

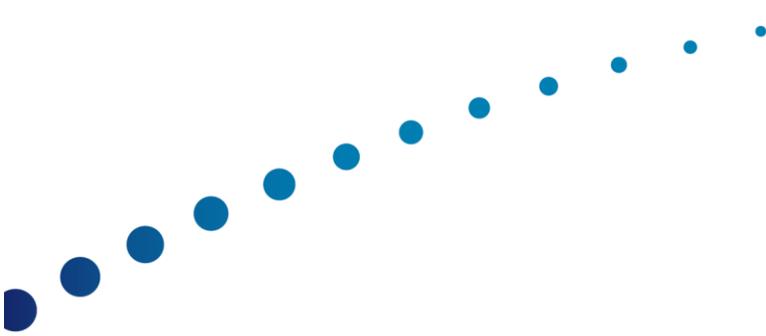


# **Citrix XenDesktop 7.1 on Microsoft Hyper-V Server® 2012 High-Level Design**

Citrix Validated Solutions

17<sup>th</sup> December 2013

Prepared by:  
Citrix Consulting



## Revision History

Revision	Change Description	Updated By	Date
1.0	Document created and updated.	APAC Citrix Consulting	17-Dec-2013

## Table of Contents

<b>1.</b>	<b>Executive Summary .....</b>	<b>5</b>
1.1	Audience .....	5
1.2	Purpose .....	5
1.3	Reference Architecture .....	5
<b>2.</b>	<b>Architecture Overview.....</b>	<b>6</b>
2.1	Citrix Virtual Desktop Types .....	6
2.2	The Pod Concept.....	6
2.3	Justification and Validation .....	7
2.4	High Level Solution Overview.....	8
2.5	Assumptions .....	10
<b>3.</b>	<b>Logical Architecture.....</b>	<b>11</b>
3.1	Logical Component Overview for HSD.....	11
3.2	Logical Component Overview for HVD.....	13
<b>4.</b>	<b>Physical Architecture.....</b>	<b>15</b>
4.1	Physical Component Overview.....	15
4.2	Physical Component Design HSD.....	18
	BoM Cisco UCS Compute Hardware for HSD.....	18
	BoM Cisco Nexus Network Hardware for HSD .....	19
	BoM Nimble Storage Array Hardware for HSD .....	19
4.3	Hardware Support and Maintenance for HSD .....	19
	Cisco UCS and Nexus for HSD .....	19
	Nimble Storage for HSD .....	20
4.4	Physical Component Design HVD.....	21
	BoM Cisco UCS Compute Hardware for HVD.....	21
	BoM Cisco Nexus Network Hardware for HVD .....	22
	BoM Nimble Storage Array Hardware for HVD .....	22
4.5	Hardware Support and Maintenance for HVD .....	22
	Cisco UCS and Nexus for HVD .....	22
	Nimble Storage for HVD .....	23
<b>5.</b>	<b>High-Level Design .....</b>	<b>24</b>
5.1	Network / Cisco Nexus .....	24
	Overview .....	24
	Key Decisions .....	25
	Design.....	25
5.2	Cisco UCS .....	27
	Overview .....	27
	Key Decisions .....	27
	Design.....	28
5.3	Nimble Storage.....	30
	Overview .....	30
	Key Decisions .....	31
	Design.....	33
5.4	Microsoft Hyper-V Server 2012 .....	38
	Overview .....	38
	Key Decisions .....	39
	Design.....	40
5.5	SMB File Services .....	43
	Overview .....	43
	Key Decisions .....	44
	Design.....	45
5.6	Citrix Provisioning Services .....	48
	Overview .....	48
	Key Decisions: PVS.....	49
	Key Decisions: DHCP .....	50
	Design.....	51
5.7	Citrix XenDesktop .....	52
	Overview .....	52

Key Decisions .....	53
Design .....	54
5.8 Virtual Desktop VM Guest Workloads .....	55
Overview .....	55
Key Decisions .....	55
Design .....	60
5.9 Citrix StoreFront.....	61
Overview .....	61
Key Decisions .....	61
Design .....	61
5.10 Citrix License Server .....	62
Overview .....	62
Key Decisions .....	62
Design .....	62
5.11 Citrix NetScaler SDX .....	63
Overview .....	63
Key Decisions .....	63
Design .....	64
5.12 User Profile Management Solution.....	65
Overview .....	65
Key Decisions .....	65
Design .....	65
5.13 Active Directory .....	66
Overview .....	66
Key Decisions .....	66
Design .....	66
5.14 Database Platform .....	68
Overview .....	68
Key Decisions .....	68
Design Considerations.....	68
<b>Appendix A. Decision Points .....</b>	<b>69</b>
<b>Appendix B. Server Inventory .....</b>	<b>71</b>

