



VMware vCloud® Architecture Toolkit™  
for Service Providers

# Certified Reference Design for VMware Cloud Providers™

Version 1.0  
October 2017





© 2017 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](http://www.vmware.com)



Contents

Introduction ..... 5

    1.1 Audience ..... 5

    1.2 Scope ..... 5

Interoperability Stack (Bill of Materials) ..... 6

    2.1 Support..... 7

    2.2 Performance Characteristics ..... 7

Sizing Guideline ..... 9

    3.1 Tenant Resource Capacity Clusters ..... 9

    3.2 Management and Resource Component Sizing ..... 9

Appendix A – Test Environment and Benchmarking Methods..... 12

Appendix B – FAQ ..... 19

List of Tables

Table 1. Bill of Materials ..... 6

Table 2. Scale Profile A..... 7

Table 3. Performance and Throughput..... 8

Table 4. API Operations Latency ..... 8

Table 5. OVF Upload and Download Times ..... 8

Table 6. Management and Resource Component Sizing ..... 9

Table 7. vCloud Director Operations (Part 1) ..... 18

Table 8. vCloud Director Operations (Part 2) ..... 18

List of Figures

Figure 1. Management Component Deployment..... 13

Figure 2. Management Cluster Networking ..... 14

Figure 3. Resource Cluster Setup ..... 15

Figure 4. Resource Cluster Networking..... 16





## 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 we do not cover every cloud configuration and size, we strive to provide 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 CRD’s intent to push vCloud Director to its absolute limits. For configuration maximums and limits, consult the *Configuration Maximums* document for vCloud Director.

Finally, we are documenting 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 tech 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 version 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™*
- *Configuration Maximums for vCloud Director*

The [VMware Product Interoperability Matrix](#) is the authoritative resource for interoperability between VMware software components.

## Interoperability Stack (Bill of Materials)

The following table lists the pre-validated set of software components recommended for Cloud Providers at the time of the CRD 1.0 launch. While VMware Cloud Providers are free to choose and pick other versions or different combinations of VMware Cloud Provider Program software products, the recommended stack guarantees a known predictable support time and specific performance and scaling characteristics. Performance and scaling information is provided later this document.

**Table 1. Bill of Materials**

Component	Version and Build	Notes
VMware vCenter Server®	6.5	See Table 6 for patch level tested.
VMware ESXi™	6.5	See Table 6 for patch level tested.
VMware NSX	6.3.3	
vCloud Director	8.20	1. Database: SQL Server 2014. 2. We recommend the latest available patch of 8.20. At the time of this document – 8.20.0.2.
VMware vRealize® Log Insight™	4.5	
VMware vCloud Usage Meter	3.6	
VMware vRealize Operations™	6.5 or 6.6	We recommend the latest available 6.6 patch release.

**Note** These are the recommended set of products, but this is not a full interoperability matrix. For example, vCloud Director 8.20 is supported with multiple versions of NSX but we are recommending a specific NSX version here. Consult the [VMware Product Interoperability Matrix](#) for full vCloud Director interoperability information.



## 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 will stay in support for at least 12 months after the CRD release date.

The following profile represents a common environment similar to the environments of approximately 60 percent of all VMware Cloud Providers. While vCloud Director is capable of a larger scale, the following profile is what we validated and benchmarked in the current CRD.

**Table 2. Scale Profile A**

Parameter	Value
Number of tenants (Organizations in vCloud Director)	200
Number of tenant workloads (mix of powered-on and off VMs)	10,000
Number of powered-on tenant VMs	7000
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
Max network latency from vCloud Director to VMware vCenter Server / VMware NSX Manager™	Network RTT latency up to 100 ms

## 2.2 Performance Characteristics

### 2.2.1 Environment Setup

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

Latency from vCloud Director cells to vCenter Server and NSX Manager components was set to 40 ms and 100 ms.

### 2.2.2 Performance and Throughput

We measured the test throughput 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 100 ms). We used a representative random sample of operations from the *List of Operations* during this test.



**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 = 100 ms
32	61	55	50
64	113	106	97
128	185	171	161

### 2.2.3 API Latency

The following 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 template	20 sec	22 sec	33 sec
Create edge gateway	55 sec	77 sec	85 sec
Create independent disk	12 sec	14 sec	23 sec

Increasing network RTT from 0.3 ms to 100 ms affects these numbers with the size of the effect varying significantly depending on the operation. The highest impact was observed with deploying a vApp where an RTT of 100 ms caused the latency to increase by a factor of 2.

