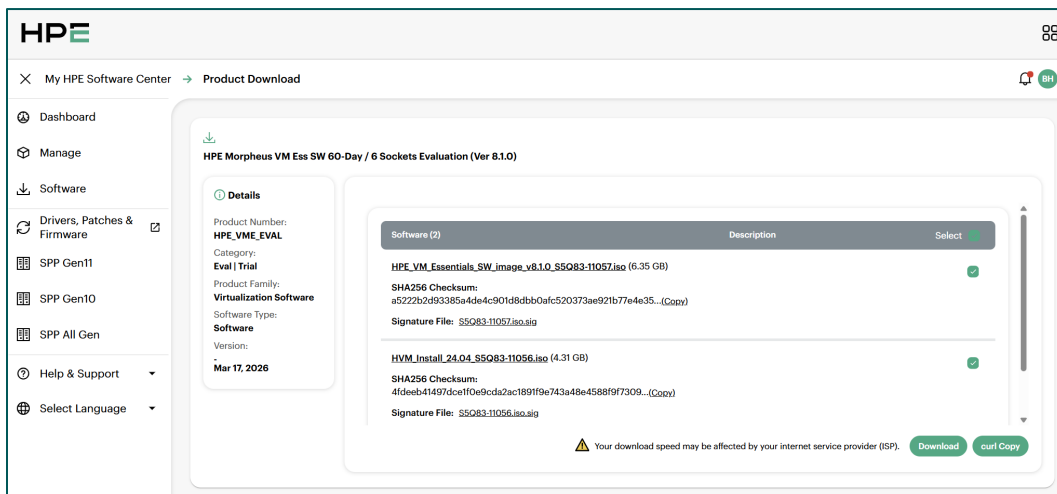
Take a look in the Qualification Matrix which hardware is supported by HPE (if you need and want support), however any hardware will do the thing as long as Ubuntu 24.04 LTS is supported and runs on your hardware.

HPE delivers a **60 day trial license** for everyone to evaluate and test the software. The 60 day trial license starts automatically once you deployed the HPE VM Essentials Manager.

The software can be downloaded straight from the HPE Software Center at

https://myenterpriselicense.hpe.com/cwp-ui/product-download-info/HPE_VME_EVAL/-/sw360_eval_customer

where you just look for VME and find the 6 Socket Evaluation Version software.



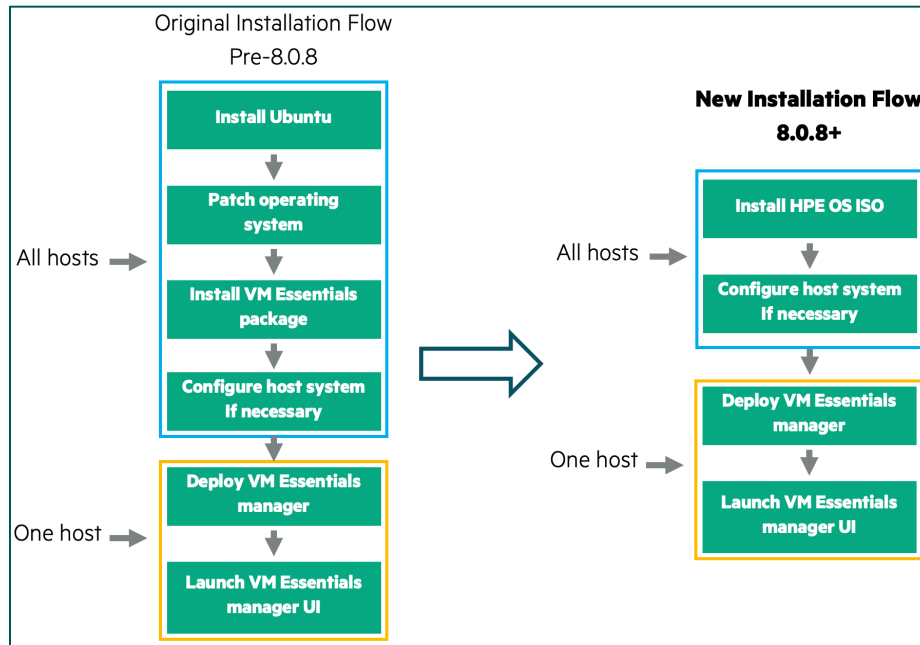
Download both files since **you will need both of them!** Find out why below!

First some history on the installation procedure.

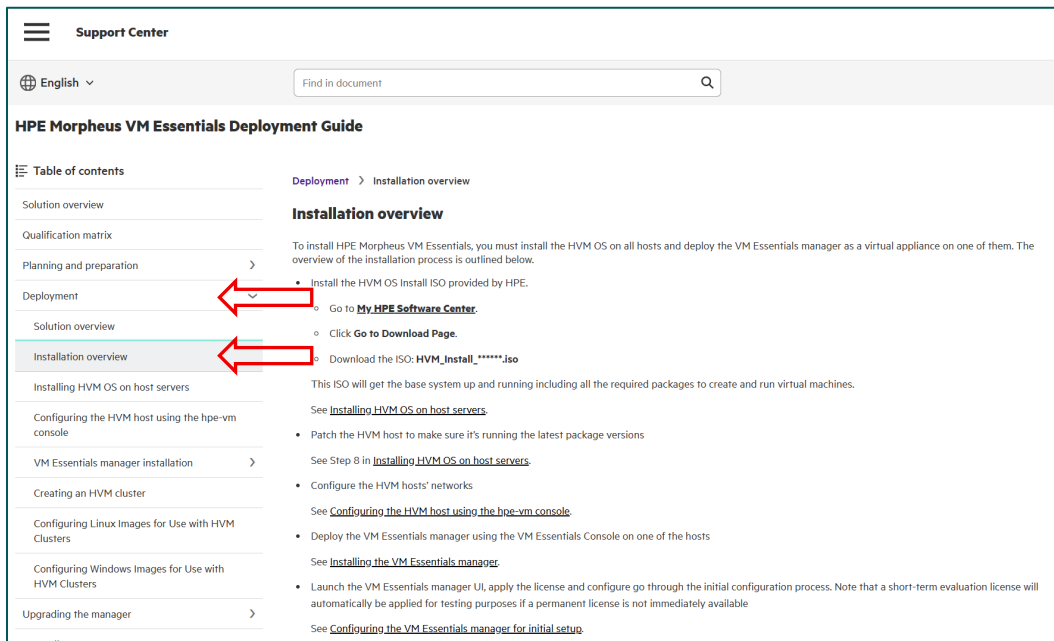
Before version 8.0.8 you needed to install yourself Ubuntu 24.04 on each host, patch the operating system and install the HPE VM Essentials hypervisor software individually, and have it configured.

From version 8.0.8 on, HPE simplified the installation experience by providing **1 unified installer ISO image** (the bottom one from the list above) that will combine the 4 steps into 1 single installation experience.

You still need to do some configuration like the networking and local storage configuration, but it was enhanced for sure, compared to the original way of working.



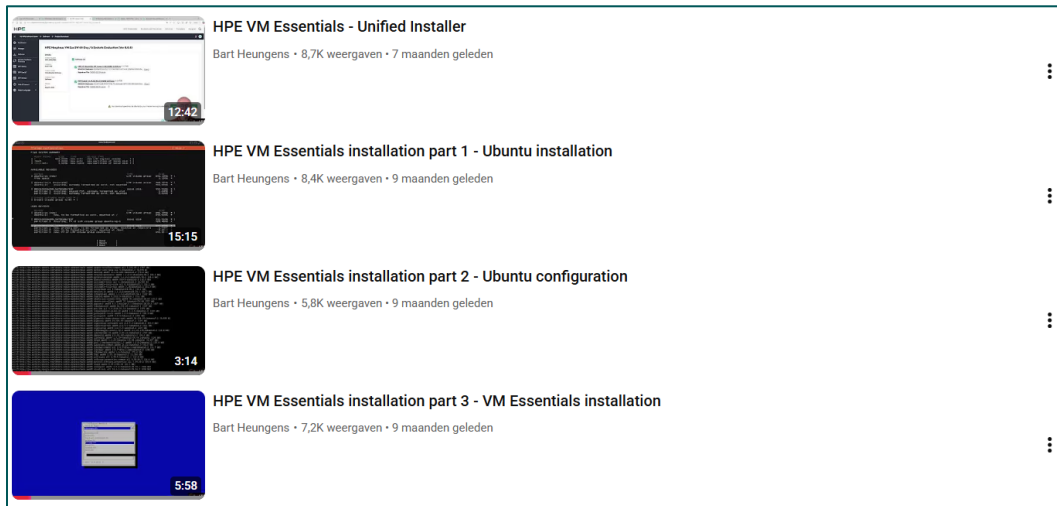
The installation procedure is nicely documented on the HPE Support pages at www.hpe.com/support/vme-docs, go to the Deployment Guide portal and navigate to the Deployment – Installation section.



For those who prefer watching video's instead of reading text, I have created some videos on the Unified Installer experience, but also the separate installation of Ubuntu, the patching of Ubuntu, and installing the HVM hypervisor bits for the die-hards.


Just go to my YouTube playlist via www.hpepedia.com or directly

<https://www.youtube.com/playlist?list=PLoEeErOXEC9VuF3n-C3myREYaU6gL7B1T>



Once the installation is done on each host, you are ready for the next step, which is the deployment of the HPE VM Essentials Manager VM.

The installation of the HPE VM Essentials Manager virtual machine is initiated on one of the HVM hosts from the *hpe-vm* application within the SSH console.




It consists of entering the required information like name, networking and location of the qcow2 source file, which you will find in the second ISO file that you downloaded earlier. Extract the file and upload it to the /tmp (or any other of your choice) folder on the HVM host using WinSCP or a similar tool to copy files towards a Linux host.

The complete installation procedure can be found on the HPE Support website so I will not repeat it here again.

For the people who prefer videos, you can find the **Part4 video in my YouTube playlist** which shows the entire procedure

The final part of the installation is the initial configuration of the VM Essentials Manager, where you need to create a Group, add a VME Cloud, which will hold the VME Cluster of 1 or more HVM hosts...

All these steps are nicely documented as well on the HPE support webpages, and otherwise look at my **video Part5**:



Once the cluster is created and configured, you are ready to deploy your first VM!

Where can I find features like vMotion? Storage-vMotion? Affinity? HA & DRS? Maintenance Mode?

In almost every demo/workshop I hosted, this must be the most popular question from the typical VMware administrator, just to be able to compare features and functions...

And I must admit the first I started deploying and evaluating HPE VM Essentials I had to look myself around a little bit, since the functions are there of course but are called differently.

Therefore, an overview when comparing with your trusted hypervisor.

Instance versus Virtual Machine

Before diving in the comparison details there is an important thing to understand in the HPE Morpheus control plane.

Remember that the underlying Morpheus software can deploy and manage bare metal, virtual machine and container objects in multiple private and public clouds. Therefore, the Morpheus team decided many years ago already to call these instances. And knowing that VM Essentials shares the same code as Morpheus Enterprise, also here you need to get used about these so-called Instances.

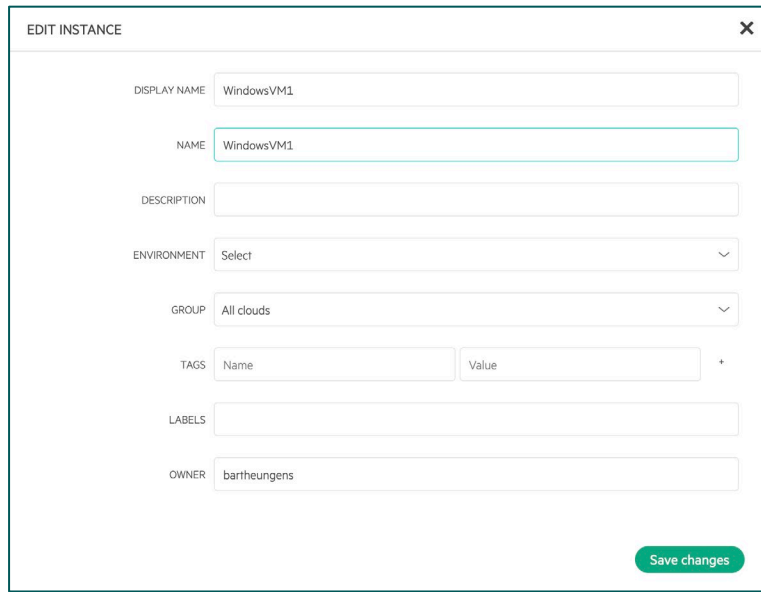
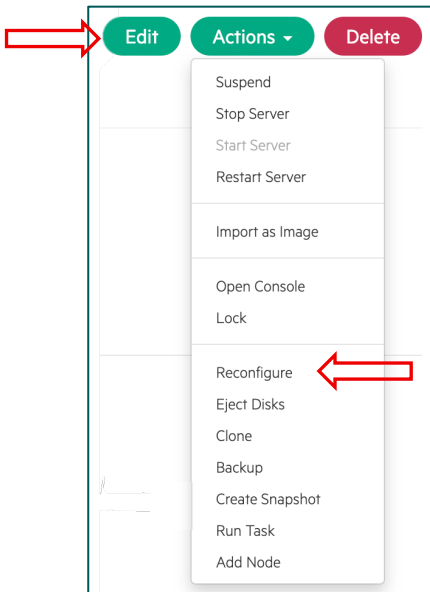
So, if you want to deploy a VM, then go to Instances in the top menu, and don't look for Virtual Machines...