## 1. Executive Summary

### 1.1 Audience

This reference architecture document is created as part of a Citrix Validated solution (CVS) and is intended to describe the detailed architecture and configuration of the components contained within. Readers of this document should be familiar with Citrix XenDesktop, its related technologies and the foundational components; Cisco UCS, Cisco Nexus, Nimble Storage and Microsoft Hyper-V Server® 2012.

### 1.2 Purpose

The purpose of this document is to provide high-level design information that describes the architecture for this Citrix Validated Solution which is based on Citrix Hosted Shared Desktop (HSD) and Citrix Hosted Virtual Desktop (HVD) FlexCast models. The solution is built on Cisco UCS compute, Cisco Nexus switching and Nimble Storage array.

### 1.3 Reference Architecture

In order to facilitate rapid and successful deployments of the Citrix XenDesktop [FlexCast](#) models described in this document, Citrix Consulting APAC have procured, built and tested a solution built on Cisco UCS, Cisco Nexus and Nimble Storage hardware. The Citrix Validated Solution provides prescriptive guidance on Citrix, Cisco and Nimble Storage design, configuration and deployment settings thereby allowing customers to quickly deliver virtual desktop workloads.

Extensive testing was performed using Login VSI to simulate real-world workloads and determine optimal configurations for the integration of components that make up the overall solution.

## 2. Architecture Overview

This Citrix Validated Solution and its components was designed, built and validated to support two distinct Citrix virtual desktop types. Each desktop type is described to support up to 1,000 user desktop sessions:

- **Hosted Shared Desktops.** Up to 1,000 individual user sessions running XenDesktop [Hosted Shared Desktops](#) on Windows Server 2008 R2 Remote Desktop Session Hosts or
- **Hosted Virtual Desktops.** Up to 1,000 individual XenDesktop [Hosted Virtual Desktops](#) running Windows 7 Enterprise x64.

Each of these desktop types is described in the Citrix [FlexCast](#) model operating as virtual machine instances on Microsoft Hyper-V Server® 2012. This architecture is a single, self-supporting modular component identified as a [Pod](#), supporting up to 1,000 users allowing customers to consistently build and deploy scalable environments.

### 2.1 Citrix Virtual Desktop Types

This Citrix Validated Solution document references Citrix Hosted Shared Desktops and Hosted Virtual Desktops (HVD). Both types of virtual desktops are discussed below for reference. For more information, refer to Citrix FlexCast delivery methods <http://flexcast.citrix.com/>

- **Hosted Shared Desktop (HSD).** A Windows Remote Desktop Session (RDS) Host using Citrix XenDesktop to deliver Hosted Shared Desktops in a locked down, streamlined and standardised manner with a core set of applications. Using a published desktop on to the Remote Desktop Session Host, users are presented a desktop interface similar to a Windows 7 “look and feel”. Each user runs in a separate session on the RDS server.
- **Hosted Virtual Desktop (HVD) aka Hosted VDI.** A Windows 7 desktop instance running as a virtual machine where a single user connects to the machine remotely. Consider this as 1:1 relationship of one user to one desktop. There are differing types of the hosted virtual desktop model (existing, installed, pooled, dedicated and streamed). This document exclusively refers to the pooled type of HVD.

This document will discuss the Citrix Validated Solution for both Hosted Shared Desktops and Hosted Virtual Desktops (pooled desktops). Throughout this document nomenclature may reference the FlexCast model as; “<FlexCast model>” which should be substituted for either HSD or HVD as appropriate to the design under consideration.

### 2.2 The Pod Concept

The term “pod” is referenced throughout this solution design. In the context of this document a pod is a known entity, an architecture that has been pre-tested and validated. A pod consists of the hardware and software components required to deliver 1,000 virtual desktops using either FlexCast model.

For clarity this document does not attempt to describe combining both FlexCast models, it specifically discusses each type as its own entity.

The pod prescribes the physical and logical components required to scale out the number of desktops in increments of 1,000 users or part thereof.

## 2.3 Justification and Validation

The construct of this Citrix Validated Solution is based on many decisions that were made during validation testing. Testing was carried out using the Login VSI virtual Session Indexer (VSI), an industry standard tool for user / session benchmarking. Login VSI allows comparisons of platforms and technologies under the same repeatable load. The “Medium” VSI workload is expected to approximate the average office worker during normal activities and was the workload used throughout testing.

