VMware Validated Design for Cloud Providers: Scale and Performance Guidelines

for vCloud Director 10 environments

Published December 2019

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

List of Tables

List of Figures

Introduction

The VMware Validated Designs (VVD) for Cloud Providers: Scale and Performance Guidelines (Scale and Performance) is an evolution of Certified Reference Design for VMware Cloud Providers™. It is a pre-validated set of software components that simplify the deployment of a VMware vCloud Director®-based multitenant cloud in a predictable and efficient manner. The intent of the Scale and Performance initiative is to document a verified stack and provide scale and performance benchmarking. It also helps reduce the complexity of figuring out dependencies between the VMware components required for a vCloud Director-based service. While this initiative does not yet involve software automation for software upgrades, it aims to present clearly what components are needed, which versions must be used, and what kind of scale and performance VMware Cloud Providers can expect.

VMware Cloud Providers get clarity and predictability about which version of each software component of the stack is recommended at a given time. Each Scale and Performance version also includes a predictable support timeframe for all underlying components, typically 12 – 18 months from the launch of the corresponding Scale and Performance release. This reduces the expense and time involved in determining what components to upgrade when and to which version, so that the entire software stack stays in support and incompatible combinations are avoided.

VMware Cloud Providers also benefit from clear guidelines for sizing hardware and software components to match their expected tenant load. While the Scale and Performance does not cover every cloud configuration and size, it provides a sizing recommendation for a “typical” cloud (a cloud size representative of a broad set of VMware Cloud Providers). Future versions of the Scale and Performance may address larger and less common environment configurations as well as more specialized use cases.

It is *not* the current intent of Scale and Performance to push vCloud Director to its absolute limits. For configuration maximums and limits, see VMware vCloud Director Configuration Maximums.

This document also includes the expected performance as observed by tenant users and VMware Cloud Provider administrators interacting with the vCloud Director user interface and API.

A vCloud Director-based platform can be properly sized by following the sizing guidelines for hardware and scale based on anticipated tenant demand.

1.1 Audience

This document is intended for VMware Cloud Provider architects and technical leads responsible for planning and executing the deployment and upgrades of a VMware-based cloud environment.

1.2 Scope

This document addresses the following aspects:

- Interoperability stack

  Provides a list of certified versions of all the component software comprising the software stack. Using the recommended versions guarantees known support life of the stack as well as performance characteristics.

- Sizing guidelines and software requirements
- Performance characteristics of the solution

The certified solution stack provides known performance and scale characteristics and includes recommendations and guidelines for hardware and scale based on anticipated tenant demand.

See the complimentary documents that are part of the *VMware vCloud® Architecture Toolkit™ for Service Providers*:

- *Architecting a VMware vCloud Director Solution for VMware Cloud Providers*
- *Architecting Tenant Networking with VMware NSX® in VMware vCloud Director*
- *Developing a Hyper-Converged Storage Strategy for VMware vCloud Director with VMware vSAN™*
- *VMware vCloud Director Configuration Maximums*

The VMware Product Interoperability Matrices is the authoritative resource for interoperability between the VMware software components.

A compliant solution must comply with all relevant security guidelines outlined in the product-specific documentation as well as security recommendations in the *VMware vCloud Architecture Toolkit for Service Providers* document.

Interoperability Stack (Bill of Materials)

The Bill of Materials table lists the pre-validated set of software components for Cloud Providers at the time of the Scale and Performance launch. While VMware Cloud Providers are free to choose and pick other versions or different combinations of VMware Cloud Provider Program software products, the specified stack guarantees a known predictable support time and specific performance and scaling characteristics. Performance and scaling information is provided later this document. Products marked "Core" are required to officially achieve VVD for Cloud Providers compliance.

Table 1. Bill of Materials

Component	Version and Build	Core/ Optional	Notes
VMware vCenter Server®	6.7 Update 3	Core	See Table 8 for patch level tested.
VMware ESXi™	6.7 Update 3	Core	See Table 8 for patch level tested.
VMware NSX-V	6.4.6	Core	See Table 8 for patch level tested.
VMware vSAN	6.7 Update 3	Core	See Note 2
vCloud Director	10	Core	10.0.0.1 Virtual Appliance with an embedded database.
vCloud Availability	3.5	Optional	
Container Service	2.5.1	Optional	Container Services Extension
VMware vRealize® Log Insight™	8.0	Optional	
VMware vRealize® Network Insight™	5.0	Optional	Apply the latest available patch.
VMware vRealize® Orchestrator™	7.6	Optional	Apply the latest available patch.
VMware vCloud Usage Meter	3.6.1 Hot Patch 3	Core	
VMware vRealize Operations™	8.0	Optional	Apply the latest available patch.

