VMware vCloud® Architecture Toolkit™
for Service Providers

Certified Reference Design for VMware Cloud Providers™

Version 2.5

August 2018

© 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at http://www.vmware.com/download/patents.html.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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

Contents

List of Tables

List of Figures

vmware
CLOUD PROVIDER PROGRAM

Introduction

The Certified Reference Design (CRD) for VMware Cloud Providers™ 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 CRD initiative is to reduce the complexity of deploying, upgrading, and managing 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 clearly present what components are needed, which versions should be used, and what kind of scale and performance VMware Cloud Providers can expect from a CRD-compliant cloud.

The CRD gives VMware Cloud Providers clarity and predictability about which version of each software component of the stack is recommended at a given time. The CRD also comes with a predictable support time frame for all underlying components, typically 12 – 18 months from the launch of the corresponding CRD 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 CRD 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 CRD may address larger and less common environment configurations as well as more specialized use cases.

It is *not* the current intent of CRD to push vCloud Director to its absolute limits. For configuration maximums and limits, see the 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.

CRD-compliant solutions 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 of Certified Reference Design:

- Interop stack

  Provides a list of certified versions of all of 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 Reference Design-based solution provides known performance and scale characteristics and comes with recommendations and guidelines for hardware and scale based on anticipated tenant demand.

Consult 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 VMware software components.

A CRD-compliant stack 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 CRD 2.5 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 CRD compliance.

Table 1. Bill of Materials

Component	Version and Build	Core/ Optional	Notes
VMware vCenter Server®	6.5	Core	See Table 8 for patch level tested.
VMware ESXi™	6.5	Core	See Table 8 for patch level tested.
VMware NSX	6.4	Core	
VMware vSAN	6.6.1	Core	See Note 2
vCloud Director	9.1	Core	1. Databases tested: • Main Database: both PostgreSQL 10.3 and Microsoft SQL Server 2016 • Metric Database: Cassandra 3.9 2. Apply the latest available patch of vCloud Director 9.1.
vCloud Director Extender	1.1	Optional	
VMware vRealize® Log Insight™	4.6	Optional	
VMware vRealize® Network Insight™	3.7	Optional	Apply the latest available patch.
VMware vCloud Usage Meter	3.6	Core	Apply the latest available patch.
VMware vRealize Operations™	6.7	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 9.1 is supported with multiple versions of NSX but the CRD recommends using a specific NSX version. 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).

2.1 Support

Each component of the CRD 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 CRD release date.

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 CRD.

Table 2. Scale Profile B

Parameter	Value
Number of tenants (Organizations in vCloud Director)	400
Number of tenant workloads (mix of powered-on and powered-off virtual machines)	15,000
Number of powered-on tenant virtual machines	9,000
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 1 management cluster
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 VM migrations to vCloud Director from tenant environments by vCloud Director Extender	100

2.2 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	111	90	83
64	186	164	143
128	274	254	208

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

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

Table 4. API Operations Latency

Operation	Concurrency		
	32	**64**	**128**
Instantiate 150 MB vApp from a template	23 s	34 s	35 s
Create edge gateway	70 s	60 s	81 s
Create independent disk	12 s	20 s	21 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 vast majority of API operations RTT increase from 0.3 ms to 150 ms caused the latency to increase by a factor of 2 or less.

NOTE: The automation tool used for performance testing in CRD 2.5 is Apache JMeter. Comparing the API latency between CRD 2.5 and earlier CRD versions is not recommended.

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)	252	255	289
OVF download time in seconds (4 GB)	123	138	125

2.2.5 vCloud Director Extender

The VM Migration to vCloud Director table represents the time for cold migration and synchronization of virtual machines of various sizes from a vCenter Server system to a vCloud Director environment by using vCloud Director Extender. 10 GB uplinks were configured between vCenter Server and vCloud Director. Network throughput was stable around 710 Mbps.