Note. All workloads were tested using the XenDesktop Template Policy “High Server Scalability” running in “Legacy Graphics mode” therefore the Bill of Materials described for each FlexCast model within this document are based on the density of users with these policy settings in place. Using these Citrix Policies allows the greatest host density for each FlexCast model.

## 2.4 High Level Solution Overview

The diagram below depicts the Citrix XenDesktop Hosted Shared Desktop technology stack.



Figure 1. Solution Stack HSD Workload

The diagram below depicts the Citrix XenDesktop Hosted Virtual Desktop technology stack.



Figure 2. Solution Stack HVD Workload

- **Citrix XenDesktop.** Two virtualised Desktop Delivery Controller servers will be deployed to support the XenDesktop Site. A single XenDesktop Site will be utilised to manage the initial 1,000 desktop pod. Additional desktop pods and supporting hardware can be deployed to scale out the XenDesktop Site to thousands of virtual desktops.
- **Virtual Desktops.** This solution will focus on the delivery of the two discrete virtual desktops types:
  - Hosted Virtual Desktops (HVD). Describing the delivery of a 1,000 Pooled Windows 7 virtual desktops powered by Citrix XenDesktop 7.1.
  - Hosted Shared Desktops (HSD). Describing the delivery of a 1,000 Shared desktop based on Microsoft Windows Server 2008 R2 Remote Desktop Session host workloads powered by Citrix XenDesktop 7.1.
- **Microsoft Hyper-V Server 2012 (Hyper-V).** The hypervisor selected to host the virtualised desktop and server instances for this solution is Microsoft Hyper-V Server 2012<sup>®</sup>. Hyper-V will be deployed onto the Cisco UCS blades and configured to boot from iSCSI SAN.
- **Virtual Desktop Provisioning.** This document describes the use of Citrix Provisioning services:
  - **Citrix Provisioning Services (PVS).** Desktop and RDS Server workloads may be streamed using Provisioning Services 7.1 using a predefined vDisk image containing the optimised operating system and Tier-1 application set.
- **Applications.** Tier-2<sup>1</sup> applications which may include line of business or customer specific applications that are not embedded as part of the Disk image may be delivered using Citrix XenDesktop RDS workloads or Microsoft App-V<sup>2</sup>.
- **Citrix StoreFront.** Virtualised StoreFront servers will be deployed to provide application and desktop resource enumeration. The StoreFront servers will be load balanced using Citrix NetScaler appliances.
- **Citrix NetScaler SDX 11500.** NetScaler SDX appliances configured with high availability virtual instances (HA) will be deployed to provide remote access capability to the Hosted Shared Desktops and server load balancing of Citrix services.
- **Citrix Performance Management.** Citrix Director, Citrix NetScaler HDX Insight and Citrix EdgeSight will provide monitoring capabilities into the virtual desktops.
- **Cisco UCS.** The hardware platform of choice for this solution is Cisco UCS consisting of the UCS 5108 chassis and UCS B200 M3 blades. Second generation Fabric Interconnect (6248UP) and Fabric Extenders (2204XP) are utilised. The Hyper-V servers will be hosted on Cisco UCS hardware.
- **Cisco Nexus.** Second generation Cisco Nexus 5548UP switches are used to provide converged network connectivity across the solution using 10GbE.

---

<sup>1</sup> The solution design for Tier-2 applications delivered by Citrix XenDesktop or Citrix XenApp is out of scope for this document.

<sup>2</sup> The solution design of Microsoft App-V components is out of scope for this document.

- **Nimble Storage.** Hypervisor operating system disks will be delivered via boot from iSCSI SAN. Shared storage in the form of iSCSI mounted volumes and Clustered Shared Volumes (CSVs) for virtual disk images.
- **Supporting Infrastructure.** The following components are assumed to exist within the customer environment and are required infrastructure components:
  - Microsoft Active Directory Domain Services
  - A suitable Microsoft SQL database platform to support the solution database requirements.
  - Licensing servers to provide Citrix and Microsoft licenses are assumed to exist.
  - CIFS SMB File sharing. This can be provisioned as part of the solution using Window Server Failover Clustering with the General Use file server role enabled. Please refer to section, [SMB File Services](#).

This design document will focus on the desktop virtualisation components which include the desktop workload, desktop delivery mechanism, hypervisor, hardware, network and storage platforms.