Note 1. These are the recommended set of products, but this is not a full interoperability matrix. For example, vCloud Director 10 is supported with multiple versions of NSX but in the current benchmarking we used a specific NSX version. Test results generally apply to all patches within

the specified major version of each component. See the VMware Product Interoperability Matrix for full vCloud Director interoperability information.

2. vSAN-based storage must be deployed in at least one cluster (either management or capacity).

Scale and Performance

The Scale Profile B table represents a common environment similar to the environments of approximately 60% of all VMware Cloud Providers. While vCloud Director is capable of a larger scale, the following profile is what is validated and benchmarked in the current Scale and Performance.

Table 2. Scale Profile B

Parameter	Value
Number of tenants (Organizations in vCloud Director)	400
Number of powered-on tenant virtual machines (with an OS installed)	10,000
Number of data centers	1
Number of vCloud Director cells	4
Number of vCenter Server instances managed by vCloud Director	1 vCenter Server for management cluster 3 vCenter Server instances for resource capacity
Number of hosts and clusters	3 resource capacity clusters, 33 hosts total 1 management cluster; 5 hosts
Maximum network latency from vCloud Director to VMware vCenter Server, VMware NSX Manager™, and ESXi hosts	Network RTT latency up to 150 ms
Concurrent API operations	Up to 128 concurrent users executing operations against the vCloud Director API
Concurrent virtual machine migrations to vCloud Director from tenant environments by vCloud Availability	100

2.1 Performance Characteristics

2.2.1 Environment Setup

The multitenant cloud environment is set up based on Scale Profile B.

Testing is performed at different levels of network latency from vCloud Director cells to vCenter Server and NSX Manager to measure the impact of network latency on performance.

2.2.2 Performance and Throughput

The test throughput is measured as the number of operations executed over 30 minutes. The test was run with different test concurrency (32, 64, and 128) and network latency (0.3 ms, 40 ms, and 150 ms). During this test, a representative random sample of operations from the List of Operations is used.

Table 3. Performance and Throughput

Concurrency (Number of concurrent users)	Throughput at RTT = 0.3 ms (Successfully completed operations per minute)	Throughput at RTT = 40 ms	Throughput at RTT = 150 ms
32	124	101	72
64	213	170	137
128	319	289	229

2.2.3 API Latency

The API Operations Latency table shows average user observed latency (in seconds) for a selection of API operations at RTT = 0.3 ms. See the

Results

Total Operations completed: 1341318

Average test throughput over 5 days: 186 Ops**/min**

Total failures over 5 days: **59**

Failure rate over 5 days: 0.004**%**

Success rate: **99.996%**

List of Operations for the full list of operations invoked during this test.

Table 4. API Operations Latency

Operation	Concurrency (seconds)

	32	64	128
Instantiate 150 MB vApp from a template	28 s	35 s	52 s
Create edge gateway	45 s	51 s	65 s
Create an independent disk	13 s	16 s	25 s

Increasing network RTT from 0.3 ms to 150 ms affects these numbers with the size of the effect varying significantly depending on the operation. With most API operations RTT increase from 0.3 ms to 150 ms caused the latency to increase by a factor of 2 or less.

2.2.4 Upload and Download Performance

The OVF upload and download times observed in the test environment vary depending on the different network latencies.

Table 5. OVF Upload and Download Times

	RTT = 0.3 ms	RTT = 40 ms	RTT= 150 ms
OVF upload time in seconds (4 GB)	359	388	392
OVF download time in seconds (4 GB)	232	235	245

2.2.5 vCloud Availability

The Time to Protect a virtual machine in vCloud Availability represents the time to establish replication of virtual machines of various sizes between a vCloud Director and a disaster recovery target vCenter Server environment using vCloud Availability. 10 GB uplinks were configured between vCenter Server and vCloud Director. Network throughput was stable around 710 Mbps.

Table 6. Time to Protect a VM

VM size	Time to Protect
1 GB	29 sec
10 GB	2 min
100 GB	27 min
500 GB	1 hr: 58 min
1 TB	3 hr: 55 min

The Network Latency Impact on Migration Performance table displays how the network latency between vCenter Server and vCloud Director impacts cold migration for virtual machines of 100 GB size.