It is important to understand this way of working by HPE Morpheus, since in most cases **1 Instance will be equal to 1 Virtual Machine.**

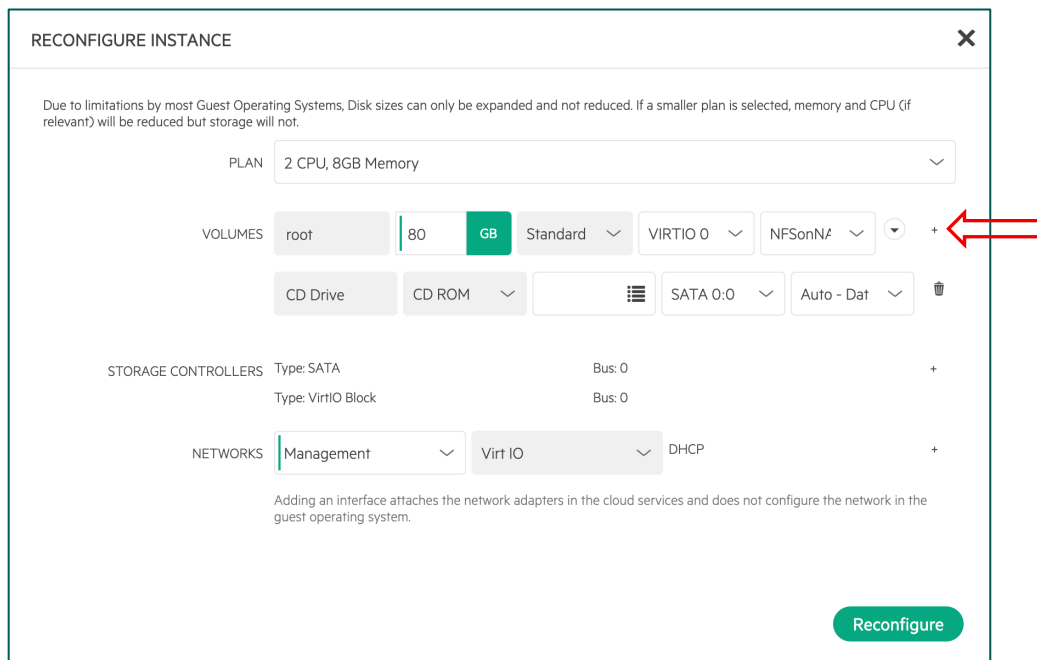
However, know that there are scenarios where you might have multiple Virtual Machines running underneath one Instance. One example might be multiple web servers hosting a web-based application running as a cluster of webhosts.

Edit a Virtual Machine

Second thing I needed to get used to when starting using VM Essentials is making changes to the properties of a Virtual Machine. In VMware just hit the Edit button, not in VM Essentials!



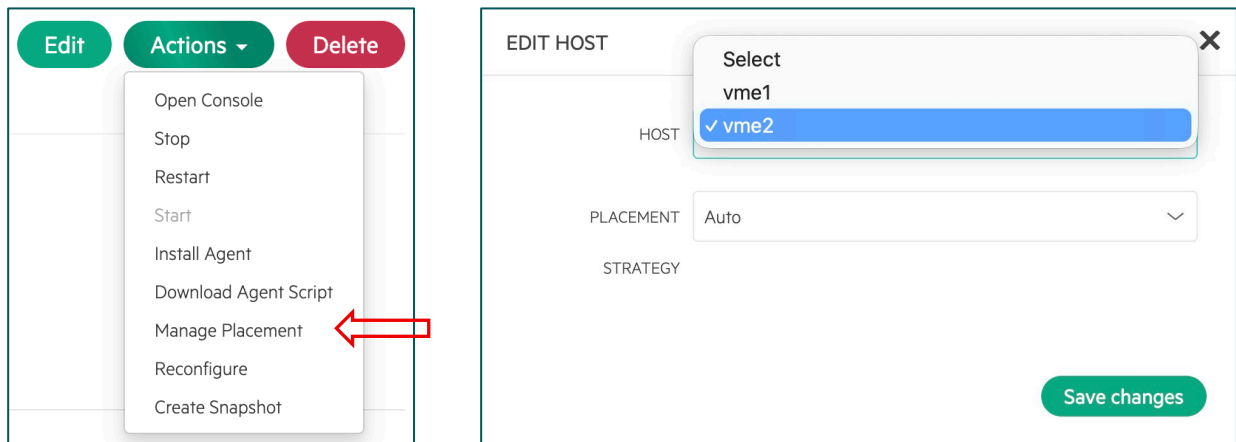
If you want to change the virtual resources of a Virtual Machine, then select **Reconfigure** instead of Edit.



It is here where you can change the typical virtual resources like CPU and memory through so called plans, existing disks and networks. Here you can add also additional storage and network devices through the (tiny) **+** icon!

vMotion

Moving a Virtual Machine online from 1 host to another host in a cluster is perfectly possible, however HPE does this through a mechanism called “Manage Placement” instead of vMotion in the VMware world.



Manage Placement is done on VM level and not on Instance level! So don't be confused when you are in the Instance pane and cannot find Manage Placement in the Actions menu. Since you don't move an Instance, you move a Virtual Machine inside an Instance.

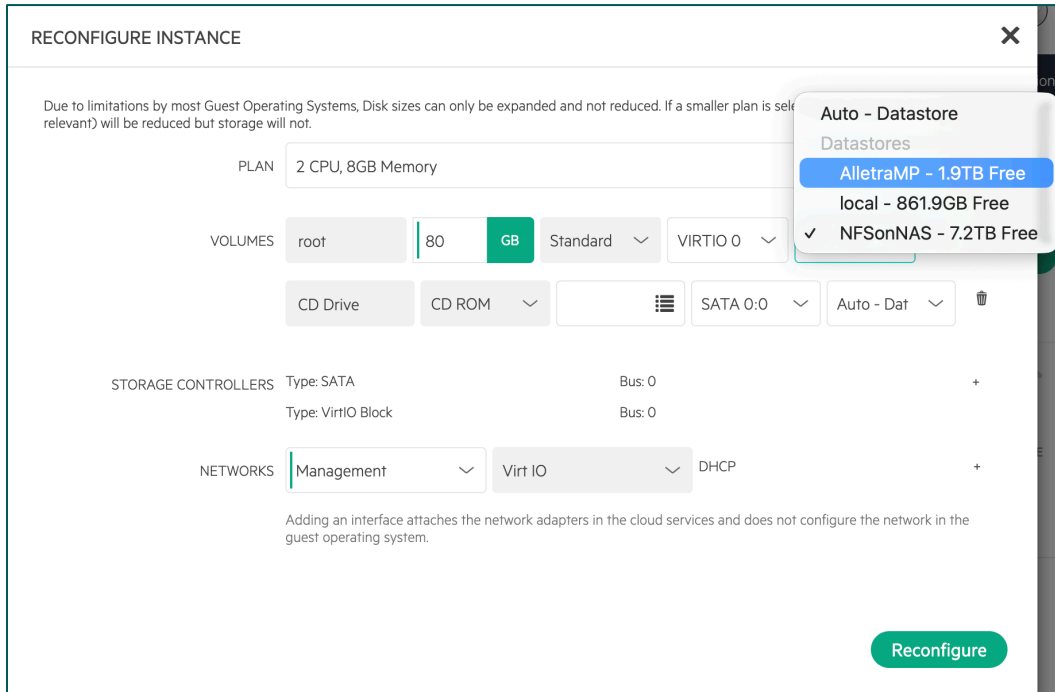
So, if you are in the Instance view, just click the Virtual Machine inside the Instance at the bottom of the Instance pane, and select the Actions button from the Virtual Machine view, which is actually called the Compute View under the Infrastructure menu.

First time it might look a little bit confusing or weird, but you will get used to it quite quickly.

Storage-vMotion

It is possible to change the location of the QCOW2 file from an existing Virtual Machine to another datastore online by going to the **Reconfigure** pane from that Virtual Machine, and

then select the dropdown box from the to-be-moved virtual disk and select another datastore from the list.



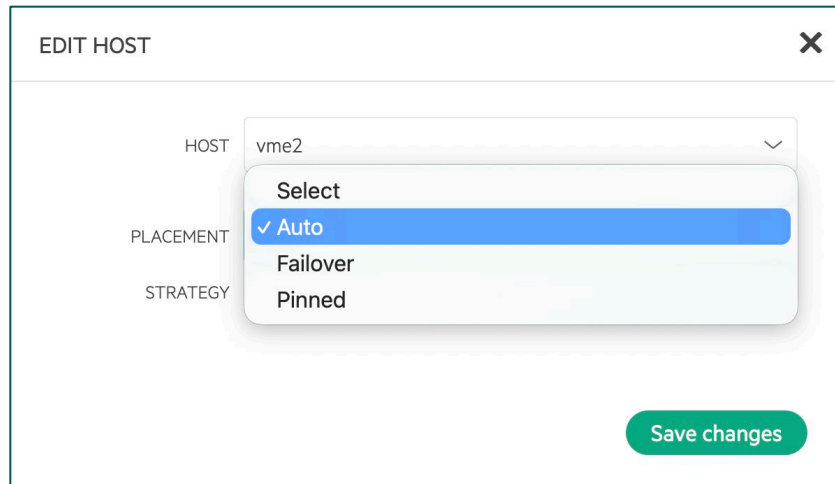
By hitting the Reconfigure button will start the Storage-vMotion-like Move process of the QCOW2 file.

HA & DRS

Another highly solicited VMware feature is **HA aka High Availability**.

And yes, it is also there. However, there is a difference compared to VMware, since the HA feature is not set at Cluster level like in VMware, but can be configured on Virtual Machine level.

It is the same for that other highly demanded feature which needed the Enterprise Plus license called DRS aka Distributed Resource Scheduler. Also, this can be configured at Virtual Machine level.



But don't look for HA or DRS, this feature is controlled through the "Manage Placement" option of a VM (look a few pages above).

The first option "**Auto**" in the dropdown box will place the Virtual Machine automatically on the host in the cluster based on load of the cluster and its hosts, which is like DRS style feature in VMware, but then on VM level instead of cluster level.

The second option "**Failover**" is then more the HA style of action inside VMware, where the Virtual Machine will be moved only when failover is necessary, for instance when the host where the VM is running is not available anymore (planned or unplanned).

There is a third option in the Manage Placement list called "**Pinned**". As the name implies, the Virtual Machine is pinned to that host and so will not move to another host when that actual host is not available. So, the Virtual Machine will stay in a down state unless you start it yourself on another host manually.

A practical use case for a pinned VM could be a hardware device in a host server that needs to be connected via USB or PCIe Passthrough with a particular VM (USB dongle with application license, or an RDX tape drive for a backup server?).

Affinity rules

Next VMware feature that is asked frequently is the availability of **Affinity Rules**.

And yes, it is there since version 8.0.9.

NEW AFFINITY GROUP ✕

NAME

TYPE
 Keep Separate
 Keep Together

ACTIVE

SERVICES

▶ Group Access

Save changes

An affinity group contains a type (either Keep Together or Keep Separate) and a list of servers which should have the rule applied. Whenever possible, servers configured to "Keep Together" will run on the same HVM Host. Servers configured to "Keep Separate" will be balanced across HVM Hosts to the maximum extent possible.

