



HPE NFV Director

Performance and Sizing Guide

Release 4.2

First Edition



Hewlett Packard
Enterprise

Notices

Legal notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Confidential computer software. Valid license from HPE 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.

The information contained herein is subject to change without notice. The only warranties for HPE 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. HPE shall not be liable for technical or editorial errors or omissions contained herein.

Printed in the US

Trademarks

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

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

Microsoft®, Internet Explorer, Windows®, Windows Server 2007®, Windows XP®, and Windows 7® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Firefox® is a registered trademark of the Mozilla Foundation.

Google Chrome® is a trademark of Google Inc.

EnterpriseDB® is a registered trademark of EnterpriseDB.

Postgres Plus® Advanced Server is a registered U.S. trademark of EnterpriseDB.

UNIX® is a registered trademark of The Open Group.

X/Open® is a registered trademark, and the X device is a trademark of X/Open Company Ltd. in the UK and other countries.

Red Hat® is a registered trademark of the Red Hat Company.

Apache CouchDB, CouchDB, and the project logo are trademarks of The Apache Software Foundation.

Node.js project. Joyent® and Joyent's logo are registered trademarks of Joyent, Inc.

Neo4j is a trademark of Neo Technology.

VMware ESX, VMWare ESXi, VMWare vCenter and VMWare vSphere are either registered trademarks or trademarks of VMware, Inc. in the United States and other jurisdictions..

Contents

Notices	1
Preface	5
Document history.....	5
Chapter 1 Introduction	6
1.1 Glossary of Terms.....	6
1.2 References.....	6
Chapter 2 Test Environment	7
Chapter 3 Performance Summary.....	10
3.1 Monolithic Deployment of N VNFs.....	10
3.2 Partitioned Deployment of N VNFs.....	10
3.3 Dependency on the Number of Compute Nodes.....	11
3.4 Parallel Deployments	11
3.5 Recommendations.....	12
Chapter 4 Sizing Guidelines	13
4.1 Small Configuration	13
4.2 Medium Configuration.....	14

List of tables

Table 1: Document history	5
Table 2: Small NFV-D sizing	14
Table 3: Medium NFV-D sizing	14

List of figures

Figure 1: NFV-D Test Deployment.....	7
Figure 2: Helion Carrier Grade Configuration.....	8
Figure 3: Simulated Compute Nodes in HCG.....	9

Preface

This document provides a summary of the results of NFV-D performance testing. It provides recommendations of how to organize VNFs in order to maximize performance. It then gives sample recommendations for hardware configurations based on the number of VNFs and VMs to be deployed.

Document history

Table 1: Document history

Edition	Date	Description
1.0	15 Feb 2017	First version.

Chapter 1

Introduction

1.1 Glossary of Terms

API	Application Programming Interface
DB	Database
LDAP	Lightweight Directory Access Protocol
UCA	Unified correlation analyzer
MWFM	Micro Work Flow Manager
WF	Workflow
IDP	Identity Provider
NFV-D	Network Function Virtualization-Director
VNF	Virtualized Network Function
VDC	Virtual Data Center

1.2 References

The following references are used for the preparation of this document. All of these documents can be found on the HPE software support site: <https://softwaresupport.hpe.com>

1	NFV-D 4.2 Installation Guide
2	hp_man_ServiceActivator6_20_Linux_Installation.pdf
3	SIS_11.31_SiteScopeDeployment.pdf
4	UCA Automation 1.2 Installation Guide
5	HP UCA for Event Based Correlation V3.1 - Installation Guide V1.0.pdf

Chapter 2

Test Environment

The tests were conducted using the HPE CMS R&D test environment. The details of this environment are illustrated in the following two figures.

The first figure describes the deployment of NFV-D itself:

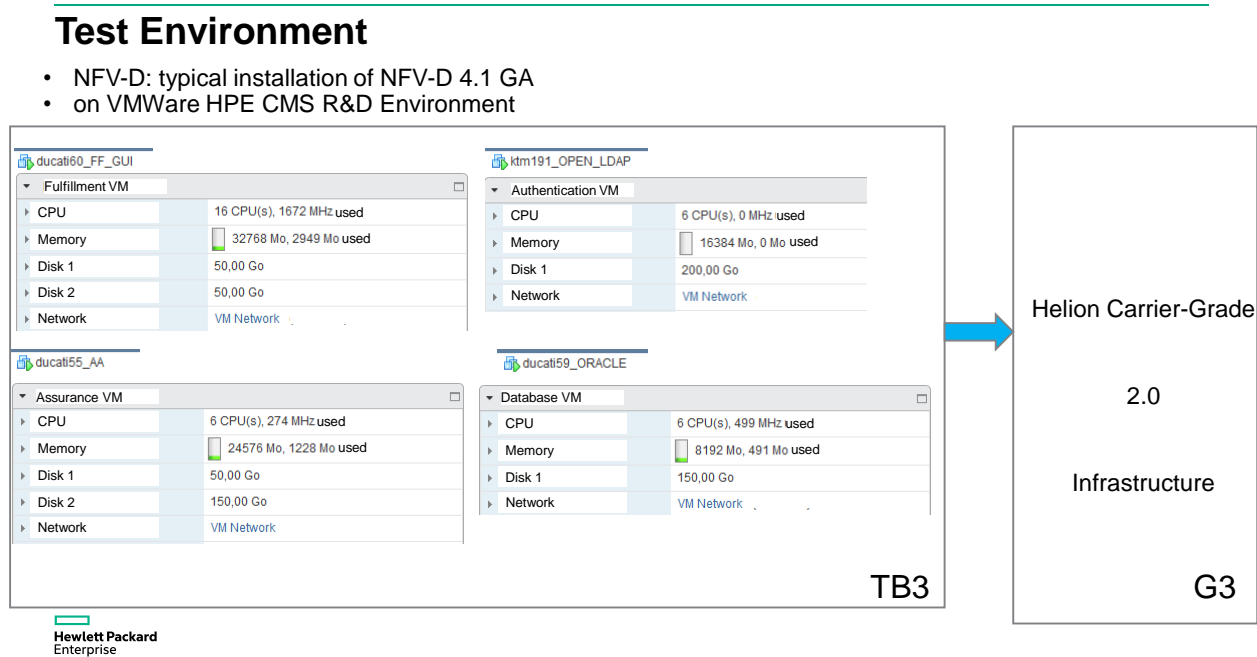


Figure 1: NFV-D Test Deployment

The next figure shows the Helion Carrier Grade VIM infrastructure:

Test Environment – Helion Carrier Grade

- Infrastructure VIM is Helion Carrier-Grade 2.0 setup in CMS R&D lab dedicated to NFV-D (G3):

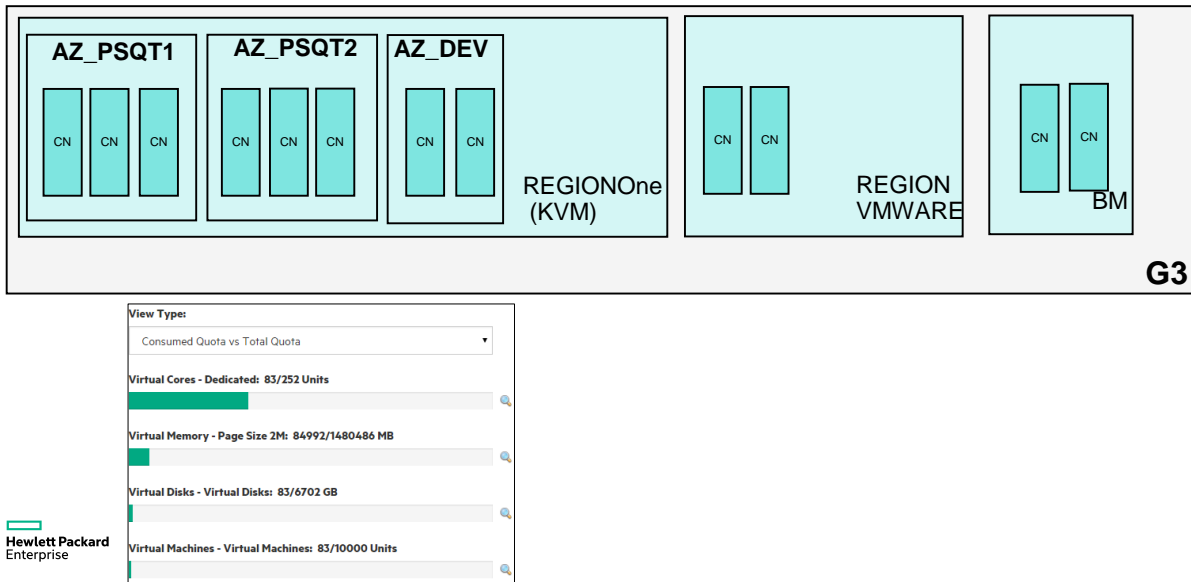


Figure 2: Helion Carrier Grade Configuration

The test environment provided only a relatively small number (8) of physical compute nodes. Thus, in order to study the influence of the number of compute nodes on the deployment time, a simulation was put in place. A large number of compute nodes were artificially “discovered” by NFV-D; however these were mapped approximately 10x1 to “real” HCG compute nodes. This allowed a characterization of the NFV-D performance in the face of a relatively large number of compute nodes, while allowing the deployments to complete within HCG.