Table 7. Network Latency Impact on Time to Protect (VM size = 100 GB)

RTT Latency between vCenter Server and vCloud Director	Time to Protect
0.3 ms	27 min
40 ms	29 min
150 ms	33 min

Sizing Guidelines

Many environment variables influence the number of hosts and CPU and memory resources required to run a cloud service based on VVD for Cloud Providers. It is impossible to give a precise formula for how much of each resource is required. The current effort focuses on demonstrating how our deployment setup behaves in terms of scale and performance under the defined test load.

The Scale Profile B captures the parameters of the load on the system in terms of number of tenants, organizations, VMs, network latency, and cloud management operations load. The results in terms of average response time, throughput, and uptime under these controlled conditions provide a starting point for Cloud Providers to estimate how much capacity is needed for their use cases. We strongly recommend that Cloud Providers extensively test each environment prior to production use to ensure that the performance meets the business SLAs.

3.1 vCloud Director Virtual Appliance

The current VVD for Cloud Providers recommends the use of a vCloud Director 10 virtual appliance. The virtual appliance includes an embedded and fully managed PostgreSQL database and built-in replication for maintaining consistency between cells. A properly configured virtual appliance cluster is also resilient to failures of individual cells.

The recommended vCloud Director virtual appliance deployment includes 1 primary cell, 2 stand-by cells and 0 or more application cells. See the vCloud Director Installation, Configuration, and Upgrade guide for more details about the deployment, configuration, and operation of the virtual appliance.

For the purposes of VVD scale and performance benchmarking, the following setup was used: 1 primary cell, 2 standby cells, 1 application cell. The Table 8 Management and Resource Component Sizing gives the details of memory and CPU capacity on each vCloud Director appliance node.

The embedded PostgreSQL database is automatically configured. However, some post-deployment tuning might be needed for best performance. See the section **PostgreSQL Tuning** for the database parameters used in this benchmarking.

3.2 Management and Resource Component Sizing

The following table summarizes sizing choices made for various management and resource components.

Table 8. Management and Resource Component Sizing

Component	Version	Size	Resources	Notes
Management vCenter Server (vCenter Server Appliance with an embedded Platform Services Controller)	6.7 Update 3	Tiny	RAM: 10 GB CPU: 2 Storage: 250 GB	1 management vCenter Server
Resource vCenter Server (vCenter Server Appliance with an embedded Platform Services Controller)	6.7 Update 3	Medium	RAM: 24 GB CPU: 8 Storage: 400 GB	3 resource vCenter Server instances
ESXi	6.7 Update 3		Cisco UCSC-C240-M5SX servers	33 hosts for resource cluster, 5 hosts for management cluster
vSAN (deployed in management cluster)	6.7 Update 3			
NSX-V for vSphere	6.4.6		RAM: 16 GB CPU: 4 Storage: 60 GB	

Component	Version	Size	Resources	Notes
vCloud Director	10.0.0.1	Virtual appliance	**1 Primary, 2 Standby cells** RAM: 32 GB CPU: 24 Storage: 170 GB + 500 GB NFS **1 application cell** RAM: 8 GB CPU; 8 Storage: 170 GB	170 GB = 120 GB for database and 50 GB for logs per cell and 500 GB of shared NFS storage for the vCloud Director transfer service.
vCloud Availability	3.5			
vCloud Availability C4 Appliance			RAM: 4 GB CPU: 2 Storage: 10 GB	
vCloud Availability Replicator			RAM: 6 GB CPU: 4 Storage: 10 GB	
vCloud Availability Tunnel Appliance			RAM: 2 GB CPU: 2 Storage: 10 GB	
vCloud Director Database	PostgreSQL 10.5		RAM: 32 GB CPU: 16 Storage: 300 GB	Database is embedded with the vCloud Director virtual appliance
Container Service Extension	2.5			
vCloud Director AMQP	RabbitMQ 3.7.9			
vCloud Director Metrics Database	Cassandra 3.11.3			