The screenshot shows the HPE Morpheus VM Essentials interface. The top navigation bar includes 'Operations', 'Provisioning', 'Library', 'Infrastructure', 'Backups', 'Tools', and 'Administration'. The 'Infrastructure' tab is active, showing 'Clusters > VME cluster'. The main content area displays the 'VME cluster' details, including a status bar with 'OK', 'Type: HVM', 'Last Sync: 01/05/2026 11:36 PM', and 'Sync Duration: 5 seconds'. Below this are several metrics: 2 HOSTS, 0 ALARMS, 4% MAX CPU, 4% MEMORY, and 69% STORAGE CAPACITY. A navigation bar at the bottom of the main content area includes 'Summary', 'Hosts', 'VMs', 'Network', 'Storage', 'Virtual Images', 'Monitoring', 'Resources', 'History', 'Wiki', and 'Addon Package'. The 'Resources' tab is selected, and a red arrow points to the 'History' link. Below the navigation bar is an 'Affinity Groups' section with a search bar and a '+ Add' button. A table lists the affinity groups:

NAME	TYPE	RESOURCE POOL	ACTIONS
Divide Domain Controllers	Keep Separate	VME cluster	ACTIONS

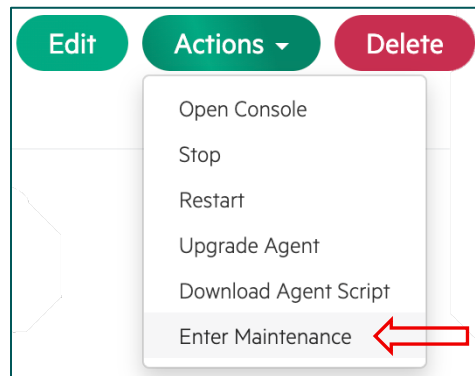
Affinity groups are listed on the **Resources tab** of the HVM Cluster detail page under Infrastructure > Clusters. From the Actions menu you can create an affinity group, or existing ones can be edited or deleted. By editing an affinity group, users may view or edit its enabled status (affinity groups which are not enabled will not be acted on).

Maintenance Mode

A common question from the typical VMware administrator is if it is also possible to put a host in Maintenance mode... Answer is of course YES.

This is for instance needed when upgrading your VM Essentials cluster to a newer version. Look further in the guide for the actual process, but Maintenance Mode is part of it.

To do this, go to one of the hosts in your cluster (notice that you are in the Compute view now) and click the Actions button. There is an option called “**Enter Maintenance**”.



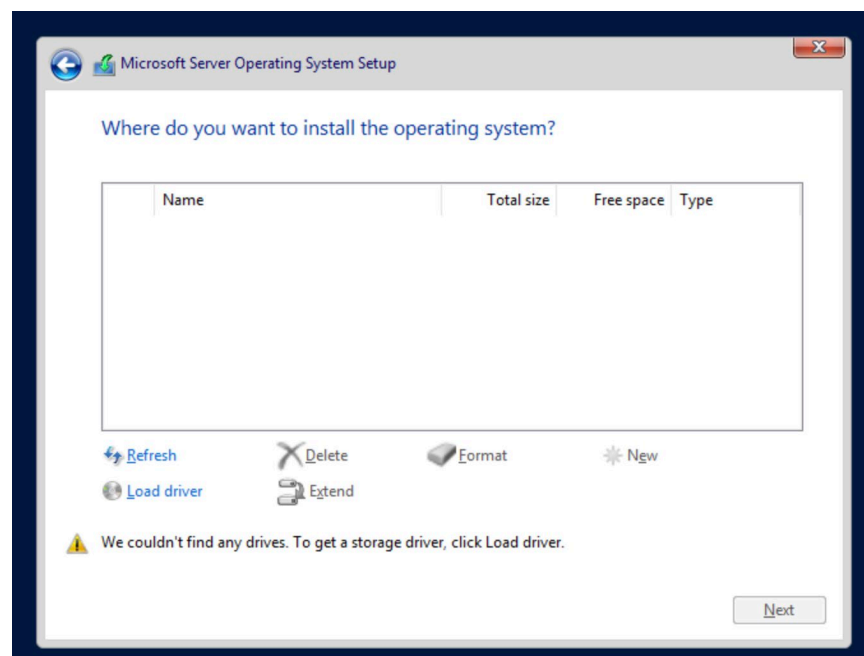
When entering maintenance mode, the host will be removed from the pool. Live VMs that can be migrated will be moved to new hosts.

When maintenance has finished, go back to the ACTIONS menu and select “Leave Maintenance.”

How to successfully deploy a Windows VM on a KVM-based hypervisor

Maybe the biggest ‘challenge’ (don’t want to call it a common issue, although when looking through the posts on Nutanix, Proxmox, Reddit and other KVM forums I think it is the number 1 mostly asked question/issue) with KVM as alternative hypervisor, is the deployment of a Microsoft Windows based Virtual Machine.

And I must admit, when I started playing around with the early versions of VM Essentials, I failed also a few times deploying that first Windows based VM, being blocked at the storage pane during the installation process, stating that no storage device can be found...



I should have read the manual, since it is nicely explained in there.

This is also why I created this YouTube video for all of you not reading this guide or not reading manuals. ;-)



Link to my YouTube playlist somewhere above in this guide and at the end in the **Useful Links section**...

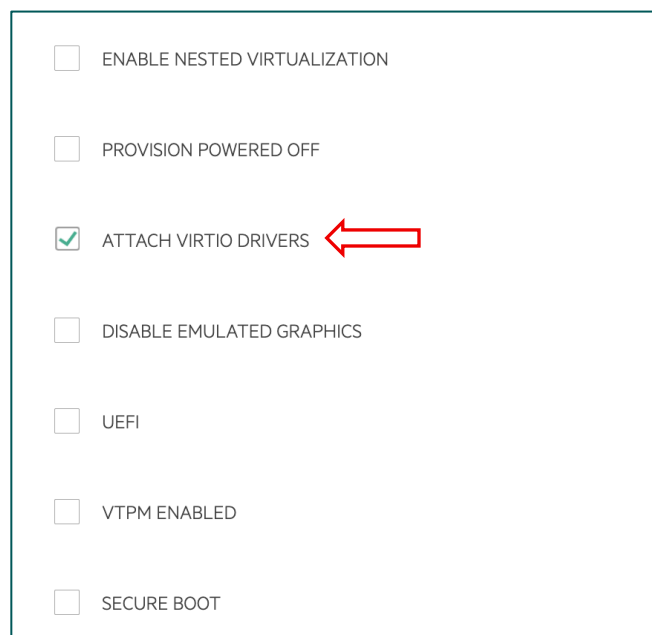
The magic trick here is actually **not that difficult, when you know...**

The issue is that KVM needs VirtIO drivers to be used inside the Virtual Machine to be able to work with the QEMU virtualization stack used by KVM and VM Essentials.

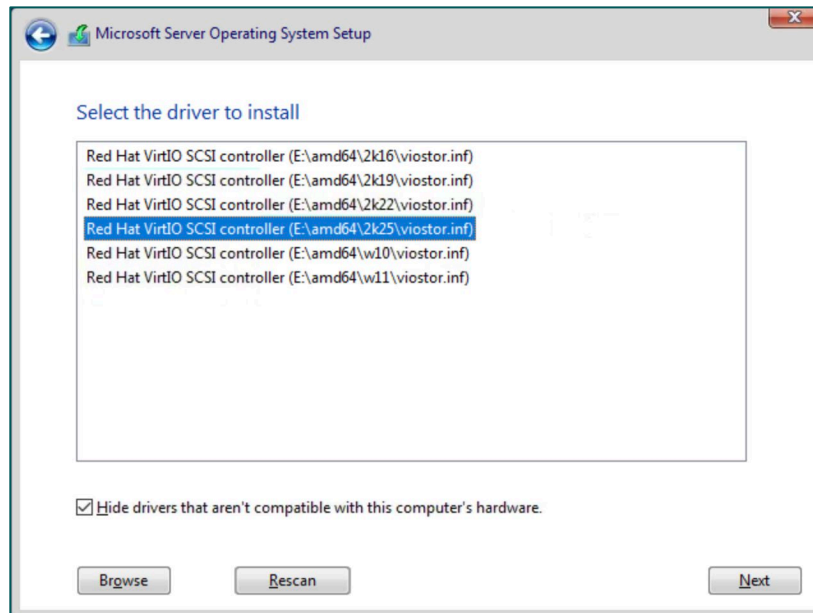
And Microsoft never added these VirtIO drivers in the sources folder from its popular OS ISO images (unlike the paravirtualization drivers from VMware). So you need to inject them yourself in the Windows ISO image, or you need to load them yourself when running the Windows installation process.

And this is what VM Essentials made quite easy for you, once you know of course.


When deploying a new Instance (look above if you don't know what is an instance compared to a Virtual Machine inside VM Essentials), you will need to go to the Advanced section in the wizard and **select the checkbox "Attach VirtIO drivers". That's it!**



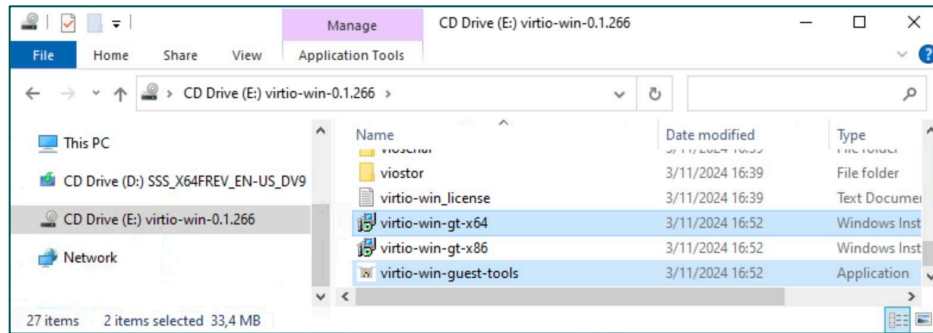
At that moment the VM Essentials deployment process will attach automatically the VirtIO ISO image to that Instance and will connect it, so that it is available during the Windows installation process and you will not get that nasty “No storage device found” error. Just select “Load Drivers”, pick the right driver, select the boot device and proceed with the installation.



Once the Windows OS is installed, you will still need to install all other drivers since you added now only the storage controller driver. Inside Device Manager you will still see a lot of unknown devices.



This is solved quite easily. Since the ISO image is still connected to the OS, once you are connected to the OS through the Console tab, you go to the E: drive and run the “**virtio-win-gt-x64.exe**” executable, that will install all drivers.

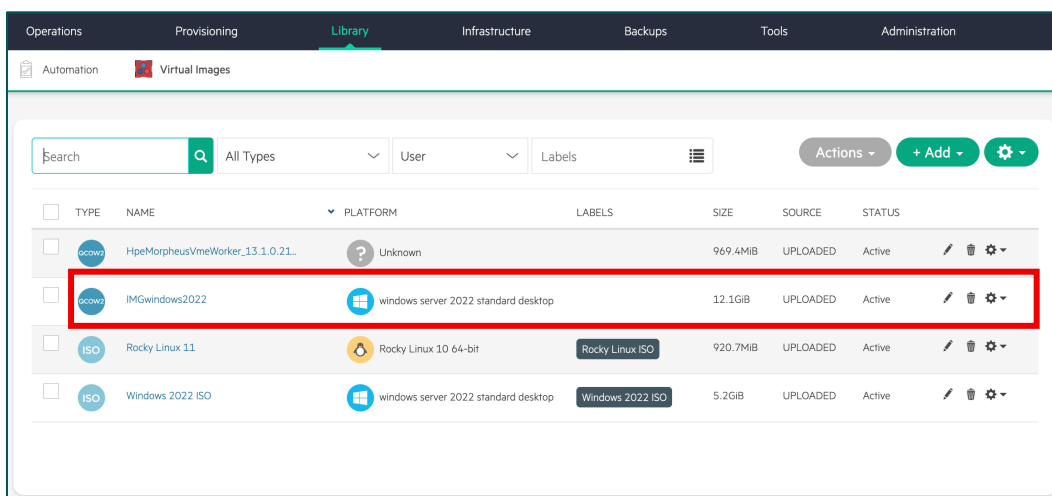


Once all drivers are installed and Device Manager is clean, the final step to do is install the “**virtio-win-guest-tools.exe**” executable, which will install the QEMU Guest Agent, very similar to the VMware Tools.

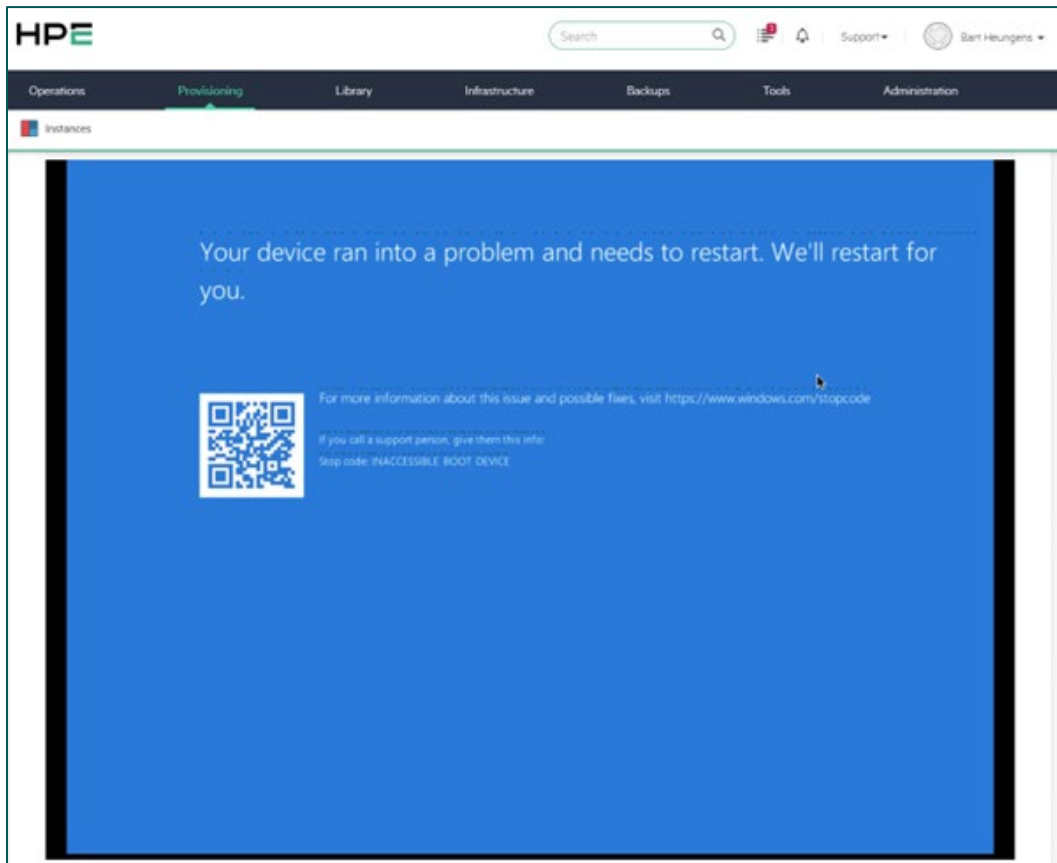
At this moment you successfully deployed a Windows Virtual Machine on VM Essentials.