Table 6. VM Migration to vCloud Director

VM size	Cold Migration	Sync Duration
1 GB	2 min	45 s
10 GB	5 min	3 min 30 s
100 GB	33 min	31 min
500 GB	1 hr 41 min	1hr 40 min
1 TB	5 hr	3 hr 19 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 Migration Performance

RTT Latency between vCenter Server and vCloud Director	Sync Duration	Total Migration Time
0.3 ms	20 minsuites	22 min 19 s
40 ms	20 min 15 s	22 min 50 s

150 ms	20 min 15 s	23 min 43 s

Sizing Guidelines

Use these guidelines to size the hardware and software components to match your anticipated tenant load. These guidelines apply to environments with size which is roughly similar to or smaller than Scale Profile B.

Sizing of the environment is done in two steps:

1. By determining the number of vCenter Server instances, ESXi hosts, and vCloud Director cells to use for handling at least 15,000 VMs with 9000 powered-on VMs.
2. By making specific configurations of resources (memory, CPU, disk) for each component in the stack.

3.1 Tenant Resource Capacity Clusters

Number of vCenter Server instances	The vCenter Server sizing guide suggests using a medium profile of vCenter Server to support 4000 VMs. Number of vCenter Server instances = number of VMs/4000 = 15,000/4000 = 4 (rounded)
Number of ESXi hosts	ESXi host count is determined based on the number of powered-on VMs, using the formula below. The formula provides a rough estimate and the actual number of hosts required depends on many parameters and the type of workloads. Number of hosts = (number of powered on VMs * Avg # of vCPUs per VM) / (sockets*cores*hyper threading*vCPU-to-pCPU ratio) = (9000*1)/(2*8*2*16) = 17 CRD used 21 hosts, with a set of 6 managed by each vCenter Server.
Number of vCloud Director cells	As with the vCloud Director design guide, the number of vCloud Director cells for this setup is calculated using following formula: Number of vCloud Director cells = (Number of VMs/4000) + 1 = (15,000/4000) + 1 = 5 (rounded)

3.2 Management and Resource Component Sizing

The following table summarizes sizing choices made for various management 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)	6.5 Update 2b	Tiny	RAM: 10 GB CPU: 2 Storage: 250 GB	1 management vCenter Server

Component	Version	Size	Resources	Notes
Resource vCenter Server (vCenter Server Appliance with an embedded Platform Services Controller)	6.5 Update 2b	Medium	RAM: 24 GB CPU: 8 Storage: 400 GB	3 resource vCenter Server instances
ESXi	6.5, patch ESXi650-201806001			
vSAN (deployed in management cluster)	6.6.1			
NSX for vSphere	6.4		RAM: 16 GB CPU: 4 Storage: 60 GB	
vCloud Director	9.1		RAM: 16GB CPU: 4 Storage: 60 GB	Allocate 60 GB of storage for each vCloud Director cell and 500 GB of shared NSX storage for the vCloud Director transfer service.
vCloud Director DB	Microsoft SQL Server 2016 Enterprise		RAM: 32GB CPU: 16 Storage: 500 GB	
vRealize Log Insight deployment	4.6	Medium	RAM: 16 GB CPU: 8 Storage: 1 TB	Use the vRealize Log Insight sizing calculator: http://www.vmware.com/go/loginsight/calculator
vCloud Usage Meter	3.6.0.1	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

Component	Version	Size	Resources	Notes
vRealize Operations Manager	6.7	Large	RAM: 48 GB CPU: 16 Storage: 400 GB	Use the vRealize Operations sizing guidelines: https://kb.vmware.com/s/article/54370
Management Pack for NSX for vSphere	3.5.1			
Management Pack for vSphere	6			
Management Pack for vRealize Log Insight	6			
Management Pack for vCloud Director	4.5			
Management Pack for Storage Devices	6.0.5			
Management Pack for vSAN	2.0			
vCloud Director Extender	1.1			

vmware
CLOUD PROVIDER PROGRAM

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 cells
  - 1 x Microsoft SQL Server 2016 Enterprise
  - 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)
  - 1 x vCloud Director Extender
- Management Cluster Resources
  - 6 physical servers with 192 GB RAM and 28 cores, each with vSAN supported SSDs
  - 10 TB vSAN, 2 TB iSCSI