The simulated environment, with a total of 98 compute nodes, is described in the figure below:

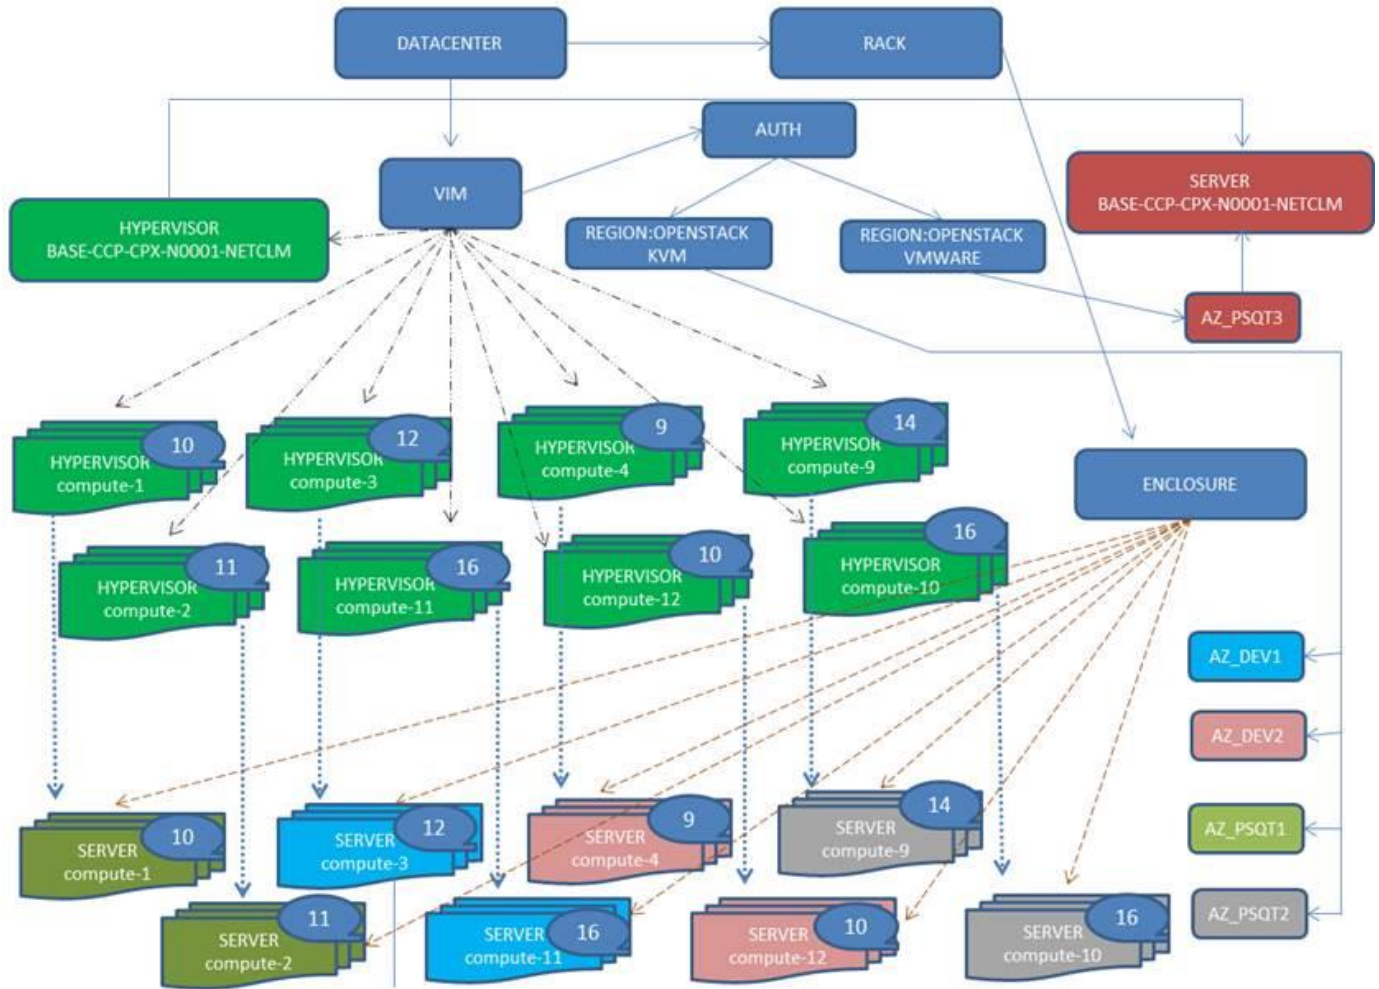


Figure 3: Simulated Compute Nodes in HCG

Chapter 3

Performance Summary

In the following tests, N identical VNFs are deployed. Each VNF consists of a single VM booting a simple Cirros image.

3.1 Monolithic Deployment of N VNFs

In this test, N VNFs are deployed sequentially. The VNFs were all deployed into a single VDC and a single VNF group.

The test used the 8 physical compute nodes in the HCG environment.

The test deployment was executed 250 times.

The time to deploy the Nth VNF was as follows:

- 1st VNF: 171 seconds
- 50th VNF: 190 seconds
- 100th VNF: 235 seconds
- 200th VNF: 356 seconds
- 250th VNF: 476 seconds

The deployment time increases only slightly up to approximately 50 VNFs, but increases rapidly after 100 VNFs. It is thus interesting to examine the behavior when the VNFs are partitioned into several VNF groups and/or VDCs.

3.2 Partitioned Deployment of N VNFs

The VNFs are typically grouped into virtual data centers (VDCs) and/or VNF Groups. The VDCs and VNF Groups can be a logical collection such as per usecase, per network domain, per customer or per Network service.

NFVD provides operators the option to organize the VNFs into VDCs VNG Groups as per their needs. It is also observed that this hierarchical distribution helps in achieving better deployment performance. The searches made by NFV-D during deployment are more efficient when the VNFs are partitioned among several VDCs or VNF Groups. This is borne out by the following test, where 200 VNFs were deployed sequentially, but split over 20 VNF groups (thus 10 VNFs per VNF group).

The test used the 8 physical compute nodes in the HCG environment.

The test deployment was executed 200 times.

The time to deploy the Nth VNF was as follows:

- 1st VNF: 164 seconds
- 50th VNF: 167 seconds
- 100th VNF: 182 seconds
- 200th VNF: 188 seconds

Thus the deployment time increases only slightly as the number of VNFs increases.

3.3 Dependency on the Number of Compute Nodes