Run Windows Update and SYSPREP afterwards, and you are ready to hit “**Import as image**”, which will create an image (very similar to a template inside VMware) that can be used to deploy all your future Windows VM’s.

These images are stored in the **Library – Virtual Images**.



By the way, if you convert a VMware VM towards VM Essentials to a QCOW2 file and import it without injecting the VirtIO drivers first (see below for more information on what you need to do particularly for migrations/conversions), you'll get this nice screen:



More on migrations in the next section!

Migrate VMware VM's to HPE VM Essentials

Of course you need to get those VM's running on your existing VMware environment into this new HPE VM Essentials cluster.

There are actually several ways how to accomplish this. Let me give you an overview.

Let me start first with the basics, and what we saw in the previous section on deploying a new Windows VM on a KVM based hypervisor. Drivers!

A VM in VMware is based on the VMDK format and so this will need to be converted into a QCOW2 format which is supported by the KVM hypervisor. Several tools exist to do this, Google is your friend for the techies who are interested in this matter, or just keep on reading.

If you have an existing Windows VM on VMware and move it towards VM Essentials (or any other KVM based hypervisor) by converting the VMDK file in a QCOW2 format, being a V2V process. Remember the P2V process you did many years ago to get these physical servers into VMware? You will do the same now but a little easier since they are virtual already. But the process remains the same.

The problem however is that, when you have converted the VM in a QCOW2 format, have it imported into VME and try to power on the VM, you will get a message like “No boot device found” or whatsoever... Since the VirtIO storage driver is not loaded/available/installed in the VM, since it was configured with the VMware paravirtual drivers.

So you will need to **inject** (not just install or copy whatsoever!) these drivers into the Windows VM BEFORE (so when it is still running inside VMware) you convert the VM into QCOW2 format...

That is why I tell my customers to get started NOW and not wait until the end of the VMware support contract...

All these steps are nicely explained in this document provided by HPE called the “**HPE Morpheus VM Essentials Software Migration Guide**”, which provides a step-by-step technical guide for users migrating VMs from VMware to HPE HVM. It includes Linux (Ubuntu, Red Hat Enterprise Linux) and Microsoft Windows Operating System migrations, and can be found here:

<https://www.hpe.com/psnow/doc/a50013873enw>

Maybe the most important step in the document is the injection of the 2 drivers called `viostor.inf` and `vioscsi.inf` into the Windows boot-start drivers by using the **DISM tool provided by Microsoft**, and as shown below nicely documented in the guide:

11. Enter the following command to inject the VirtIO storage drivers into the Windows [boot-start drivers](#) from the mounted ISO:

```
dism /image:C:\ /add-driver:D:\viostor\2k22\amd64\viostor.inf
```

```
dism /image:C:\ /add-driver:D:\vioscsi\2k22\amd64\vioscsi.inf
```

- a. If the OS volume letter assigned was not **C**, then use the appropriate letter.
- b. If the CD-ROM/DVD-ROM volume letter assigned was not **D**, then use the appropriate letter.
- c. In the previous example, **2k22** is used, but select the appropriate OS version as needed.
- d. If successful:

```
X:\Windows\system32>dism /image:C:\ /add-driver:D:\viostor\2k22\amd64\viostor.inf
```

```
Deployment Image Servicing and Management tool
```

```
Version: 10.0.17763.1
```

```
Image Version: 10.0.17763.771
```

```
Found 1 driver package(s) to install.
```

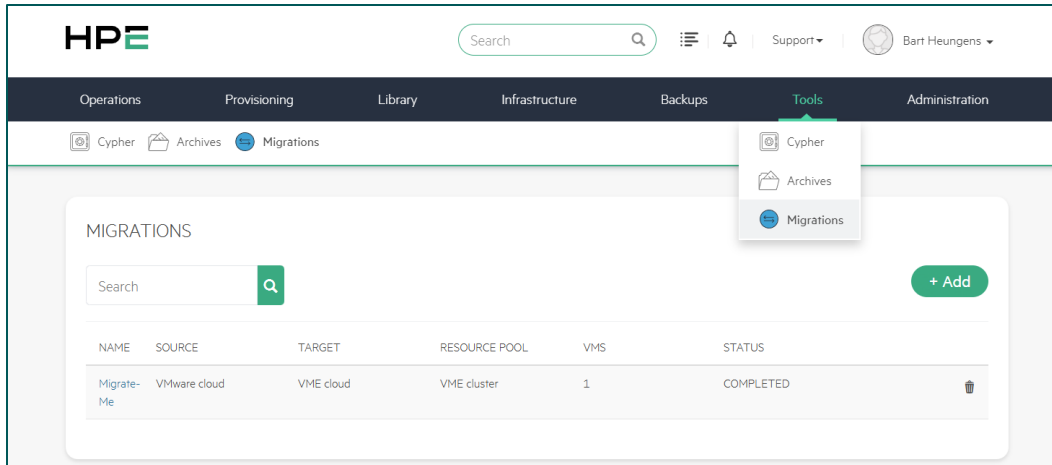
```
Installing 1 of 1 - oem21.inf: The driver package was successfully installed.
```

```
The operation completed successfully.
```

Once you injected the boot-start drivers you can start actually with the conversion of the VM into the QCOW format and import it as an instance into HPE VM Essentials.

HPE Bulk Migration Tool

To do this conversion and import, since version 8.0.8 HPE makes it easier for you with the availability of a bulk migration tool inside the HPE VM Essentials GUI to migrate existing VMs running on integrated VMware vCenter Clouds to HVM Clusters. You can find the tool in the Tools section from the main menu.



The bulk migration tool is nicely documented in the Software Documentation Guide (link elsewhere in the document) in the Tools section, it is also here that the migration requirements are listed, including the ones on injecting drivers when migrating a Windows-based VM.


The overall process as such is not that complex. You create a migration plan where you verify the current network and storage configuration of the selected VM(s), and define the target network and datastore on the VME cluster.

Once you created the plan, you can run it at the appropriate moment.

Be aware that this migration is an offline process, and so the VM will be shut down during the conversion. The process exports the VMware VM into an OVF template, which will be used then to do the actual conversion into a QCOW2 file. When something might happen during the conversion, there will be always an easy failback plan by powering on the original VM on VMware, since the original VM is not touched. Smart! Always good to have a failback plan during migrations...

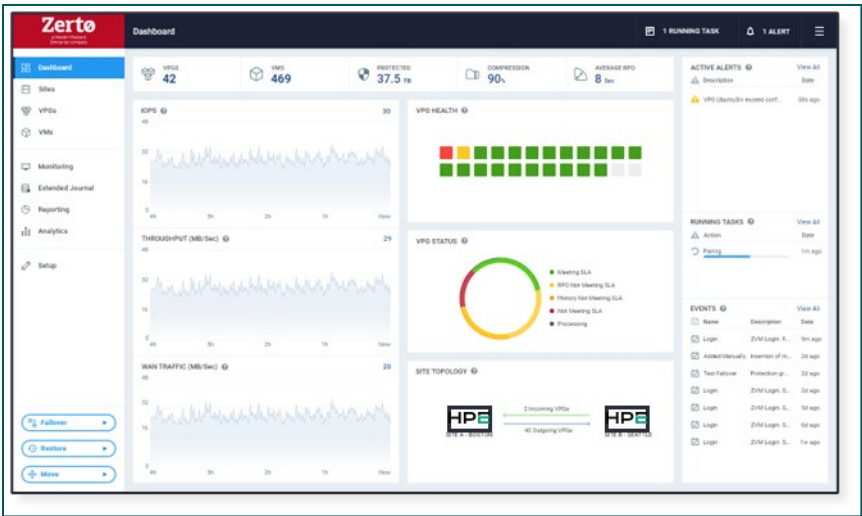
The length of time it will take for the Migration to run depends on many factors like for instance the size of the VM so the number of gigabytes that needs to be transferred. Take this into account. If downtime is a concern, keep on reading since there are alternatives!

I have also created a video of the entire process to migrate a Windows VM from my VMware environment into my HPE VM Essentials cluster, and also this video can be found of course in my VME playlist on YouTube:



HPE Zerto as a migration tool

At the moment of writing this guide, HPE released an updated version of HPE Zerto that supports VM Essentials as hypervisor.



The setup is very similar to a VMware environment. You deploy a ZVM Zerto Virtual Manager virtual machine that will deliver the web based management UI, and from there you deploy a VRA Virtual Replication Adapter on each host in the cluster.

VRAS		DATASTORES					
2 Installed		5 Available					
<input type="text" value="Search"/> View: General <input type="checkbox"/> Show only hosts with VRA installed							
Cluster / Host Address	Host Vers...	VRA Name	Status	VRA Version	VRA Addre...	# VPGs	# VMs
<input type="checkbox"/> VME cluster (2)							
<input type="checkbox"/> vme1	24.04	Z-VRA-58-v...	Installed	Latest	192.168.131...	0	0
<input checked="" type="checkbox"/> vme2	24.04		Installing... (25%)	NA	192.168.131...	0	0

<input type="checkbox"/>	Z-VRA-58-vme1 HVM	IP addr: 192.168.131.175 Version: 1.0 Virtual Machines: 1	Group: All clouds Cluster: VME cluster		2	0	0
				STATUS	MAX CPU	MEMORY	STORAGE
<input type="checkbox"/>	Z-VRA-59-vme2 HVM	IP addr: 192.168.131.176 Version: 1.0 Virtual Machines: 1	Group: All clouds Cluster: VME cluster		0	0	0
				STATUS	MAX CPU	MEMORY	STORAGE
<input type="checkbox"/>	Zerto ZVM HVM	IP addr: 192.168.130.48 Version: 1.0 Virtual Machines: 1	Group: All clouds Cluster: VME cluster		6	1	0
				STATUS	MAX CPU	MEMORY	STORAGE

At this moment HPE supports Zerto for Continuous Data Protection of VM's from VME to VME, and from VMware to VMware. The roadmap shows that later in 2026 cross-hypervisor will be supported as well, which means that Zerto technology can be used to protect VMware VM's on VME and vice versa. As soon as this will be available, at that moment you will be able to use the Zerto technology as well as a **migration tool**, where the downtime of the VM during the migration will be minutes instead of hours (or even longer) when doing the conversion/migration of the VM.

Since Zerto is using journaling technology, the to-be-migrated VM will be running while all data will be copied over to the target VM. Only when all data is in sync, you can do the actual switch by powering down the VM on the VMware cluster and starting it instantly afterwards on the VME cluster. This is how it works now on VMware, we assume it will be exactly the same on VM Essentials.

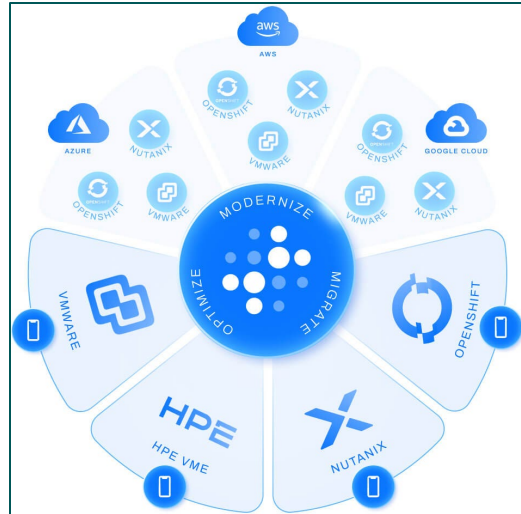
Stay tuned for further updates from this guide and HPE announcements of course.