### 2.2.4 Upload/Download Performance

The following table shows OVF upload and download times observed in our test environment under different network latencies.

**Table 5. OVF Upload and Download Times**

	RTT = 40 ms	RTT= 100 ms
OVF upload time in seconds (4GB OVF)	530	543
OVF download time (4GB OVF)	136	147





## Sizing Guideline

This section includes guidelines to help you size hardware and software components to match your anticipated tenant load. These guidelines apply to environments whose size is roughly similar to Scale Profile A (that is, within a factor of 2 of profile A).

Sizing of the environment was done in two steps:

1. We determined the number of vCenter Server instances, ESXi hosts, and vCloud Director cells to use to handle at least 10,000 VMs with 7000 powered-on VMs.
2. For each component in the stack, specific configurations of resources (memory, CPU, disk) were made.

### 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 =  
10,000/4000 = 3 (rounded)

#### Number of ESXi hosts

ESXi host count was determined based on the number of powered-on VMs, using the following formula:

Number of hosts = number of powered on VMs /  
(sockets\*cores\*hyperthreading\*processors) = 7000/(2\*8\*2\*16) = 14

We used 15 hosts, with a set of 5 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 arrived at using following formula:

Number of vCloud Director cells = (Number of VMs/4000) + 1 =  
(10,000/4000) + 1 = 2.5 + 1 = 4

### 3.2 Management and Resource Component Sizing

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

**Table 6. Management and Resource Component Sizing**

Component	Version	Size	Resources	Notes
Management vCenter Server  (vCenter virtual appliance with embedded DB and VMware Platform Services Controller™)	6.5 EP1 (6.5 a)	Tiny	RAM: 10 GB  CPU: 2  Storage: 400 GB	



Component	Version	Size	Resources	Notes
Resource vCenter (vCenter virtual appliance with embedded DB and Platform Services Controller)	6.5 EP1 (6.5 a)	Medium	RAM: 24 GB CPU: 8 Storage: 400 GB	
ESXi	6.5 EP1 (6.5 a)			
NSX	6.3.1		RAM: 16 GB CPU: 4 Storage: 60 GB	
vCloud Director	8.20.0.1		RAM: 16GB CPU: 4 Storage: 300GB + 500GB NFS	
vCloud Director DB	MSSQL 2014 Enterprise		RAM: 32GB CPU: 16 Storage: 1 TB	
vRealize Log Insight deployment	4.5	Medium	RAM: 16 GB CPU: 8 Storage: 1 TB	For a vRealize Log Insight sizing calculator, see <a href="http://www.vmware.com/go/oginsight/calculator">http://www.vmware.com/go/oginsight/calculator</a>
Content Pack for NSX	3.6			
Content Pack for vSAN	2			
Content Pack for vCloud Director	8.8			
Content Pack for vRealize Operations	1.8			
Content Pack for VMware vSphere®	4.5			



Component	Version	Size	Resources	Notes
vCloud Usage Meter	3.6.0.1	Standard	RAM: 4 GB CPU: 2 Storage: 100 GB	The deployment requirements for vCloud Usage Meter 3.6 are published here: <a href="https://www.vmware.com/support/vcloud-usage-meter/doc/vcloud-usage-meter-36-interop.html">https://www.vmware.com/support/vcloud-usage-meter/doc/vcloud-usage-meter-36-interop.html</a>
vRealize Operations	6.5		RAM: 48 GB CPU: 16 Storage: 400 GB	vRealize Operations 6.5 sizing guidelines are published in a VMware Knowledge Base article: <a href="http://kb.vmware.com/kb/2148829">http://kb.vmware.com/kb/2148829</a>
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			



## Appendix A – Test Environment and Benchmarking Methods

### Test Environment

The test environment is broadly divided into three main setups:

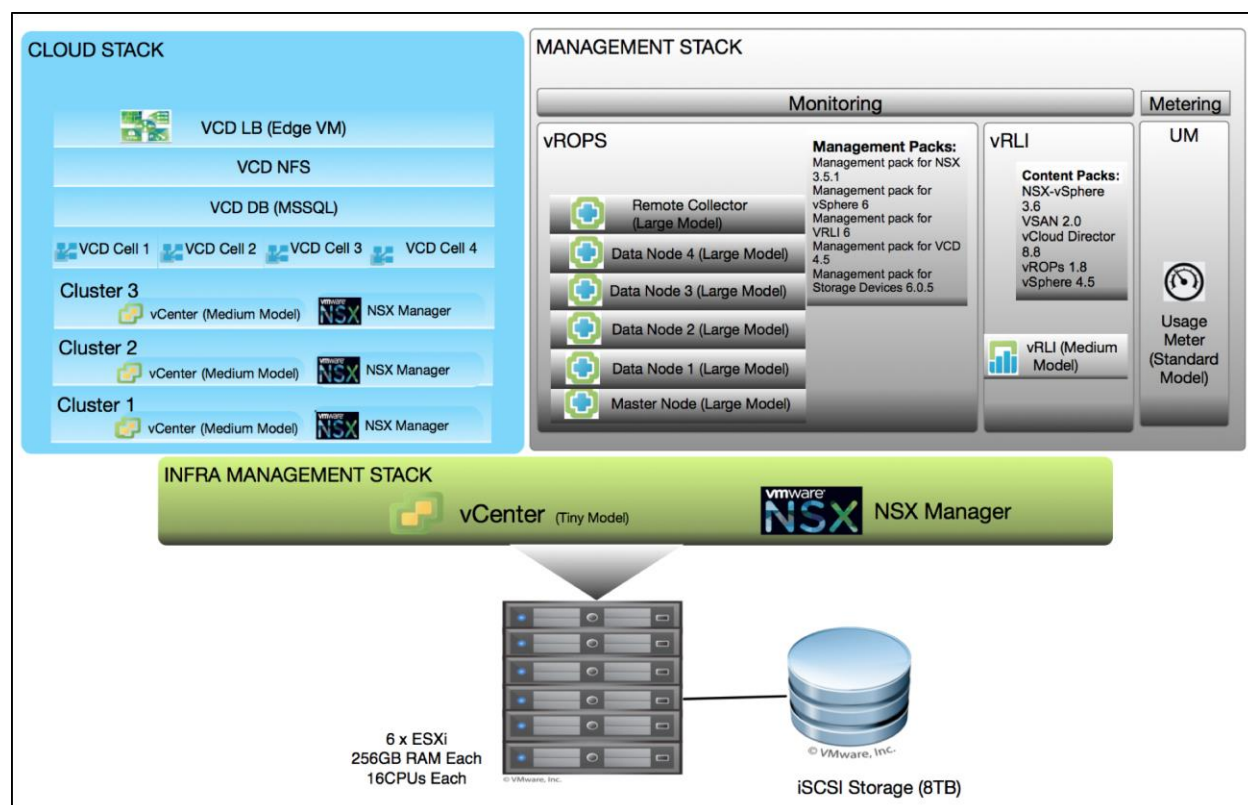
- Management cluster
- Resource cluster
- Test driver

### Management Cluster

This is where all the management components were deployed.