vmware® CLOUD PROVIDER PROGRAM

1. Management Component Deployment

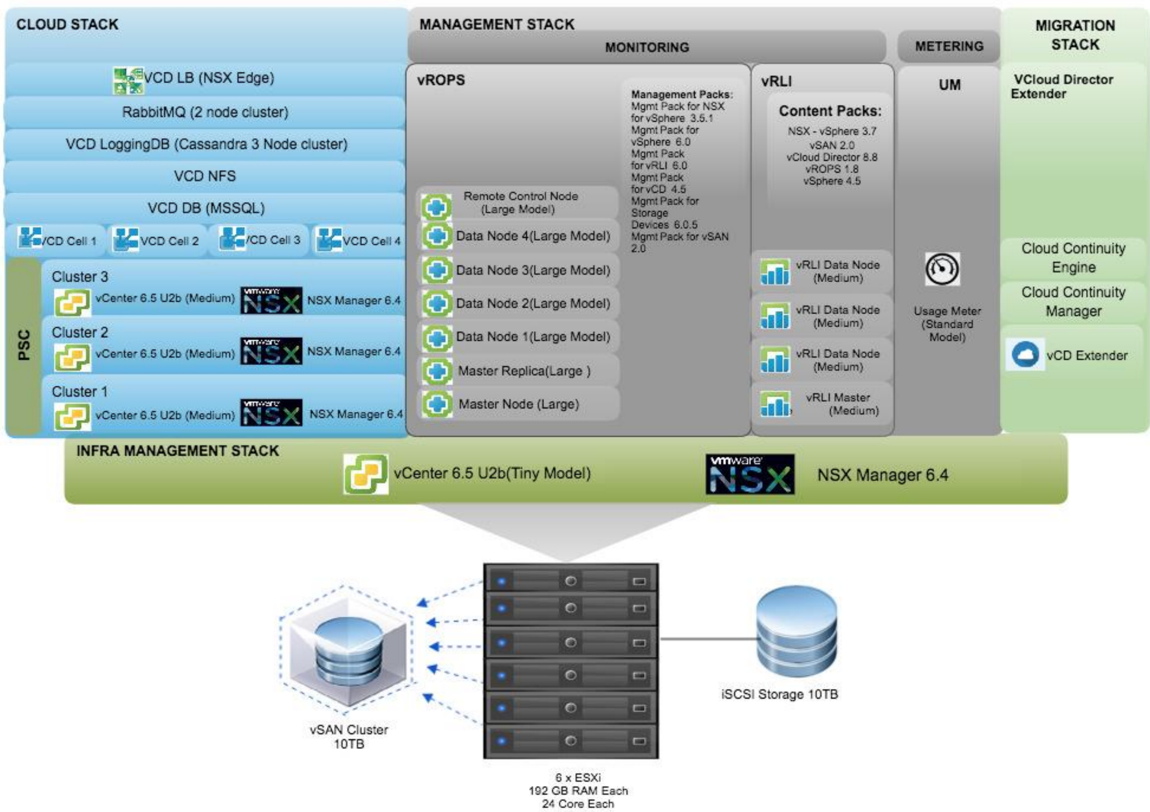

Figure 2. Management Cluster Networking

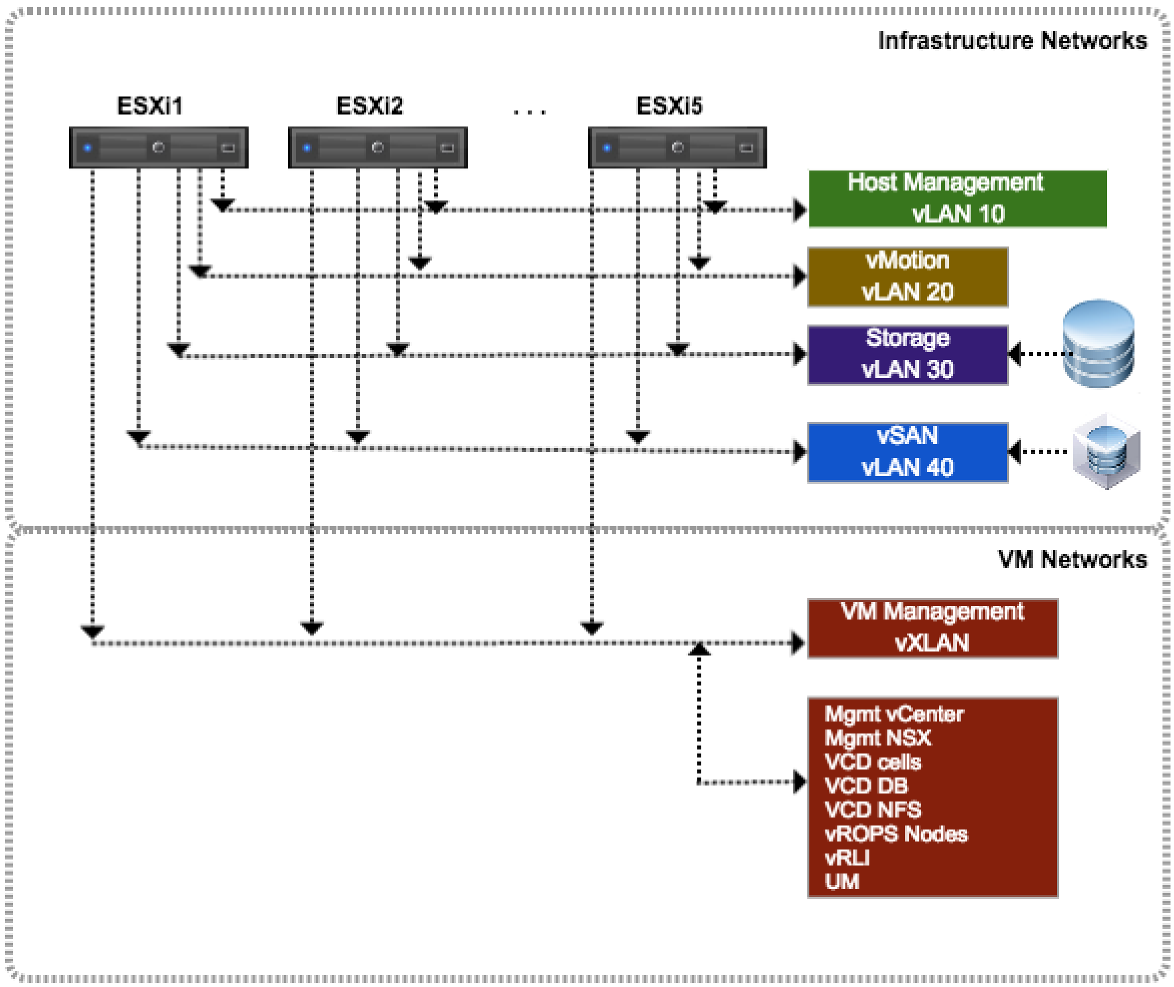

Resource Cluster

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

- Resource cluster resources
  - 21 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