RiverMeadow

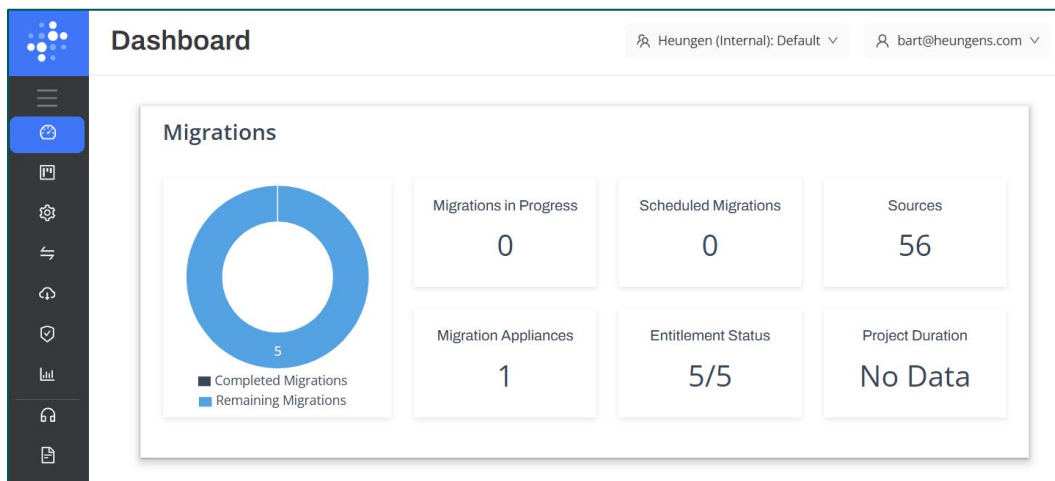
If you are looking for the Swiss army knife for migration tools, then you might want to take a look at RiverMeadow.

<https://www.rivermeadow.com/>

They provide a platform that accelerate workload migration and modernization to consolidate and simplify tooling, processes, and governance by supporting most of the cloud platforms out there (public and private) AND they support already HPE VM Essentials as a target!



Practically you will get an online dashboard on their website from where you can setup migration projects and deploy migration appliances on the target environment, being here the VM Essentials cluster.



Stay tuned for more details, I am planning to make a dedicated video and article on this migration tool.

Others (Veeam – StarWind – ...)

Like I mentioned earlier, Google is your friend when looking for VMDK to QCOW2 conversion tools.

I will not dive deeper into this matter here but tools like Veeam will be able to help you as well to migrate a VM from VMware onto VM Essentials by doing some kind of backup & restore process (don't forget the VirtIO drivers! – see above) or third-party tools like StarWind V2V Converter that will help you as well to get that QCOW2 file from your beloved virtual machine. Check out their websites for more information.

How to update HPE VM Essentials?

Another popular question/concern that I often get is regarding updates.

In 2025 every month there was a new version of the HPE VM Essentials software!

Answer: Of course, everybody wanted new features and functions...

But I don't want to update my virtualization stack every month! You know how long it takes to update my VMware stack?

Answer: The procedure is quite simple. You only have to update the HPE VM Essentials Manager by performing 3 CLI commands. After having updated the HPE VM Essentials Manager you proceed with updating each node in the HVM cluster by upgrading the Host Agent from within the VME Manager UI.

These 3 CLI commands are nicely documented on the HPE VM Essentials Software Documentation pages under Getting Started – Installation – Upgrading

The screenshot shows the 'HPE Morpheus VM Essentials Deployment Guide' documentation page. The 'Table of contents' on the left has red arrows pointing to 'Upgrading the manager' and 'Upgrade process'. The main content area is titled 'Upgrade process' and includes an 'IMPORTANT' warning about a 10-15 minute downtime. A red box highlights the following CLI commands:


```
sudo morpheus-ctl stop morpheus-ui
sudo dpkg -i xxxx.deb
sudo dpkg -i xxxx.supplemental.deb # Optional -- Only for offline upgrades
sudo morpheus-ctl reconfigure
```

Below the commands is a 'NOTE' section with the following points:

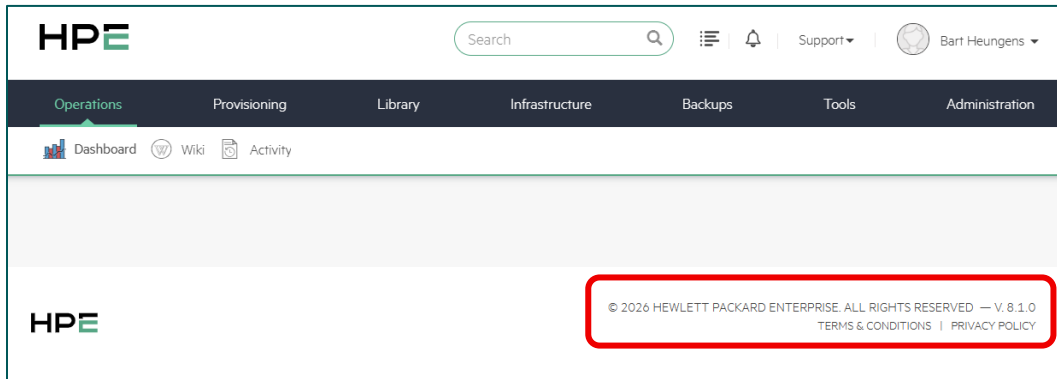
- Services will be stopped during package installation and started during the reconfigure process, including the `morpheus-ui` service.
- If the reconfigure process is interrupted or fails, the `morpheus-ui` service may need to be manually started or restarted.
- In certain situations if another service hangs on starting during reconfigure, run `systemctl restart morpheus-runsvidir` then reconfigure and restart `morpheus-ui` if successful.

At the bottom, there is a note about upgrading the HVM Host Agent.

To show the ease of the process you can find a video in my YouTube playlist showing the entire procedure.

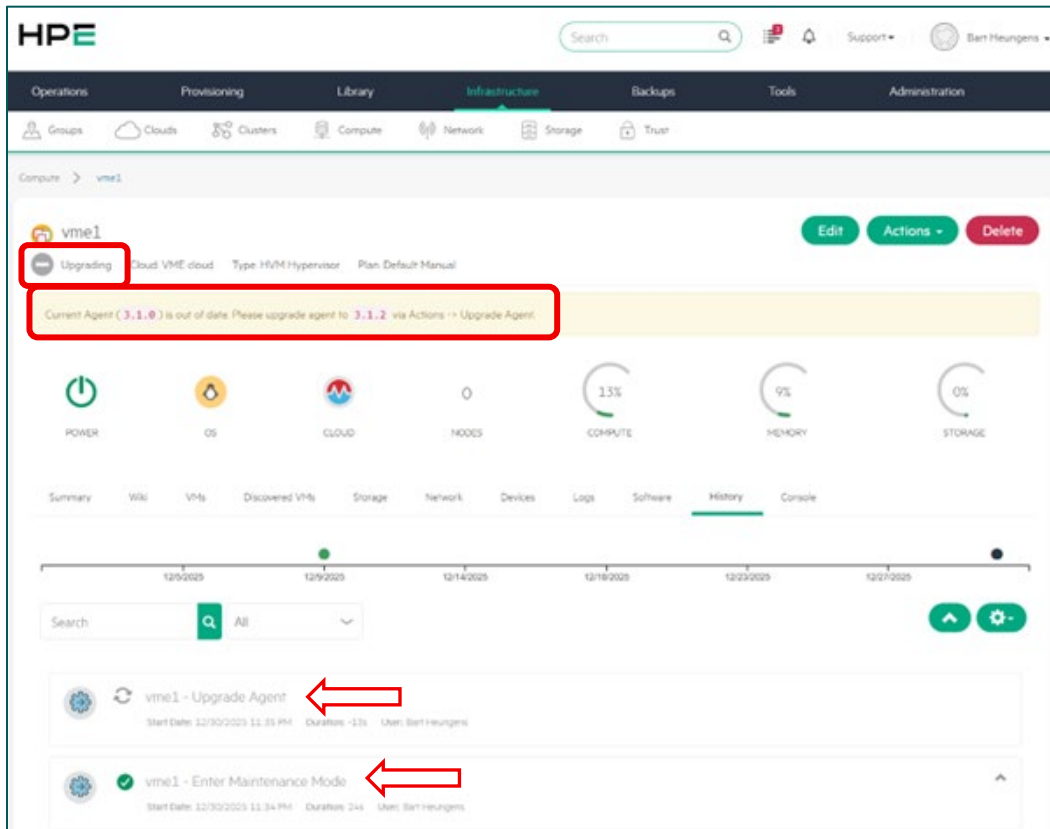


You can check the successful upgrade by checking the running version by logging in into the HPE VM Essentials Manager UI at the bottom right and verify the version displayed.

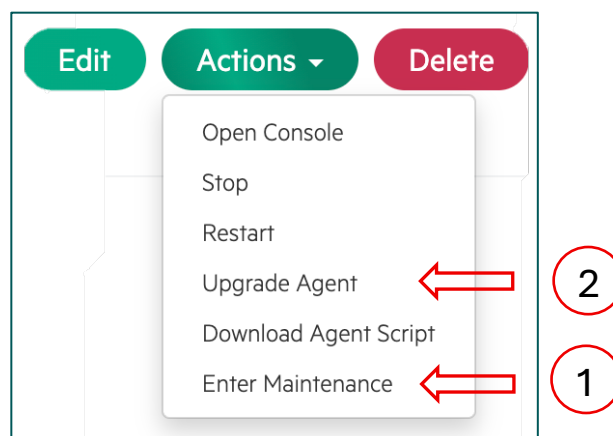


After you upgraded the HPE VM Essential Manager the next step is to **upgrade the Host Agent** software on each host of the cluster. This is only needed when a new agent version is part of the upgrade. This can be verified in the Release Notes of the update you installed.

Since version 8.0.12 it is also shown in the VM Essentials GUI in the host view.



The procedure is quite straight forward. Normally it is not required to put the server in Maintenance Mode, says the management guide, but just to be 100% safe I put the HVM host in Maintenance Mode via the Actions menu on Host level and select Update Agent from the Actions menu.



Verify if the Agent is updated before removing Maintenance Mode and proceed with the next host in the cluster.

vme1 Provisioned Cloud: VME cloud Type: HVM Hypervisor Plan: Default Manual Edit Actions Delete

POWER OS CLOUD NODES COMPUTE MEMORY STORAGE

Summary Wiki VMs Discovered VMs Storage Network Devices Logs Software History Console

Cloud: VME cloud	Cluster: VME cluster	Type: VM Host
Created: 12/09/2025 08:20 PM	Platform: ubuntu 24.04	Operating System: Linux
Hostname: vme1	iSCSI Initiator Name: iqn.2024-12.com.hpevme1:25120	Remote Host: 192.168.131.120:22
Internal IP: 192.168.131.120	External IP: 192.168.131.120	Hardware Vendor: HPE
Hardware Name: ProLiant MicroServer Gen11	CPU Model: Intel(R) Xeon(R) E E-2434	CPU Frequency: 1499.844
Cores: 8	Total Memory: 62.7GiB	Total Storage: 1.8TiB
Agent Version: 3.1.2		

Is HPE VM Essentials supported on SimpliVity?

Short answer: YES!

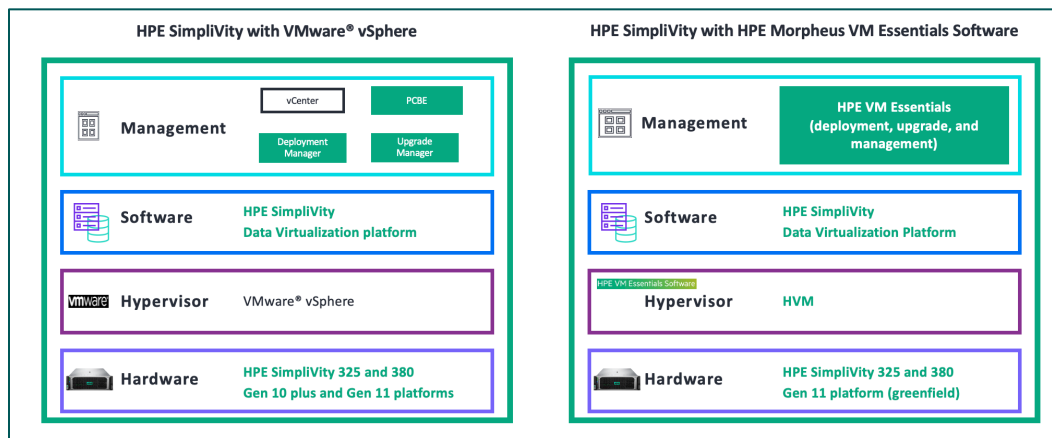
HPE released SimpliVity version 6.0.0 software in August being the first version of SimpliVity software to work with the VM Essentials as hypervisor, and started shipping systems since then. Meanwhile, when writing this guide, version 6.2.0 is already available, which has a few major updates including a more recent version of VM Essentials software 8.0.11 (was 8.0.7 with 6.0.0), which unlocks many features like 4-node cluster support, more flexible upgrade deployment and options, and a Windows version of the Arbiter software. All details can be found in the Release Notes on the Support portal (link at the end of this guide)

Always check on the Release Notes pages on the HPE Support website which is the most recent available version.