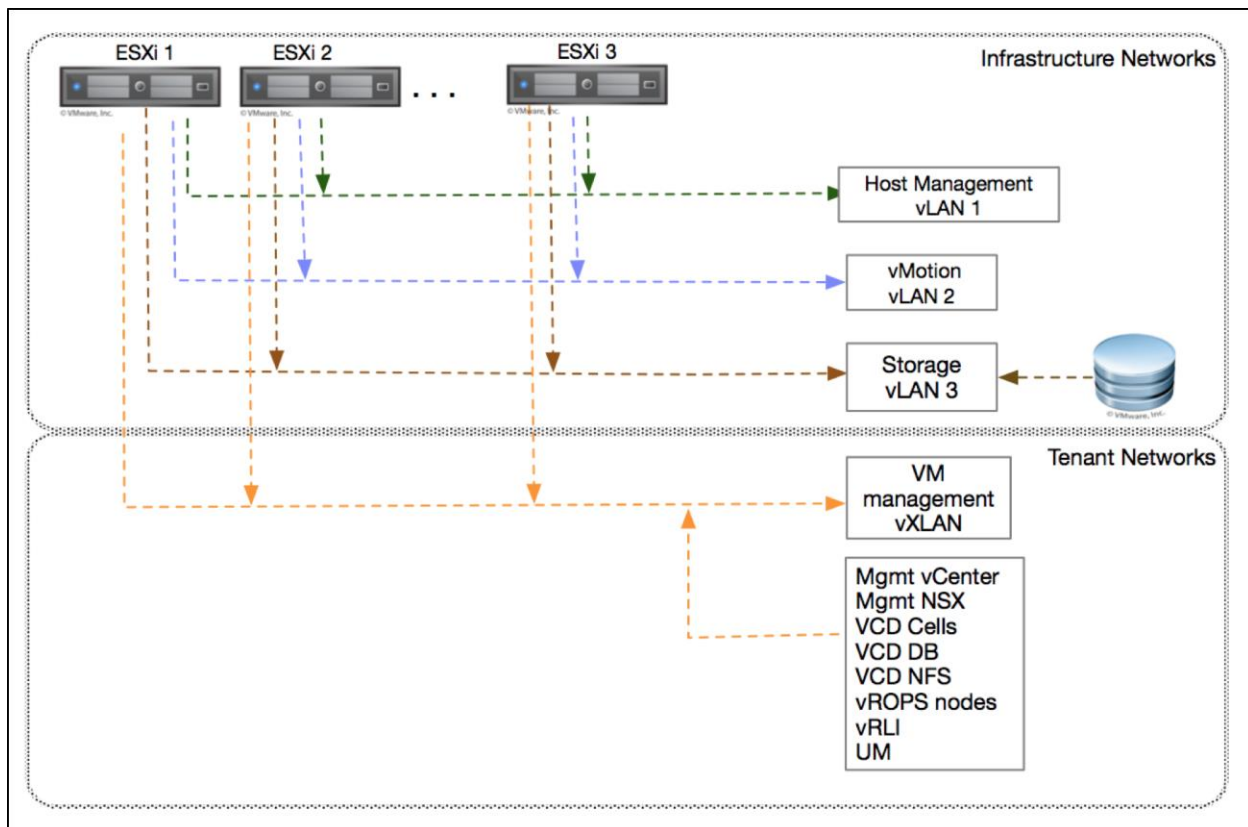
- Management components
  - 1 x Management vCenter Server (Tiny)
  - 4 x vCloud Director cells
  - 1 x MSSQL 2014 Enterprise
  - 3 x Resource vCenter Server (Medium)
  - 3 x NSX Manager
  - 1 x Management NSX vCloud Director edge
  - 1 x vRealize Log Insight (Medium)
  - 1 x vRealize Operations (Large)
  - 1 x vCloud Usage Meter (Standard)
- Management Cluster Resources
  - 6 x Supermicro SYS-F627R3-RTB+02-VI009
    - 16 CPUs 2.6 GHz
    - 256 GB memory
    - 8 TB iSCSI LUNs from VNX array