Component	Version	Size	Resources	Notes
vRealize Log Insight deployment	8.0	Medium	RAM: 16 GB CPU: 8 Storage: 500 GB	Use the vRealize Log Insight sizing calculator: http://www.vmware.com/go/loginsight/calculator
vRealize Network Insight deployment	5.0	Large		
vCloud Usage Meter	3.6.1 HP3	Standard	RAM: 4 GB CPU: 2 Storage: 100 GB	Use the deployment requirements for vCloud Usage Meter 3.6: https://www.vmware.com/support/vcloud-usage-meter/doc/vcloud-usage-meter-36-interop.html
vRealize Orchestrator	7.6		RAM: 6 GB CPU: 2 Storage: 20 GB	vRealize Orchestrator 7.6 was used as opposed to 8.0 (in the BOM) because of an issue affecting vRealize Orchestrator 8.0.
vRealize Orchestrator plug-in for vCloud Director	10			
vRealize Operations Manager	8.0	Medium	RAM: 32 GB x 3 CPU: 8 x 3 Storage: 300 GB x 3	1 Master, 1 Master Replica, 1 Data node Use the vRealize Operations sizing guidelines: https://kb.vmware.com/s/article/75162

Component	Version	Size	Resources	Notes
Management Pack for NSX for vSphere	3.5.2			
Management Pack for vSphere	8.0			
Management Pack for vRealize Log Insight	7.5			
Management Pack for vCloud Director	5.1			
Management Pack for vSAN	8.0			
vRealize Operations Tenant App for vCloud Director	2.3			

Appendix A – Test Environment and Benchmarking Methods

Test Environment

The test environment is broadly divided into three main setups:

- Management cluster
- Resource cluster (30% of workloads on vSAN, 70% on iSCSI storage)
- Test driver

Management Cluster

This is where all the management components were deployed.

- Management components
  - 1 x Management vCenter Server (Tiny)
  - 4 x vCloud Director virtual appliance cells (1 primary, 2 stand-by, 1 application)
  - 3 x Resource vCenter Server (Medium)
  - 3 x NSX Manager
  - 1 x Management NSX vCloud Director edge
  - 1 x vRealize Log Insight (Medium)
  - 1 x vRealize Operations (Large)
  - 1 x vCloud Usage Meter (Standard)
- Management Cluster Resources
  - 5 physical servers with 192 GB RAM and 28 cores, each with vSAN supported SSDs
  - 10 TB vSAN, 10 TB iSCSI