I have already some customers meanwhile running HPE SimpliVity with HPE VM Essentials as hypervisor in production here in Belgium (February 2026).

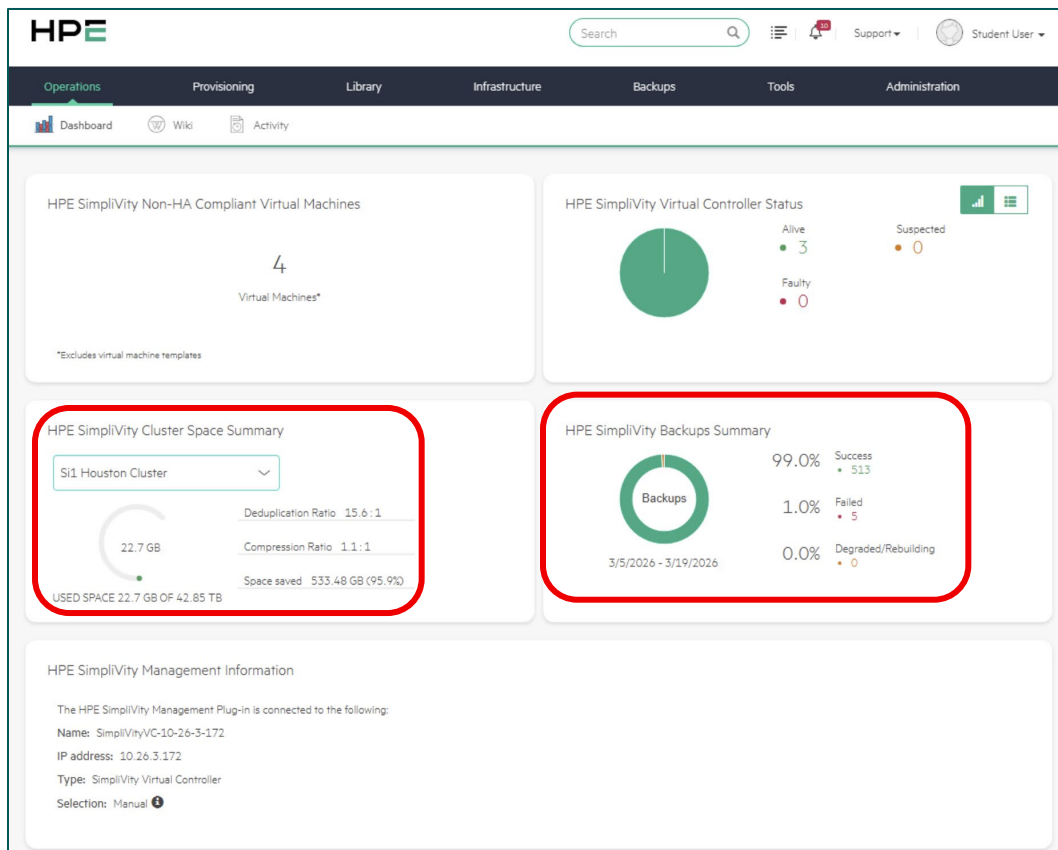
The management experience of the standalone deployment of the HPE VM Essentials Manager on a 3-tier environment is exactly the same as when deployed on HPE SimpliVity.

The ‘only’ difference is the SimpliVity Virtual Controller virtual appliance VM that is running on each HPE SimpliVity node in the cluster to run the Data Virtualization Platform stack.



If we compare both HPE SimpliVity solutions with VMware and HPE VM Essentials as hypervisor you see that only the hypervisor and management stack is changed, being the HPE Morpheus software stack.

Unlike the management experience with VMware, where a plug-in was used inside VMware vCenter, and a separate Windows-based tool was used for the deployment and upgrade of the HPE SimpliVity cluster, in the HPE VM Essentials solution you will see everything will happen through the HPE VM Essentials Manager UI.

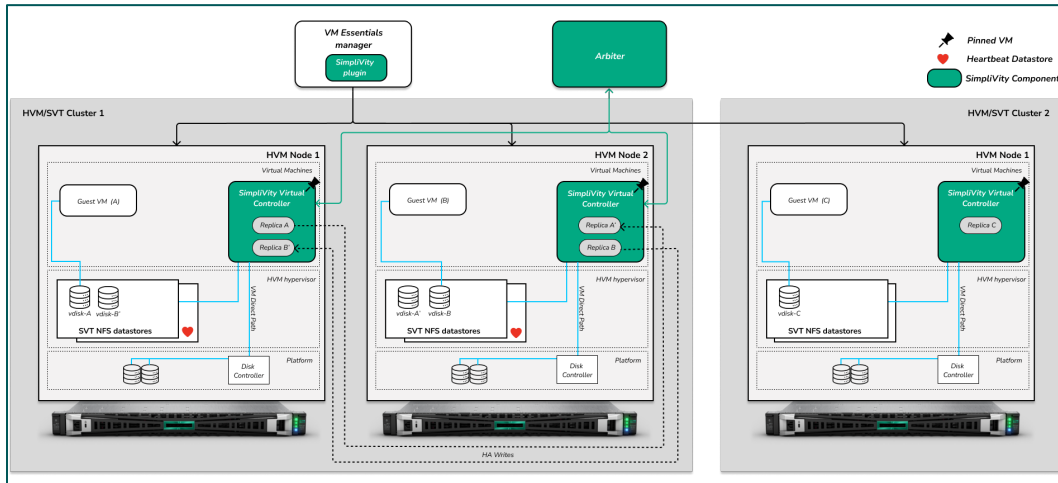


All the typical HPE SimpliVity features like data efficiency (deduplication and compression), use of NFS datastores and integrated data protection remain exactly the same, independent of the hypervisor being used under the hood.

Some important things to know however when sizing and installing HPE VM Essentials in combination with HPE SimpliVity.

With the initial version 6.0.0 and now version 6.2.0 released (at the moment of publication of this guide), it is required that the HPE VM Essentials Manager must be hosted outside the HVM cluster, so on a third server (assuming a 2-node cluster).

Note that multiple clusters are supported, like in the following picture which is a 2+1 deployment (very popular since HA and DR are optimally used in a 3node setup):



This extra server will be running Ubuntu 22.04 and will be used to run the HPE VM Essentials Manager and also the arbitrer software that will act as a quorum witness, very similar to a VMware based SimpliVity deployment.

Note however that the arbitrer software is not a Windows service, but is based on Linux now. There is also a Windows arbitrer available with the release of 6.2.0.

These are the arbitrer requirements:

Feature	Requirement
CPU	1.1 GHz or higher (recommended : 2 GHz)
Number of cores	Single core / dual thread or better
Memory	1 GB (for Arbiter) - consult Ubuntu documentation for OS memory requirements (minimum : 2 GB, recommended : 4 GB)
Persistent storage	Minimum: 25 GB, recommended : 40 GB Set policy to write-through or write-cache disabled
Network	1x 1Gb, network latency to all nodes: max. 300ms RTT Maximum latency between HPE SimpliVity nodes in local cluster: 2 ms
Supported operating system	Ubuntu 22.04.5 LTS (Jammy Jellyfish) 64-bit
Port	Port 22122 / tcp

Definitely keep on checking upcoming releases, I will update this section accordingly when more information on features and functions becomes available.

Useful Linux commands

This is a list of Linux commands that I used during the exploration and troubleshooting of HPE VM Essentials through all the last months so far...

Hope you will never need them of course but let it be a reference in your VM Essentials journey.

A lot of these commands require sudo permissions. Or you can add in front of each command sudo, or just run once **sudo -i** and you are OK for that entire session.

Special thanks to Peter Jones and Dean Colpitts for assistance with this list!

VIRSH

	Command	Possible output
Virsh information		
Display libvirt version	virsh version	# virsh version Compiled against library: libvirt 10.0.0 Using library: libvirt 10.0.0 Using API: QEMU 10.0.0 Running hypervisor: QEMU 8.2.2
Display libvirtd version	virsh version --daemon	# virsh version --daemon Compiled against library: libvirt 10.0.0 Using library: libvirt 10.0.0 Using API: QEMU 10.0.0 Running hypervisor: QEMU 8.2.2 Running against daemon: 10.0.0
Display virsh help	virsh help	
VM information		
Display VM ID number	virsh domid <vm-name>	# virsh domid WindowsVM1 2497

Display VM UUID	virsh domuuid <vm-name>	# virsh domuuid WindowsVM1 8c7e5e96-7a7c-4b41-91e0-ec66d22bbc06
List running VMs on hypervisor	virsh list	# virsh list Id Name State ----- 2497 WindowsVM1 running 2498 LinuxVM1 running
List all VMs on hypervisor	virsh list --all	# virsh list --all Id Name State ----- 2497 WindowsVM1 running 2498 LinuxVM1 running - VeeamWorker shut off
List inactive VMs on hypervisor	virsh list --inactive	# virsh list --inactive Id Name State ----- - VeeamWorker shut off
Get status or VM - running or stopped	virsh domstate <vm-name>	# virsh domstate WindowsVM1 running
Display connection state to VM	virsh domcontrol <vm-name>	# virsh domcontrol WindowsVM1 ok
Get general VM information	virsh dominfo <vm-name>	# virsh dominfo WindowsVM1 Id: 2497 Name: WindowsVM1 UUID: 8c7e5e96-7a7c-4b41-91e0-ec66d22bbc06 OS Type: hvm State: running CPU(s): 2 CPU time: 1655.6s Max memory: 8388608 KiB Used memory: 8388608 KiB Persistent: yes Autostart: disable Managed save: no Security model: apparmor Security DOI: 0

		Security label: libvirt-8c7e5e96-7a7c-4b41-91e0-ec66d22bbc06 (enforcing) Messages: tainted: custom configuration parameters specified
Dump XML config of VM	virsh dumpxml <vm-name>	# virsh dumpxml WindowsVM1 <domain type='kvm' id='2497'> <name>WindowsVM1</name> <uuid>8c7e5e96-7a7c-4b41-91e0-ec66d22bbc06</uuid> <maxMemory slots='240' unit='KiB'>131966976</maxMemory> <memory unit='KiB'>16777216</memory> OUTPUT OMITTED
VM CPU Information		
VM CPU info	virsh vcpuinfo <vm-name>	# virsh vcpuinfo WindowsVM1 VCPU: 0 CPU: 6 State: running CPU time: 909.5s CPU Affinity: yyyyyyyy VCPU: 1 CPU: 7 State: running CPU time: 640.9s CPU Affinity: yyyyyyyy
VM CPU affinity info	virsh vcpupin <vm-name>	# virsh vcpupin WindowsVM1 VCPU CPU Affinity ----- 0 0-7 1 0-7 2 0-7 3 0-7 OUTPUT OMITTED
VM Memory Information		

VM memory info	virsh dommemstat <vm-name>	# virsh dommemstat WindowsVM1 actual 8388608 last_update 1767648240 rss 8485188
VM Disk Device Information		
VM block device list	virsh domblklist <vm-name>	# virsh domblklist WindowsVM1 Target Source ----- ----- vda /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/WindowsVM1/hvm_15-disk-0-ebfb4c4 sda -
VM block device list and associations	virsh domblklist <vm-name> --details	# virsh domblklist WindowsVM1 --details Type Device Target Source ----- ----- file disk vda /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/WindowsVM1/hvm_15-disk-0-ebfb4c4 file cdrom sda -
VM block device info	virsh domblkinfo <vm-name> <target> --human	# virsh domblkinfo WindowsVM1 vda --human Capacity: 80.000 GiB Allocation: 12.400 GiB Physical: 12.400 GiB
VM block device stats	virsh domblkstat <vm-name> <target> --human	# virsh domblkstat WindowsVM1 vda --human Device: vda number of read operations: 32033 number of bytes read: 1391828992 number of write operations: 80971 number of bytes written: 2392014848 number of flush operations: 41541 total duration of reads (ns): 105948196985 total duration of writes (ns): 130871332062 total duration of flushes (ns): 6811998742