**Figure 1. Management Component Deployment**





**Figure 2. Management Cluster Networking**



## Resource Cluster

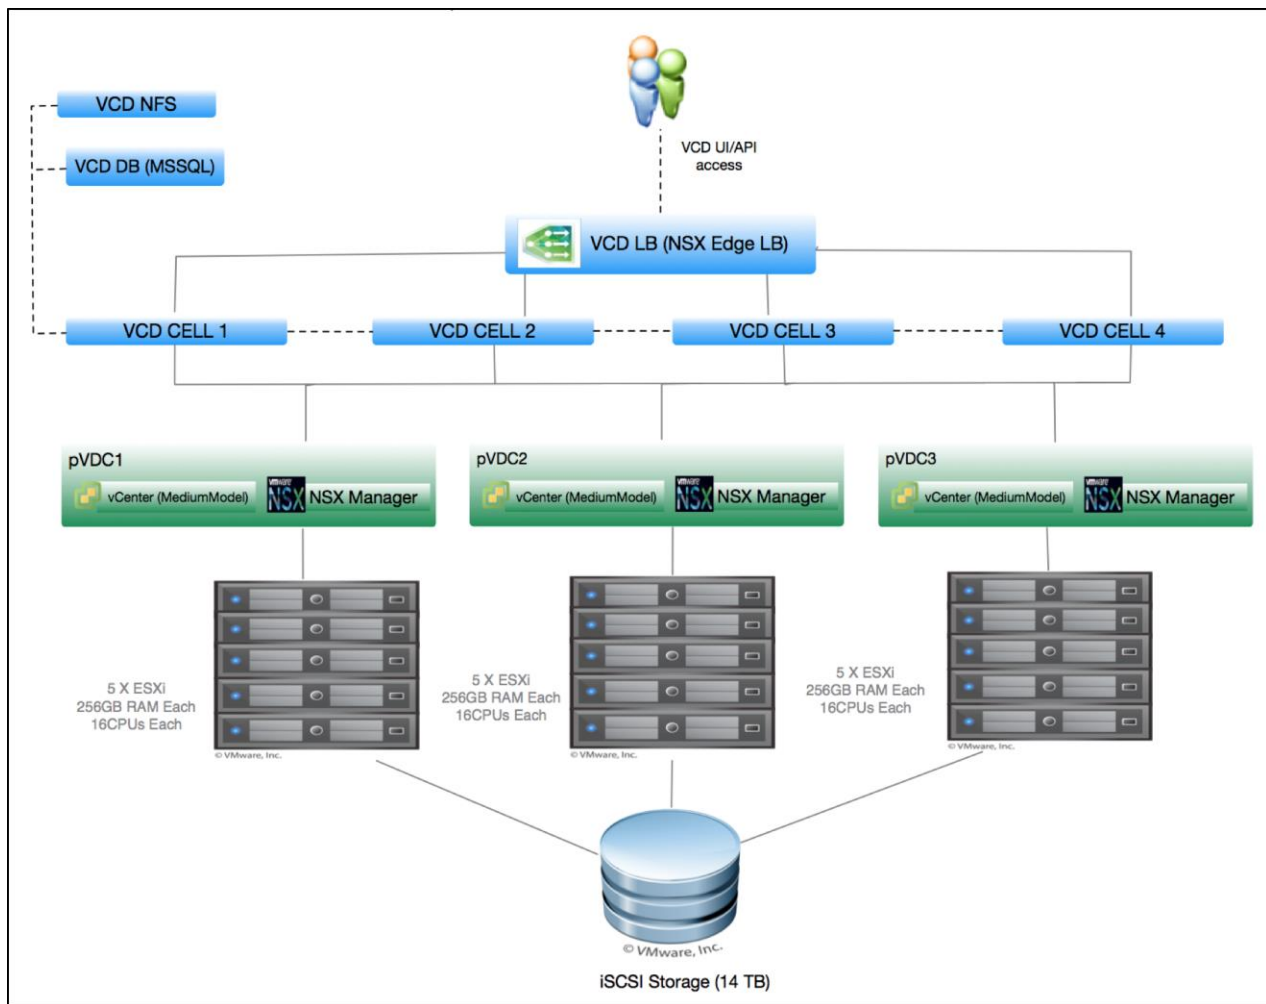
This is where Tenant Organizations and workload VMs were created.

- Resource cluster resources
  - 15 x Supermicro SYS-F627R3-RTB+02-VI009 (5 hosts for each of the 3 vCenter Server instances)
    - 16 CPUs 2.6 GHz
    - 256 GB memory
    - 14 TB iSCSI LUNs from VNX array