Figure 1. Management Component Deployment

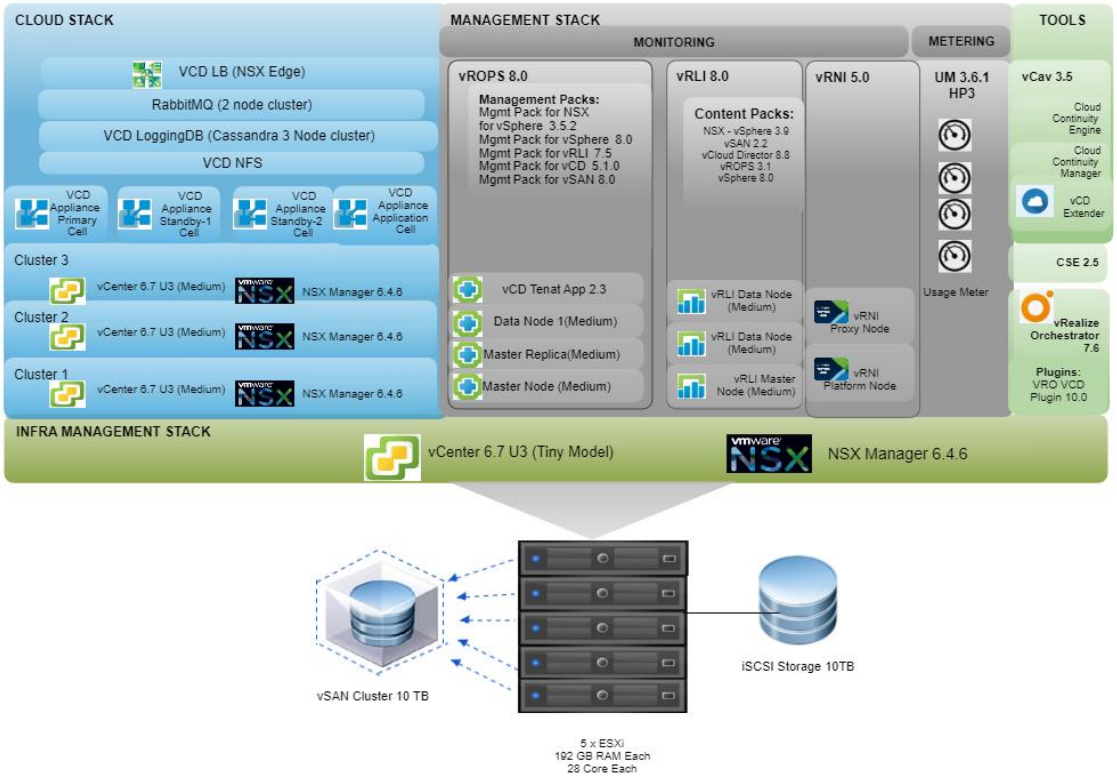

Figure 2. Management Cluster Networking

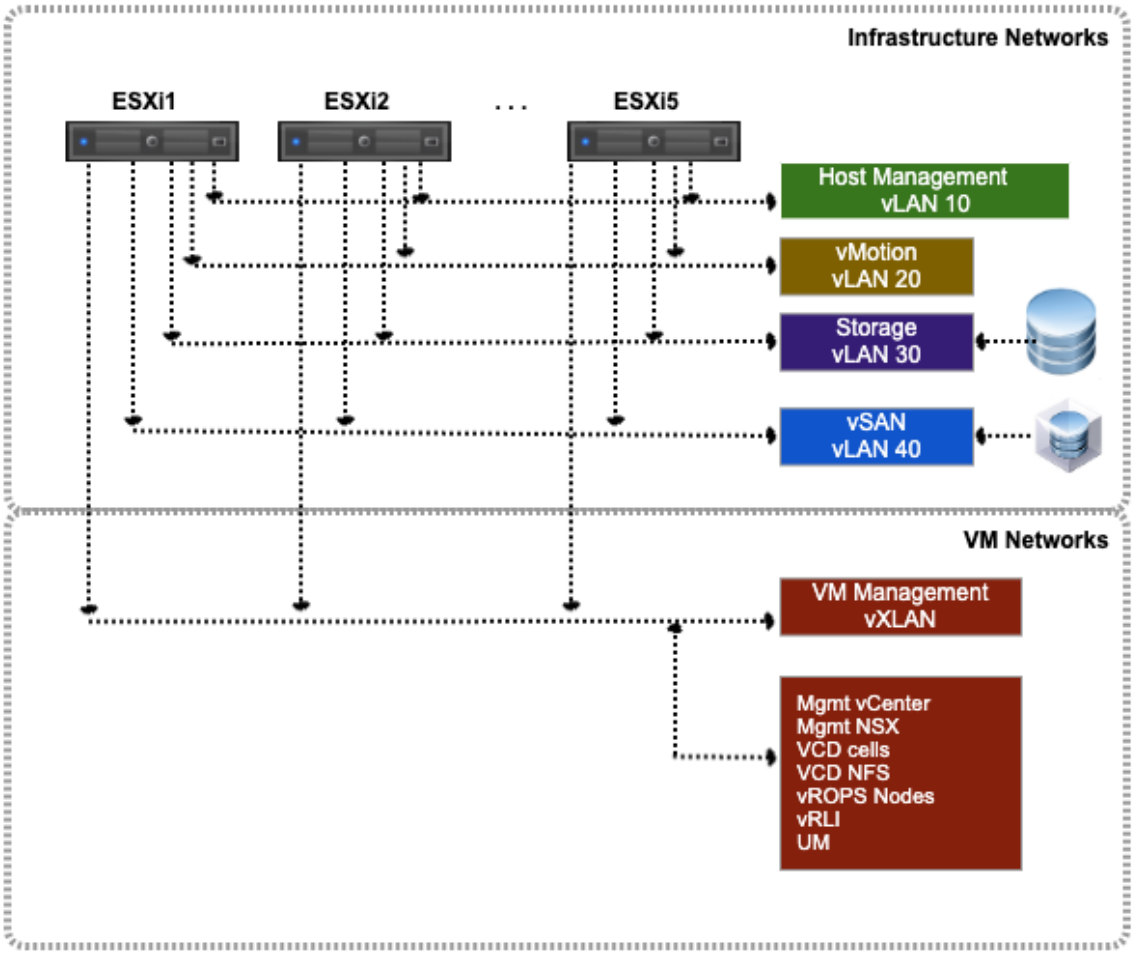

Resource Cluster

This is where Tenant Organizations and workload virtual machines were created.