A natural question is to study the influence of the number of compute nodes on the time taken to deploy the VNFs.

The HCG test environment consisted of a relatively small number (8) of compute nodes. Thus, in order to study the influence of the number of compute nodes on the deployment time, a simulation was put in place. A large number of compute nodes were artificially “discovered” by NFV-D; however these were mapped approximately 10x1 to “real” HCG compute nodes. This allowed a characterization of the NFV-D performance in the face of a relatively large number of compute nodes, while allowing the deployments to complete within HCG.

The test was run with a total of 98 simulated compute nodes.

The time to deploy the Nth VNF was as follows:

- 1st VNF: 206 seconds
- 50th VNF: 237 seconds
- 100th VNF: 272 seconds
- 200th VNF: 501 seconds

3.4 Parallel Deployments

In general, operations on VNFs are sequential. However, in certain circumstances, parallel operations may become necessary. For example:

- Complex VNF deployments with multiple VMs
- Disaster recovery, with increased load on secondary VNFs requiring massive scale-outs

Thus it is interesting to quantify the performance of NFV-D in the face of multiple simultaneous operations.

Note that there are two configuration parameters that can be used to control the maximum degree of parallelism allowed by NFV-D:

- In the “sosa_conf.xml” file (in /opt/OV/ServiceActivator/EP/SOSA/conf/), in the Queue com.hp.sosa.modules.sosamodule.queues.basic.BasicQueue, the parameter “queue.max.parallelism” controls the total number of requests (deployments, scales, etc.) that can be run in parallel. The default value is 10.
- The “mwfm.xml” file contains parameters controlling the maximum number of workflows to run in parallel. The parameter Max-Work-List-Length (default 512) specifies the maximum number of workflows that can run at once. The global parameter

A test was run with 10 VNF deployments in parallel, for a total of 100 VNFs, spread across 10 VNF groups.

The test used the 8 physical compute nodes in the HCG environment.