VM block device errors	virsh domblkerror <vm-name>	# virsh domblkerror WindowsVM1 No errors found
VM Network Information		
VM network interface stats	virsh domiflist <vm-name>	# virsh domiflist WindowsVM1 Interface Type Source Model MAC ----- vnet2496 bridge Management virtio 52:54:00:0a:ab:b7
Confirm if specific network interface is up	virsh domif-getlink <vm-name> <MAC>	# virsh domif-getlink WindowsVM1 vnet2496 vnet2496 up
VM network interface stats for device	virsh domifstat <vm-name> ,interface>	# virsh domifstat WindowsVM1 vnet2496 vnet2496 rx_bytes 246220710 vnet2496 rx_packets 1247942 vnet2496 rx_errs 0 vnet2496 rx_drop 0 vnet2496 tx_bytes 2583193 vnet2496 tx_packets 24838 vnet2496 tx_errs 0 vnet2496 tx_drop 0
Start and Stop VMs		
Start a VM	virsh start <vm-name>	
Shutdown a VM (gracefully)	virsh shutdown <vm-name>	
Shutdown a VM (hard)	virsh destroy <vm-name>	
Suspend a VM	virsh suspend <vm-name>	
Resume a suspended VM	virsh resume <vm-name>	

Reboot a VM	virsh reboot <vm-name>	
Reset a VM	virsh reset <vm-name>	
Host Information		
Display hypervisor hostname	virsh hostname	# virsh hostname vme2
Display XML representation of hypervisor's system information	virsh sysinfo	# virsh sysinfo <sysinfo type='smbios'> <bios> <entry name='vendor'>HPE</entry> <entry name='version'>2.30</entry> <entry name='date'>08/07/2025</entry> <entry name='release'>2.30</entry> </bios> <system> <entry name='manufacturer'>HPE</entry> <entry name='product'>ProLiant MicroServer Gen11</entry> <entry name='version'>Not Specified</entry> <entry name='serial'>CZJD1L01B3</entry> <entry name='uuid'>38383650-3132-5a43-4a44-314c30314233</entry> <entry name='sku'>P68821-421</entry> <entry name='family'>ProLiant</entry> </system> <baseBoard> <entry name='manufacturer'>HPE</entry> <entry name='product'>ProLiant MicroServer Gen11</entry> OUTPUT OMITTED
Hypervisor node information	virsh nodeinfo	# virsh nodeinfo CPU model: x86_64 CPU(s): 8 CPU frequency: 1400 MHz CPU socket(s): 1 Core(s) per socket: 4

		Thread(s) per core: 2 NUMA cell(s): 1 Memory size: 65744096 KiB
Display total amount of available memory in all NUMA cells	virsh freecell	# virsh freecell Total: 88049968 KiB
Display amount of available memory in individual cells	virsh freecell --all	# virsh freecell --all 0: 38539316 KiB 1: 49495276 KiB ----- Total: 88034592 KiB
Display number of CPUs in host	virsh nodecpumap	# virsh nodecpumap CPUs present: 8 CPUs online: 8 CPU map: yyyyyyyy
Display node cpu stats	virsh nodecpustats	# virsh nodecpustats user: 96305890000000 system: 92628020000000 idle: 1920821604000000 iowait: 16594060000000
Display node cpu stats for specific CPU as a percent	virsh nodecpustats <cpu_number> <percent> --percent	# virsh nodecpustats 4 --percent usage: 1.0% user: 1.0% system: 0.0% idle: 99.0% iowait: 0.0%
List the devices on the host	virsh nodedev-list --tree	# virsh nodedev-list --tree computer +- net_lo_00_00_00_00_00_00 +- net_mgmt_2a_90_2e_33_6f_43 +- net_ovs_system_7e_82_30_63_88_3c +- net_vnet2496_fe_54_00_0a_ab_b7 +- net_vnet2497_fe_54_00_17_7b_50 +- pci_0000_00_00_0

		<pre> +- pci_0000_00_14_0 +- usb_usb2 +- usb_2_0_1_0 +- usb_2_3 +- usb_2_3_1_0 OUTPUT OMITTED </pre>
Displays XML representation of hypervisor connection and host physical machine capabilities	virsh capabilities	<pre> # virsh capabilities <capabilities> <host> <uuid>38383650-3132-5a43-4a44-314c30314233</uuid> <cpu> <arch>x86_64</arch> <model>Skylake-Client-noTSX-IBRS</model> <vendor>Intel</vendor> <microcode version='303'/> <signature family='6' model='183' stepping='1'/> <counter name='tsc' frequency='341760000' scaling='yes'/> <topology sockets='1' dies='1' cores='4' threads='2'/> <maxphysaddr mode='emulate' OUTPUT OMITTED </pre>
Display the XML representation of qemu-kvm capabilities	virsh domcapabilities	<pre> # virsh domcapabilities <domainCapabilities> <path>/usr/bin/qemu-system-x86_64</path> <domain>kvm</domain> <machine>pc-i440fx-noble</machine> <arch>x86_64</arch> <vcpu max='255'/> <iothreads supported='yes'/> <os supported='yes'> <enum name='firmware'> <value>efi</value> </enum> <loader supported='yes'> <value>/usr/share/OVMF/OVMF_CODE_4M.fd</value> <enum name='type'> </pre>

		OUTPUT OMITTED
Host Volumes		
List storage pools	virsh pool-list	<pre># virsh pool-list Name State Autostart ----- local active yes morpheus-cloud-init active yes morpheus-images active yes NFSONNAS active yes</pre>
List all storage pools that are both active and inactive	virsh pool-list --all	<pre># virsh pool-list --all Name State Autostart ----- local active yes morpheus-cloud-init active yes morpheus-images active yes NFSONNAS active yes</pre>
List volumes in a storage pool	virsh vol-list <pool-name>	<pre># virsh vol-list NFSONNAS Name Path ----- ----- Migrate-Me /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/Migrate-Me morpheus-virtual-images /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/morpheus-virtual-images VeeamWorker /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/VeeamWorker virtio-win-2025.iso /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/virtio-win-2025.iso vmem /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/vmem WindowsVM1 /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887/WindowsVM1</pre>
Host Networks		
List available networks and state	virsh net-list --all	<pre># virsh net-list --all Name State Autostart Persistent -----</pre>

		Management active yes yes
Display the XML representation of a named network	virsh net-dumpxml <network-name>	# virsh net-dumpxml Management <network connections='2'> <name>Management</name> <uuid>f816d40a-a195-4c03-b98e-08da3488ad2d</uuid> <forward mode='bridge'/> <bridge name='mgmt'/> <virtualport type='openvswitch'/> </network>

General Linux commands

	Command	Possible output
System Information		
Provides metrics on memory, block IO and CPU	vmstat	<pre>administrator@vme1:/\$ vmstat procs -----memory----- --swap-- ----io---- -system-- -----cpu----- r b swpd free buff cache si so bi bo in cs us sy id wa st gu 0 0 768 11502116 1025868 51474296 0 0 0 12 2143 4 0 0 98 0 0 1 administrator@vme1:/\$</pre>
Display kernel version, architecture and system information	uname -a	<pre>administrator@vme1:/\$ uname -a Linux vme1 6.8.0-84-generic #84-Ubuntu SMP PREEMPT_DYNAMIC Fri Sep 5 22:36:38 UTC 2025 x86_64 x86_64 x86_64 GNU/Linux administrator@vme1:/\$</pre>
Shows how long system has been running	uptime	<pre>administrator@vme1:/\$ uptime 07:30:26 up 41 days, 23:57, 1 user, load average: 0.13, 0.24, 0.29 administrator@vme1:/\$</pre>

Show who is logged in and when they logged in	who	<pre>administrator@vmel:/\$ who administrator pts/0 2026-01-21 03:31 (192.168.131.99) administrator@vmel:/\$</pre>
Show who is logged in, from where and what they are doing	w	<pre>administrator@vmel:/\$ w 07:28:06 up 41 days, 23:54, 1 user, load average: 0.28, 0.28, 0.31 USER TTY FROM LOGIN@ IDLE JCPU PCPU WHAT administ pts/0 192.168.131.99 03:31 21days 0.00s 0.02s sshd: administrator [priv] administrator@vmel:/\$</pre>
Disks		
list block devices. Shows installed disks and partitions	lsblk	<pre>administrator@vmel:/\$ lsblk NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINTS sda 8:0 0 894.3G 0 disk ├─sda1 8:1 0 1G 0 part /boot/efi ├─sda2 8:2 0 2G 0 part /boot ├─sda3 8:3 0 891.2G 0 part │ └─ubuntu--vg-ubuntu--lv 252:0 0 890G 0 lvm / sdb 8:16 0 931.5G 0 disk ├─sdb1 8:17 0 1G 0 part ├─sdb2 8:18 0 2G 0 part ├─sdb3 8:19 0 928.5G 0 part │ └─ubuntu--vg--1-ubuntu--lv 252:1 0 900G 0 lvm sde 8:64 0 40G 0 disk ├─360002ac000000000000004960002ab4b 252:3 0 40G 0 mpath ├─360002ac000000000000004960002ab4b-part1 252:2 0 1M 0 part ├─360002ac000000000000004960002ab4b-part2 252:4 0 1G 0 part ├─360002ac000000000000004960002ab4b-part3 252:5 0 39G 0 part ├─r1-swap 252:6 0 3.9G 0 lvm └─r1-root 252:7 0 35.1G 0 lvm sdf 8:80 0 40G 0 disk ├─360002ac000000000000004960002ab4b 252:3 0 40G 0 mpath ├─360002ac000000000000004960002ab4b-part1 252:2 0 1M 0 part ├─360002ac000000000000004960002ab4b-part2 252:4 0 1G 0 part ├─360002ac000000000000004960002ab4b-part3 252:5 0 39G 0 part ├─r1-swap 252:6 0 3.9G 0 lvm └─r1-root 252:7 0 35.1G 0 lvm administrator@vmel:/\$</pre>
Prints all SMART (Self-Monitoring, Analysis and Reporting Technology) information about the disk including disk device information, health and errors. For a given device	smartctl -a	
Disk space usage		

Displays disk space usage on all mounted filesystems	df	<pre> administrator@vmel:/\$ df Filesystem 1K-blocks Used Available Use% Mounted on tmpfs 6574412 1264 6573148 1% /run efivarfs 494 271 219 56% /sys/firmware/efi/efivars /dev/mapper/ubuntu--vg-ubuntu--lv 917458932 13880992 856899924 2% / tmpfs 32872048 33000 32839048 1% /dev/shm tmpfs 5120 0 5120 0% /run/lock tmpfs 32872048 0 32872048 0% /run/gemu /dev/sda2 1992552 102976 1768336 6% /boot /dev/sda1 1098628 6288 1092340 1% /boot/efi 192.168.131.102:/volume1/NFS4VME 2996444392 22411043200 7553400192 75% /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887 tmpfs 6574408 16 6574392 1% /run/user/1000 administrator@vmel:/\$ </pre>
Displays disk space usage on all mounted filesystems in a human readable format	df -h	<pre> administrator@vmel:/\$ df -h Filesystem Size Used Avail Use% Mounted on tmpfs 6.3G 1.3M 6.3G 1% /run efivarfs 494K 271K 219K 56% /sys/firmware/efi/efivars /dev/mapper/ubuntu--vg-ubuntu--lv 875G 14G 810G 2% / tmpfs 32G 32M 32G 1% /dev/shm tmpfs 5.0M 0 5.0M 0% /run/lock tmpfs 32G 0 32G 0% /run/gemu /dev/sda2 2.0G 101M 1.7G 6% /boot /dev/sda1 1.1G 6.2M 1.1G 1% /boot/efi 192.168.131.102:/volume1/NFS4VME 28T 21T 7.1T 75% /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887 tmpfs 6.3G 16K 6.3G 1% /run/user/1000 administrator@vmel:/\$ </pre>
Display File System Type in use	df -T	<pre> administrator@vmel:/\$ df -T Filesystem Type 1K-blocks Used Available Use% Mounted on tmpfs tmpfs 6574412 1264 6573148 1% /run efivarfs efivarfs 494 271 219 56% /sys/firmware/efi/efivars /dev/mapper/ubuntu--vg-ubuntu--lv ext4 917458932 13880472 856900444 2% / tmpfs tmpfs 32872048 33000 32839048 1% /dev/shm tmpfs tmpfs 5120 0 5120 0% /run/lock tmpfs tmpfs 32872048 0 32872048 0% /run/gemu /dev/sda2 ext4 1992552 102976 1768336 6% /boot /dev/sda1 vfat 1098628 6288 1092340 1% /boot/efi 192.168.131.102:/volume1/NFS4VME nfs 2996444392 22411043200 7553400192 75% /mnt/cfe5caf8-ece0-422e-9b05-bd762f03f887 tmpfs tmpfs 6574408 16 6574392 1% /run/user/1000 administrator@vmel:/\$ </pre>
Folder size breakdown for that given folder where you are located	du -hs *	
Swap		
Print swap devices	swapon --summary	<pre> administrator@vmel:/\$ swapon --summary Filename Type Size Used Priority /swap.img file 8388604 768 -2 administrator@vmel:/\$ </pre>
Check the swap columns si and so. si is amount of memory swapped in from disk, and so is amount of memory swapped to disk	vmstat	<pre> administrator@vmel:/\$ vmstat procs-----memory----- --swap-- -----io----- -system-- -----cpu----- r b swpd free buff cache si so bi bo in cs us sy id wa st gu 0 0 0 768 11502116 1025868 51474296 0 0 0 12 2143 4 0 0 98 0 0 1 administrator@vmel:/\$ </pre>
Display total used free swap in GB	free -h	<pre> administrator@vmel:/\$ free -h total used free shared buff/cache available Mem: 62Gi 22Gi 10Gi 20Gi 50Gi 40Gi Swap: 8.0Gi 768Ki 8.0Gi administrator@vmel:/\$ </pre>