- Resource cluster resources
  - 33 physical servers with 192 GB RAM and 28 cores, each with vSAN supported SSDs
  - 30 TB vSAN, 10 TB iSCSI

Figure 3. Resource Cluster Setup

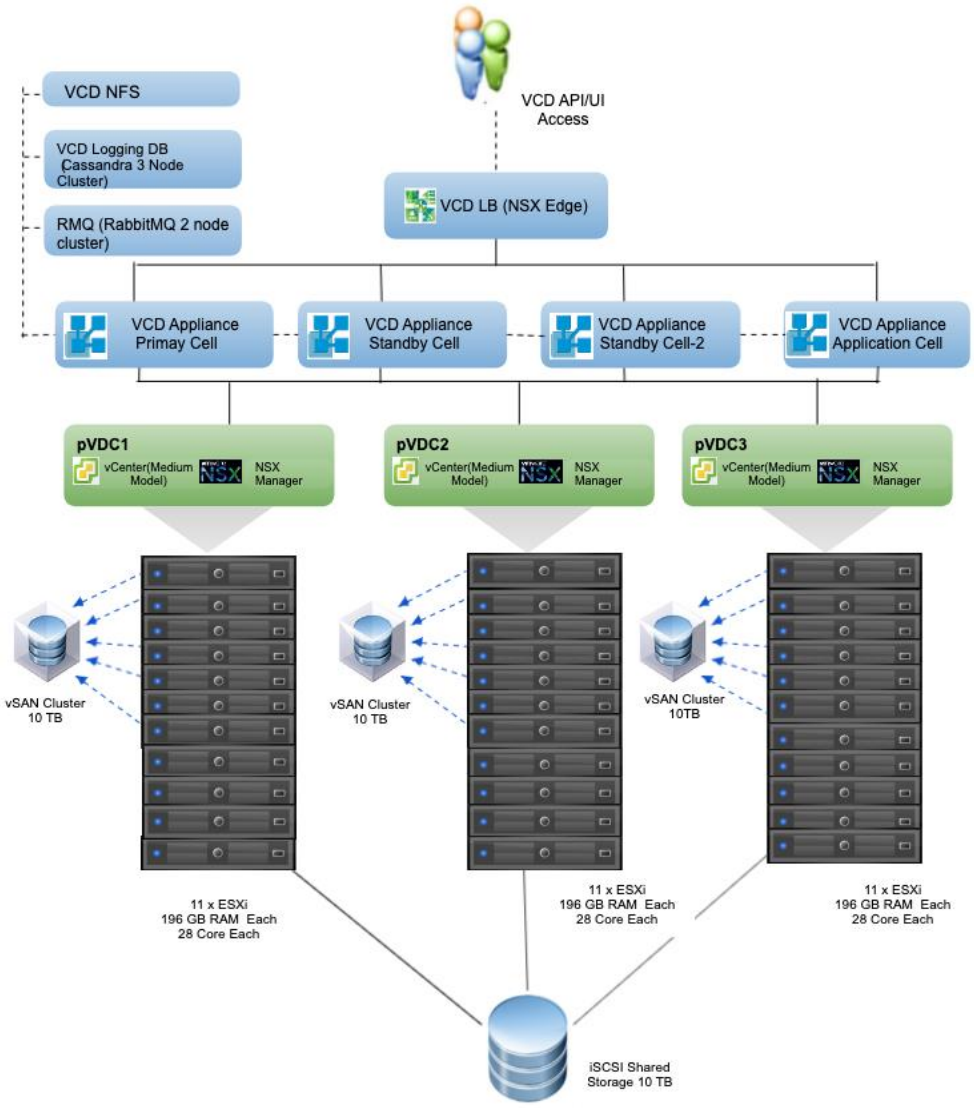

Figure 4. Resource Cluster Networking

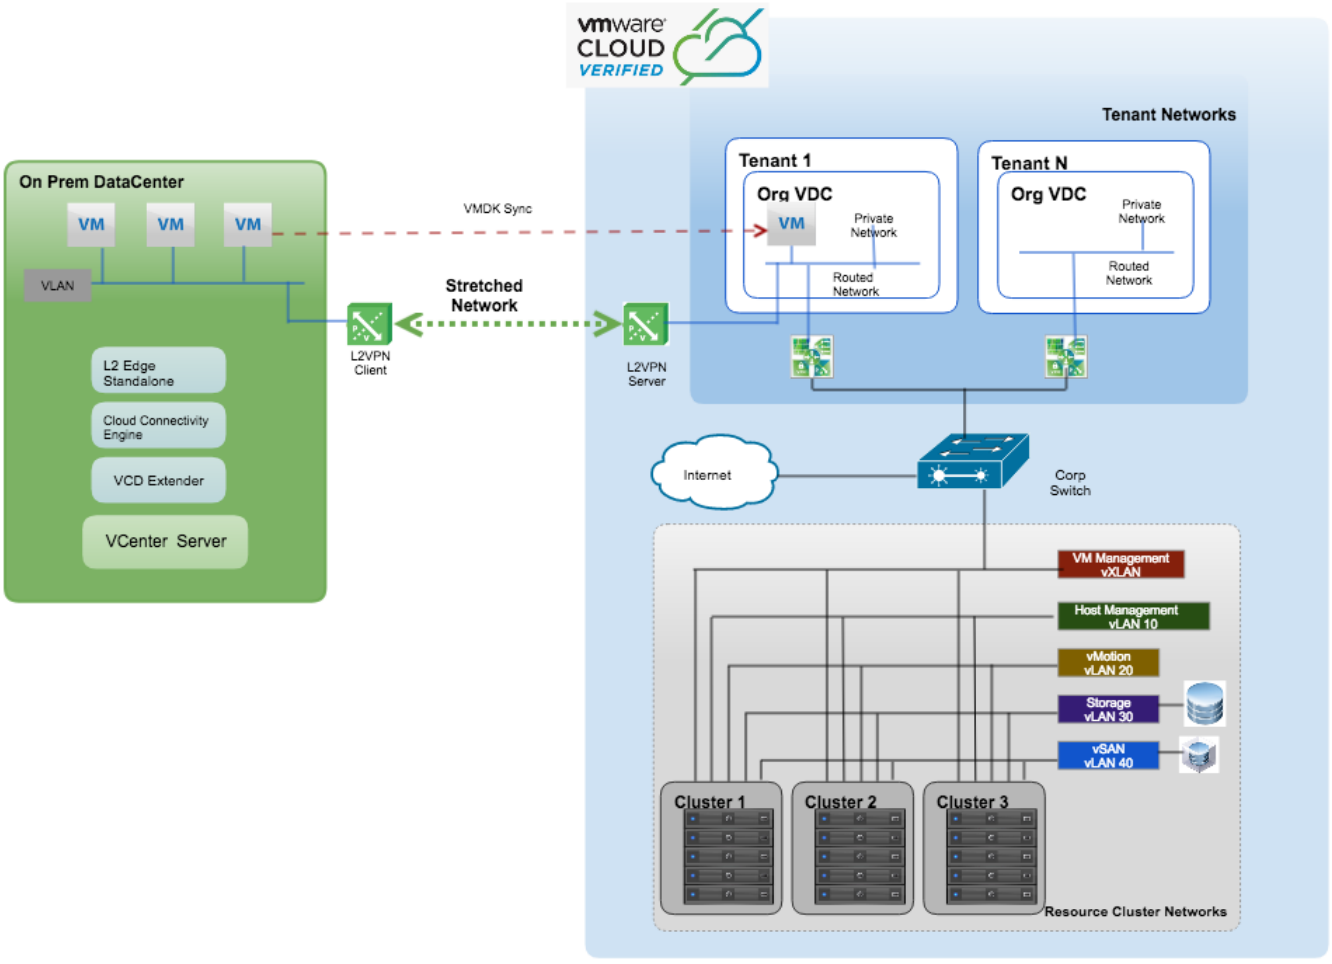

PostgreSQL Tuning

PostgreSQL database parameters were set as follows:

shared_buffers = 8GB
effective_cache_size = 24GB
work_mem = 8MB
maintenance_work_mem = 512MB
max_parallel_workers_per_gather = 0
max_worker_processes = 24

See How to Modify PostgreSQL Configuration from the vCloud Director documentation.

Test Driver

The test driver suite is executed from this environment.

- 4 CPU, 8 GB memory, CentOS 7.3

Benchmarking Methods

The testing process is focused primarily on verifying and measuring environment behavior for:

- Scale – Verify whether the environment meets the Scale Profile B requirement of 10,000 powered-on virtual machines.
- Performance – Measure operation latency and throughput when the environment is running at scale (10,000 powered-on virtual machines).
- Uptime – Verify that the environment can operate at scale with reasonable performance for a long time.

The remainder of this section details the exact methods used for test execution and measurement.

Scale Test

Scale was carried out with a mix of manual operations and JMeter test tool-based script operations by using the following steps:

1. Create 400 Tenant Organizations in vCloud Director.
2. Create and power on 10,000 virtual machines across these 400 Tenant Organizations. All virtual machines were running Ubuntu OS with 2 GB disk, 1 GB memory
3. A sample of vCloud Director operations were carried out to verify that system behaves normally at this scale.