The time to deploy the Nth VNF was as follows:

- 1 – 10 VNFs: 249 seconds
- 41 - 50 VNFs: 260 seconds
- 91 – 100 VNFs 270 seconds

As in Section 3.2, the fact that the deployments are spread across several VNF Groups allows for a relatively constant deployment time as a function of the number of VNFs.

The increase in time versus sequential deployment (roughly 50% more) is due to resource contention between the workflows – a locking mechanism prohibits simultaneous access to resource allocation, in order to eliminate erroneous duplicate allocation of resources. However, this should be compared to the time to sequentially deploy the VNFs. 10 VNFs in parallel execute in 249 seconds, whereas 10 VNF sequentially would take roughly 1640 seconds.

3.5 Recommendations

Based on the above performance results, the following recommendations can be made:

- VNFs should be spread among a number of VDCs and/or VNF groups. Ideally, there should be approximately 20 – 50 VMs per group, in order to maximize both GUI performance and deployment / scaling performance.
- The SOSA max.parallelism parameter should not be increased much past its default value of 10. Higher values will result in a much higher number of maximum workflows running; the resulting overhead will tend to decrease the performance.
- The Max-Work-List-Length parameter in the file mwfm.xml should be set to at least 50 times higher than the SOSA max.parallelism parameter.

Chapter 4

Sizing Guidelines

A typical NFV-D deployment supporting High Availability and Geo-Redundancy is spread across four primary (and four secondary) VMs:

- User Interface VM
 - HPE Unified OSS Console
 - CouchDB
 - Image Upload Service
- Fulfillment VM
 - Fulfillment API
 - Service Activator (MWFM)
 - SOSA
 - Lock Manager
- Assurance VM #1
 - SiteScope
- Assurance VM #2
 - Assurance Gateway
 - Discovery
 - Unified Correlation Analyzer (UCA)
 - Neo4J Graph DB

In addition, NFV-D requires an external database server (Oracle RAC) and an LDAP server (OpenLDAP or ActiveDirectory).

The sizing of each of the VMs is dependent on specific parameters relevant to the applications running on that VM:

- User Interface VM: Number of concurrent users
- Fulfillment VM: Number of concurrent operations (deployment, scale in/out) + size of environment
- Assurance VM #1: Number, type, and frequency of SiteScope monitors;
- Assurance VM #2: Number of events/second, (Correlation); size of environment (Discovery)

4.1 Small Configuration

A typical “small” configuration (up to 100 VMs) can be summarized as follows:

- User Interface VM
 - 5 concurrent users
- Fulfillment VM
 - 5 parallel deployments (deploy/undeploy, scale in/out)
- Assurance VM #1
 - 5 monitors/VM, each monitor activated once per minute
- Assurance VM #2
 - 1 events/second
 - Discovery of 50 physical servers

The recommended sizing for such a system is as follows:

Small	VM	GUI VM	Fulfilment VM	Database VM	Assurance VM #1 (SiteScope)	Assurance VM #2 (Correlation)
	KPI	# concurrent users	# concurrent operations	# deployed VMs	# monitors	# events per second
	Range	up to 5	up to 5	up to 100	up to 500 (1 min frequency)	up to 1
cpu intel speed 2.0 GHz or more	Cores	4	4	4	4	4
	Ram (GB)	16	16	16	16	16
	Disk	256 GB	256 GB	256 GB	256 GB	256 GB

Table 2: Small NFV-D sizing

4.2 Medium Configuration

A typical “medium” configuration (up to 1000 VMs) can be summarized as follows:

- User Interface VM
 - 10 concurrent users
- Fulfillment VM
 - 10 parallel deployments (deploy/un-deploy, scale in/out)
- Assurance VM #1
 - 5 monitors/VM, frequency 1 minute per monitor (5000 monitors total)
- Assurance VM #2
 - 10 events/second
 - Discovery of 200 physical servers

The recommended sizing for such a system is as follows:

Medium	VM	GUI VM	Fulfilment VM	Database VM	Assurance VM #1 (SiteScope)	Assurance VM #2 (Correlation)
	KPI	# concurrent users	# concurrent operations	# deployed VMs	# monitors	# events per second
	Range	up to 10	up to 10	up to 1000	up to 5000 (1 min frequency)	up to 10
cpu intel speed 2.0 GHz or more	Core	8	8	8	8	8
	Ram (GB)	32	32	32	32	32
	disk	256 GB	256 GB	512 GB	512 GB	256 GB

Table 3: Medium NFV-D sizing