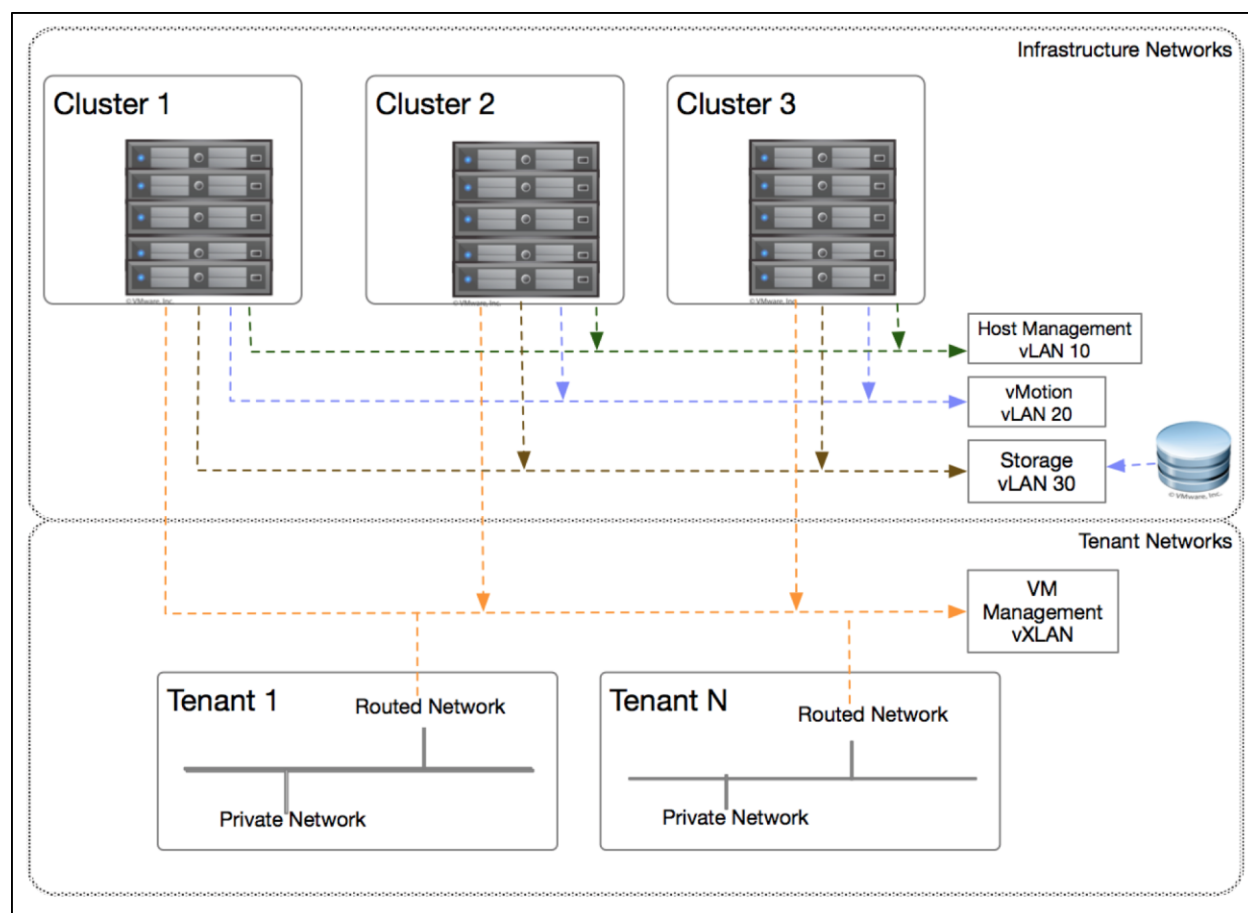


**Figure 3. Resource Cluster Setup**





**Figure 4. Resource Cluster Networking**



## Test Driver

The test suite is executed from this environment.

- 4 CPU, 8 GB memory, Windows 2014

## Benchmarking Methods

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

- **Scale** – Verify whether the environment meets the Scale Profile A requirement of 10,000 VMs and 7000 powered-on VMs.
- **Performance** – Measure operation latency and throughput when the environment is running at scale (10,000 VMs with 7000 powered-on VMs).
- **Uptime** – Verify that the environment can operate at scale with reasonable performance for a long duration.

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 to using the following steps:

1. Create 200 Tenant Organizations in vCloud Director.



2. Create 10,000 VMs across these 200 Tenant Organizations.
  - 90% of the VMs were Dummy Small Tiny VMs with 4 MB disk, 4 MB memory
  - 10% actual VMs
    - 7% Linux VMs
    - 3% Windows VMs
3. Power on 7000 VMs with similar distribution (90%:10%) of dummy to actual VMs.
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 internal test tool. (A complete operation list is included in the last section of this appendix.)

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/delete gateway, routed networks and firewall configurations.
  - 5% create/delete Orgs, users, catalog, and virtual data centers.
3. Operations were executed using vCloud Director local users of different roles (vApp Author, Org Admin, System Admin) with 10 percent admin roles and 90 percent user operations.
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
  - 100 ms

### Uptime Tests

Uptime tests involved executing operations carried out during performance tests, with following changes to execution duration and concurrency:

1. Tests ran continuously for 7 days.
2. 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 7. 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 8. 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 VCF?

- VCF 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](mailto: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.