<p>Display total used free swap in header of output</p>	<p>top</p>	<pre> administrator@vmel:/\$ top top - 07:38:24 up 42 days, 5 min, 1 user, load average: 0.01, 0.09, 0.18 Tasks: 241 total, 2 running, 239 sleeping, 0 stopped, 0 zombie %Cpu(s): 0.7 us, 0.8 sy, 0.0 ni, 98.4 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st MiB Mem : 64203.2 total, 11222.3 free, 22965.6 used, 51269.9 buff/cache MiB Swap: 8192.0 total, 8191.2 free, 0.8 used, 41237.6 avail Mem PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND 3320248 libvirt+ 20 0 17.7g 16.1g 16.0g S 7.0 25.7 39.41 qemu-system-x86 3319386 libvirt+ 20 0 8480212 4.1g 4.0g S 1.3 6.5 320:41.09 qemu-system-x86 1013 root 10 -10 901168 117352 13952 S 0.7 0.2 180:44.23 ovs-vsitchd 1647 root 0 574280 181552 65788 S 0.7 0.3 6.38 corosync 1833 haclust+ 20 0 70016 14940 11484 S 0.3 0.0 15:09.10 pacemaker-attd 1835 haclust+ 20 0 90184 22336 12936 S 0.3 0.0 16:49.08 pacemaker-contr 1847 root 20 0 487972 68436 25984 S 0.3 0.1 226:46.31 pcsd 2512403 morpheu+ 20 0 6828168 458280 22976 S 0.3 0.7 45:30.03 java 3319398 root 20 0 0 0 0 S 0.3 0.0 4:24.01 kvm-pit/3319386 1 root 20 0 23088 13800 9192 S 0.0 0.0 10:05.77 systemd 2 root 20 0 0 0 0 S 0.0 0.0 0:00.60 kthreadd 3 root 20 0 0 0 0 S 0.0 0.0 0:00.00 pool workqueue release </pre> <p>OUTPUT OMITTED</p>
<p>Displays total and used swap in header of output (user friendly display)</p>	<p>htop</p>	<pre> administrator@vmel:/\$ htop </pre> <p>OUTPUT OMITTED</p>
<p>Memory</p>		
<p>Shows total used free memory in bytes</p>	<p>free -m</p>	<pre> administrator@vmel:/\$ free -m total used free shared buff/cache available Mem: 64203 22974 11210 20516 51272 41228 Swap: 8191 0 8191 administrator@vmel:/\$ </pre>
<p>Run free command every 5 seconds</p>	<p>free -s 5</p>	<pre> administrator@vmel:/\$ free -s 5 total used free shared buff/cache available Mem: 65744100 23526080 11479908 21008892 52502728 42218020 Swap: 8388604 768 8387836 total used free shared buff/cache available Mem: 65744100 23520508 11485476 21008892 52502728 42223592 Swap: 8388604 768 8387836 ^C administrator@vmel:/\$ </pre>

<p>Sort processes by memory</p>	<p>top -o +%MEM</p>	<pre> administrator@vmel:/\$ top -o +%MEM top - 07:42:04 up 42 days, 8 min, 1 user, load average: 0.21, 0.17, 0.19 Tasks: 241 total, 1 running, 240 sleeping, 0 stopped, 0 zombie %Cpu(s): 1.6 us, 1.9 sy, 0.0 ni, 96.5 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st MiB Mem : 64203.2 total, 11214.1 free, 22971.4 used, 51272.2 buff/cache MiB Swap: 8192.0 total, 8191.2 free, 0.8 used, 41231.8 avail Mem PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND 3320248 libvirt+ 20 0 17.7g 16.1g 16.0g S 15.3 25.7 39.42 gemu-system-x86 3319386 libvirt+ 20 0 8480212 4.1g 4.0g S 1.3 6.5 320:44.75 gemu-system-x86 2512403 morpheu+ 20 0 6828168 458280 22976 S 0.0 0.7 45:30.35 java 1647 root rt 0 574280 181552 65788 S 0.7 0.3 6.38 corosync 1013 root 10 -10 901168 117352 13952 S 0.3 0.2 180:44.90 ovs-vsitchd 1688 root 20 0 3062948 75764 35456 S 0.0 0.1 230:18.25 libvirtd 1651 root 20 0 811292 71432 14976 S 0.0 0.1 22:14.19 ruby3.2 1847 root 20 0 487972 68436 25984 S 0.7 0.1 226:47.09 pcsd 1877 root 20 0 1075644 46904 6144 S 0.3 0.1 82:54.60 pcsd 1907 root 20 0 190692 46184 5504 S 0.0 0.1 0:00.00 pcsd 1889 root 20 0 190692 46172 5504 S 0.0 0.1 0:00.00 pcsd 1919 root 20 0 190692 46072 5376 S 0.0 0.1 0:00.00 pcsd 1914 root 20 0 190692 46068 5376 S 0.0 0.1 0:00.00 pcsd 1904 root 20 0 190692 46056 5376 S 0.0 0.1 0:00.00 pcsd 1912 root 20 0 190692 46056 5376 S 0.0 0.1 0:00.00 pcsd 1887 root 20 0 190692 46044 5376 S 0.0 0.1 0:00.00 pcsd 1899 root 20 0 190692 45924 5248 S 0.0 0.1 0:00.00 pcsd 1892 root 20 0 190692 45916 5248 S 0.0 0.1 0:00.00 pcsd 1895 root 20 0 190692 45916 5248 S 0.0 0.1 0:00.00 pcsd 1268 root 20 0 265876 30160 15360 S 0.0 0.0 14:03.25 tuned 485 root rt 0 420732 27904 8576 S 0.0 0.0 6:32.46 multipathd 1830 haclust+ 20 0 82160 27008 16512 S 0.0 0.0 15:03.75 pacemaker-based 1835 haclust+ 20 0 90184 22336 12936 S 0.0 0.0 16:49.14 pacemaker-contr 426 root 19 -1 84668 22292 20028 S 0.0 0.0 6:07.51 systemd-journal </pre>
<p>OUTPUT OMITTED</p>		
<p>Shows information about memory</p>	<p>cat /proc/meminfo</p>	<pre> administrator@vmel:/\$ cat /proc/meminfo MemTotal: 65744100 kB MemFree: 11486032 kB MemAvailable: 42224152 kB Buffers: 1025904 kB Cached: 50076424 kB SwapCached: 12 kB Active: 29203840 kB Inactive: 22671952 kB Active(anon): 21807600 kB Inactive(anon): 296 kB Active(file): 7396240 kB Inactive(file): 22671656 kB Unevictable: 297388 kB Mlocked: 294316 kB SwapTotal: 8388604 kB SwapFree: 8387836 kB Zswap: 0 kB Zswapped: 0 kB Dirty: 176 kB </pre>
<p>OUTPUT OMITTED</p>		
<p>CPU</p>		

<p>Show CPU architecture information. Shows number of CPUs, cores and sockets</p>	<p>lscpu</p>	<pre> administrator@vmel:/\$ lscpu Architecture: x86_64 CPU op-mode(s): 32-bit, 64-bit Address sizes: 46 bits physical, 48 bits virtual Byte Order: Little Endian CPU(s): 8 On-line CPU(s) list: 0-7 Vendor ID: GenuineIntel Model name: Intel(R) Xeon(R) E E-2434 CPU family: 6 Model: 183 Thread(s) per core: 2 Core(s) per socket: 4 Socket(s): 1 Stepping: 1 CPU(s) scaling MHz: 26% CPU max MHz: 5000.0000 CPU min MHz: 800.0000 BogoMIPS: 6835.20 Flags: fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clfl s bts rep_good nopl xtopology nonstop_tsc cpuid aperfperf tsc_known_freq pni movbe popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm 3dnowp e tsc_adjust bmi1 avx2 smep bmi2 erms invpcid rdseed adx smap clflushopt clw ush_lid arch_capabilities Virtualization features: Virtualization: VT-x Caches (sum of all): L1d: 192 KiB (4 instances) L1i: 128 KiB (4 instances) L2: 8 MiB (4 instances) L3: 12 MiB (1 instance) NUMA: NUMA node(s): 1 NUMA node0 CPU(s): 0-7 Vulnerabilities: gather_data_sampling: Not affected </pre> <p>OUTPUT OMITTED</p>
<p>Shows information about each individual CPU</p>	<p>cat /proc/cpuinfo</p>	<pre> administrator@vmel:/\$ cat /proc/cpuinfo processor : 0 vendor_id : GenuineIntel cpu family : 6 model : 183 model name : Intel(R) Xeon(R) E E-2434 stepping : 1 microcode : 0x12f cpu MHz : 1189.732 cache size : 12288 KB physical id : 0 siblings : 8 core id : 0 cpu cores : 4 apicid : 0 initial apicid : 0 fpu : yes fpu exception: yes cpuid level : 32 wp : yes flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush dts a p_good nopl xtopology nonstop_tsc cpuid aperfperf tsc_known_freq pni pclmulqdq dtes64 monitor ds_cpl v s xsave avx f16c rdrand lahf_lm abm 3dnowprefetch cpuid_fault epb sbsd ibrs lbrb stibp ibrs_enhanced tpr p clflushopt clwb intel_pt sha_ni xsaveopt xsavec xgetbv1 xsaves split_lock_detect user_shstk avx_vnmi gfni vaes vpclmulqdq tme rdpid movdiri movdir64b fsrm md_clear serialize_pconfig arch_lbr ibt flush_ll vmx flags : vnmi preemption_timer posted_intr invvpid ept_x_only ept_ad ept_lgb flexpriority apic exec tsc_scaling usr_wait_pause bugs : spectre_v1 spectre_v2 spec_store_bypass swapgs eibrs pbrsb bh1 bogomips : 6835.20 clflush size: 64 cache_alignm: 64 address sizes: 46 bits physical, 48 bits virtual power management: processor : 1 vendor_id : GenuineIntel cpu family : 6 model : 183 model name : Intel(R) Xeon(R) E E-2434 </pre> <p>OUTPUT OMITTED</p>
<p>Print the number of processors</p>	<p>nproc</p>	<pre> administrator@vmel:/\$ nproc 8 administrator@vmel:/\$ </pre>

Check services are running		
	corosync	<pre>root@me1:~# corosync Jan 21 07:58:55 notice [MAIN] Corosync Cluster Engine 3.1.7 starting up Jan 21 07:58:55 info [MAIN] Corosync built-in features: dbus monitoring watchdog susegs systemd xmlconf vsnism nozlie snmp pie reiro bindnow root@me1:~#</pre>
	systemctl grep running grep ovs	<pre>root@me1:~# systemctl grep running grep ovs ovs-vsftchd.service loaded active running Open vSwitch Forwarding Unit ovsdb-server.service loaded active running Open vSwitch Database Unit root@me1:~#</pre>
	systemctl grep running grep libvirt	<pre>root@me1:~# systemctl grep running grep libvirt libvirtd.service loaded active running libvirt legacy monolithic daemon virtlogd.service loaded active running libvirt logging daemon virtlogd.socket loaded active running libvirt logging daemon admin socket libvirtd-admin.socket loaded active running libvirt legacy monolithic daemon admin socket libvirtd-ro.socket loaded active running libvirt legacy monolithic daemon read-only socket libvirtd.socket loaded active running libvirt legacy monolithic daemon socket virtlogd-admin.socket loaded active running libvirt logging daemon admin socket virtlogd.socket loaded active running libvirt logging daemon socket root@me1:~#</pre>

Useful links

Main product page

<https://www.hpe.com/us/en/morpheus-vm-essentials-software.html>

General Support page

<https://www.hpe.com/support/vme-docs>

Software Documentation Guide

<https://hpevm-docs.morpheusdata.com/en/latest/>

Qualification Matrix

<https://www.hpe.com/support/VME-Compatibility-Matrix>

Reference Architecture

<https://www.hpe.com/support/morpheus-vm-essentials-deployment-guide>

Migration Guide

<https://www.hpe.com/psnow/doc/a50013873enw>

30-day 6-socket trial license

https://myenterpriselicense.hpe.com/cwp-ui/product-details/HPE_VME_EVAL/-/sw360_eval_customer

My video's on HPE VM Essentials

<https://www.bitcon.be/multimedia/>

An up-to-date list can be always found on the VM Essentials page at www.hpepedia.com.

Versioning:

1.0 Reference document based on HPE Morpheus VM Essentials 8.1.0