Performance Test

Performance tests were done by executing a well-known distribution of vCloud Director operations with the help of an internal test tool. For the complete operation list, see List of Operations.

The following were the key steps in execution and measurement of the operations:

1. Scaled up the environment as outlined in the previous section.
2. After the environment was at scale, executed continuous stream of operations for 30 minutes with following distribution:
    - 35-40% vApp operations such as instantiate, deploy, edit, clone, and delete.

- 25% storage-centric operations such as create, attach, detach, and delete disk.
- 15% networking-related operations, such as create and delete gateway, routed networks and firewall configurations.
- 5% create and delete Orgs, users, catalogs, and virtual data centers.

3. Operations were executed using vCloud Director local users of different roles (vApp Author, Org Admin, System Admin) with 10% admin roles and 90% user operation roles.
4. Given that most of the operations are asynchronous, the test tool monitors the task returned by vCloud Director to get completion status and execution time details.
5. Steps 2 to 4 were repeated with 32, 64, and 128 concurrent users to ascertain the ability of the system to deal with concurrent operation invocation.
6. Step 5 was repeated for following latency (between vCloud Director and vCenter Server) values (achieved by artificial latency injection with a tool):
   - 0.3 ms (default)
   - 40 ms
   - 150 ms

Uptime Tests

Uptime tests involved running the environment (based on Scale Profile B) for 5 days and executing a constant stream of API calls covering a representative set of operations. The purpose of the test is to establish the API call success rate and system uptime.

1. Tests ran continuously for 5 days.
2. API workflows were triggered by 100 concurrent clients, each client would invoke an operation roughly every 20 seconds. 10,000 powered on VMs
3. No artificial latency injection was done.

Results

Total Operations completed**: 1341318**

Average test throughput over 5 days: **186 Ops/min**

Total failures over 5 days: **59**

Failure rate over 5 days: **0.004%**

Success rate: **99.996%**

List of Operations

For performance benchmarking, API test clients executed a predetermined distribution across different types of vCloud Director operations as described in the following tables.

Table 9. vCloud Director Operations (Part 1)

vApp Operations	Network Operations	Management Operations

Instantiate vApp	Deploy a fenced vApp	Create an org
Deploy (power on)	Undeploy a fenced vApp	Create a user
Edit vApp	Create an isolated network	Create an Org VDC
Compose vApp	Delete an isolated network	Create a direct VDC network
Clone vApp	Create a gateway	Create a catalog
Power off vApp	Create a routed Org network	Delete a catalog
Delete vApp	Instantiate a vApp in that network	Delete a VDC network
	Deploy a vApp	Delete an Org VDC
	Undeploy a vApp	Delete a user
	Delete a vApp	Delete an Org
	Delete a routed Org network	
	Delete a gateway	

Table 10. vCloud Director Operations (Part 2)

NSX Management Operations	Datastore Operations	OVF Operations
Convert edge to Advanced edge	Create a disk	OVF upload
Edge routing services	Instantiate a vApp	OVF download
Edge firewall services	Attach a disk to a vApp	
Edge NAT services	Detach a disk from a vApp	
Distributed firewall services	Delete a disk	
Load balancer services	Delete a vApp	

Appendix B – FAQ

How frequently will the Scale and Performance be updated?

- We expect to release an updated Scale and Performance with every major vCloud Director release.

How is this document related to the VMware interoperability matrix?

- The benchmarked stack is a subset of the full interoperability matrix and reflects the exact components we validated and benchmarked in this exercise. The full interoperability includes many more products and versions than what is tested in this exercise

How is the Scale and Performance related to VMware Cloud Foundation?

- VMware Cloud Foundation is not currently part of the benchmarking, however the stack we are testing is closely aligned with the current VCF BOM.

Is Scale and Performance suitable for greenfield environments or brownfield environments?

- Any environment can be made compliant by simply upgrading all its components to versions listed in the Scale and Performance Bill of Materials. There is no other qualification.

How can we provide input/recommendations for future versions of this doc?

- Contact the vCloud Director team at vcd-feedback@vmware.com or reach out to your VMware account team and pass your feedback through them.

What is the support model for an environment configured according to these guidelines?

- Each component of the Cloud Provider Platform stack is supported according to its support lifecycle. A cloud deployment compliant with the Bill of Materials is in support for at least 12 months after the Scale and Performance release date.