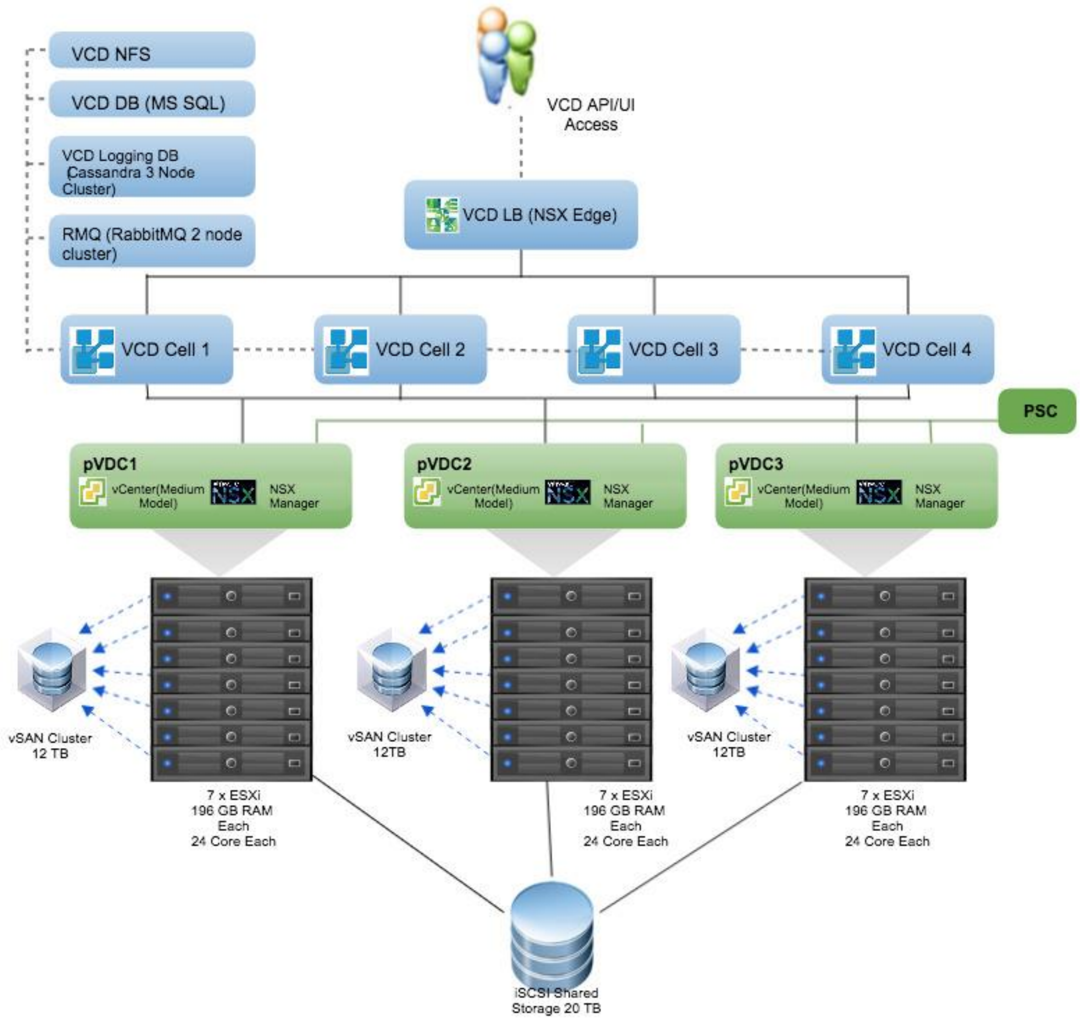

vmware
CLOUD PROVIDER PROGRAM

Figure 4. Resource Cluster Networking

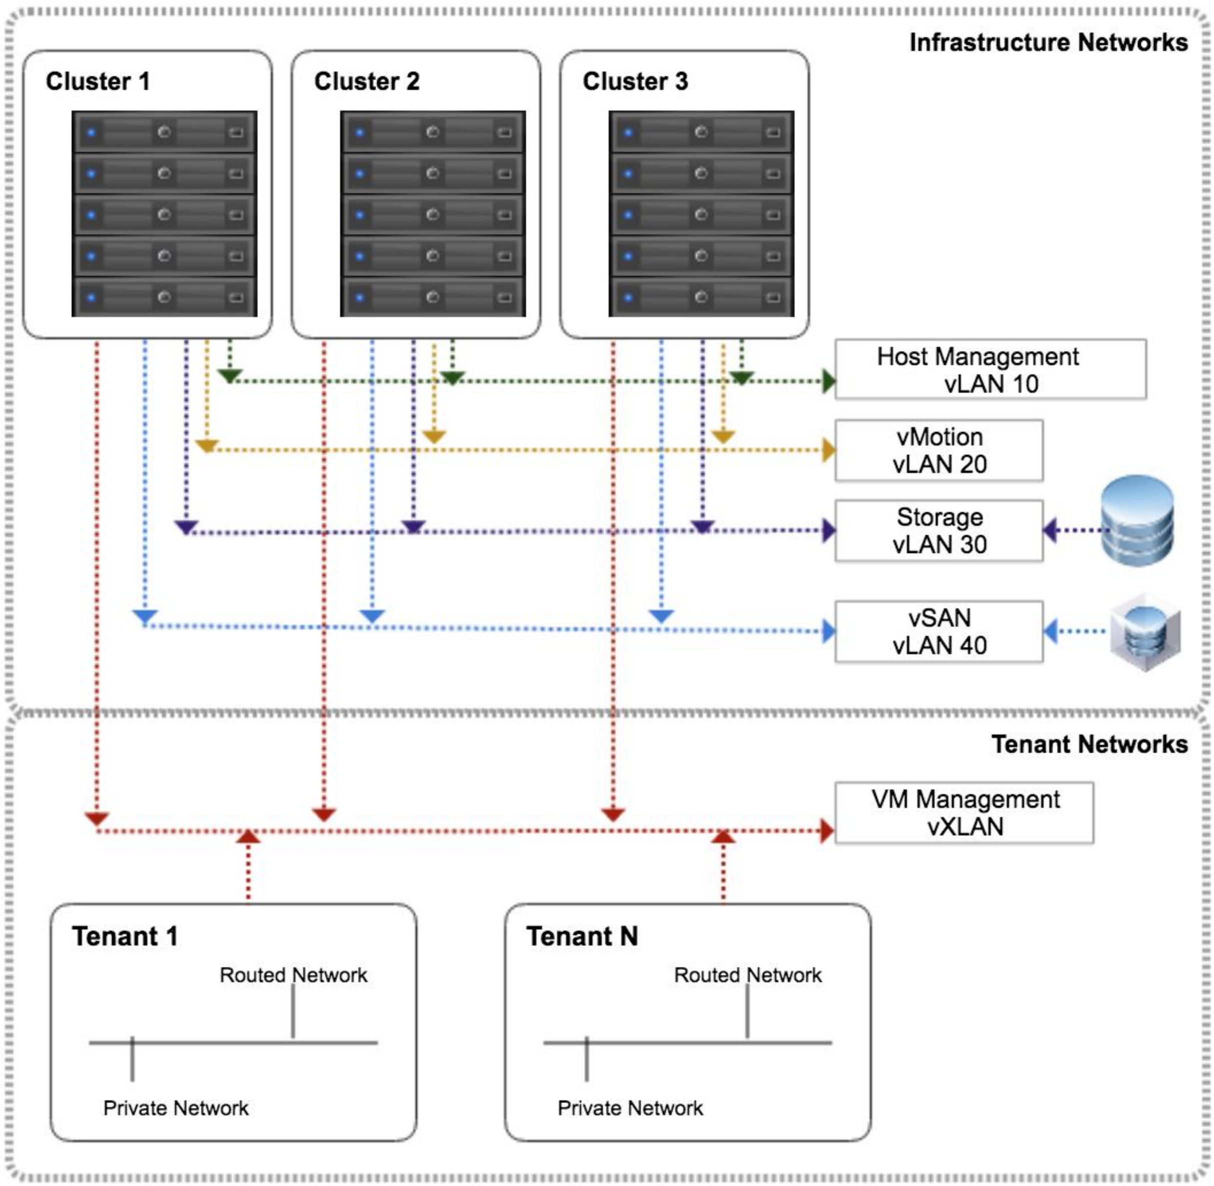

Test Driver

The test suite is executed from this environment.

- 4 CPU, 8 GB memory, Cent OS 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 15,000 virtual machines and 9000 powered-on virtual machines.
- Performance – Measure operation latency and throughput when the environment is running at scale (15,000 virtual machines with 9000 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 15,000 virtual machines across these 400 Tenant Organizations.
    - 90% of the virtual machines were Dummy Small Tiny virtual machines with 4 MB disk, 4 MB memory
    - 10% actual virtual machines
        - 7% Linux virtual machines
        - 3% Windows virtual machines
3. Power on 9000 virtual machines with similar distribution (90%:10%) of dummy to actual virtual machines.
4. Some simple 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.
5. 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.
6. 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.
7. 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 executing operations carried out during performance tests, with following changes to exaction duration and concurrency:

1. Tests ran continuously for 7 days.
2. Of 100 concurrent users each invoked an operation at every 20 seconds.
3. No artificial latency injection was done.

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

	Delete routed Org network	
	Delete gateway	

Table 10. vCloud Director Operations (Part 2)

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

Appendix B – FAQ

How frequently will the CRD be updated?

- We expect to release an updated CRD with every major vCloud Director release.

How is this document related to the VMware interoperability matrix?

- The recommended CRD stack is a subset of the full interoperability matrix and reflects the exact components we validated and benchmarked in this exercise.

How is the CRD related to VMware Cloud Foundation?

- VMware Cloud Foundation is not considered as part of this CRD.

Is CRD suitable for greenfield environments or brownfield environments?

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

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

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

Is there CRD-specific support?

- No. Each component of the CRD is supported through its existing support arrangement.