## 2.5 Assumptions

The following assumptions have been made:

- Required Citrix and Microsoft licenses and agreements are available.
- Required power, cooling, rack and data centre space is available.
- No network constraints that would prevent the successful deployment of this design.
- Microsoft Windows Active Directory Domain services are available.
- Microsoft SQL Database platform is available.

### 3. Logical Architecture

#### 3.1 Logical Component Overview for HSD

The logical components that make up the requirements to deliver a 1,000 user XenDesktop Hosted Shared Desktop solution are described in the illustration below:

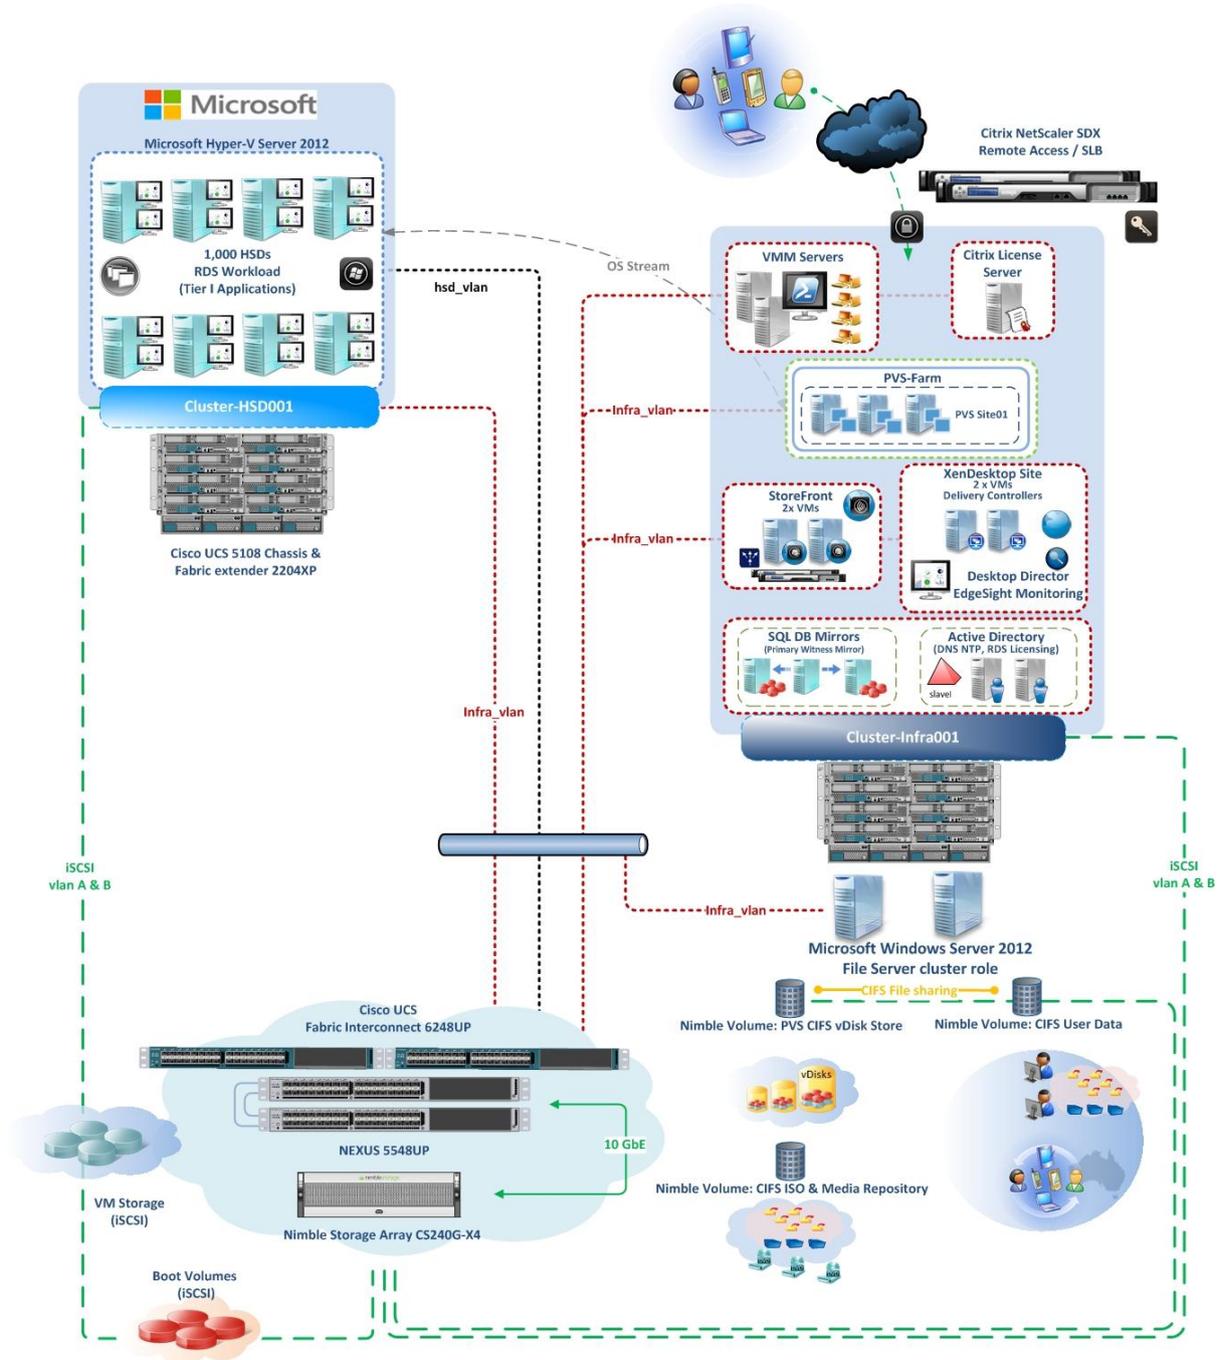


Figure 3. Hosted Shared Desktops - Logical Component View

The following Citrix components are required:

- Citrix XenDesktop – Hosted Shared Desktop virtualisation platform.
- Citrix Provisioning Services – workload delivery platform for Hosted Shared Desktops.
- Citrix User Profile Management - user personalisation.
- Citrix StoreFront – XenDesktop resource enumeration.
- Citrix License Server – pooled management of Citrix licenses.
- Citrix NetScaler SDX 11500 – remote access to the desktop instances and server load balancing capabilities for the Citrix StoreFront servers and other Citrix services.
- Performance Management – Citrix Director EdgeSight and NetScaler HDX Insight.

### 3.2 Logical Component Overview for HVD

The logical components that make up the requirements to deliver a 1,000 user XenDesktop Hosted Virtual Desktop solution are described in the illustration below:

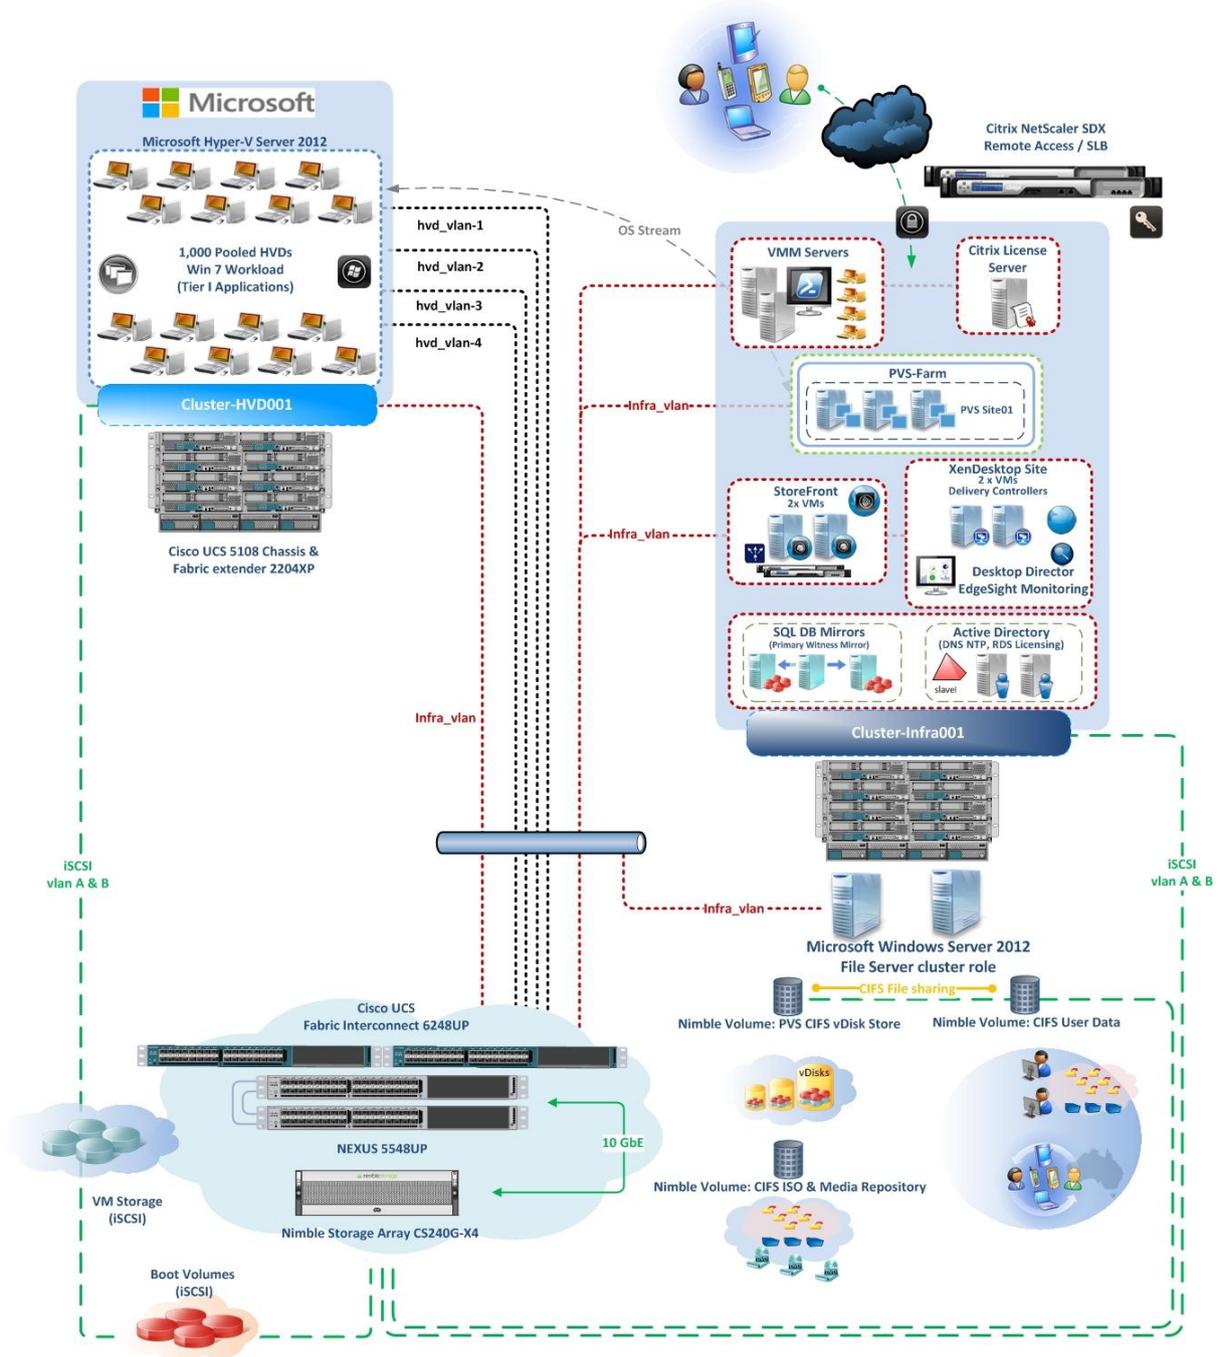


Figure 4. Hosted Virtual Desktops - Logical Component View

The following Citrix components are required:

- Citrix XenDesktop – Hosted Virtual Desktop virtualisation platform.
- Citrix Provisioning Services – workload delivery platform for Hosted Virtual Desktops.
- Citrix User Profile Management - user personalisation.
- Citrix StoreFront – XenDesktop resource enumeration.
- Citrix License Server – pooled management of Citrix licenses.
- Citrix NetScaler SDX 11500 – remote access to the desktop instances and server load balancing capabilities for the Citrix StoreFront servers and other Citrix services.
- Performance Management – Citrix Director EdgeSight and NetScaler HDX Insight.

## 4. Physical Architecture

### 4.1 Physical Component Overview

This Citrix Validated Solution is built on a Cisco Unified Computing System (UCS), Cisco Nexus switches and a Nimble Storage array, these components define the overall hardware architecture.

#### HSD Component Overview:

Figure 5. defines the Cisco and Nimble Storage array hardware components required to provide the 1,000 Hosted Shared Desktop [pod](#) delivered by Citrix XenDesktop.

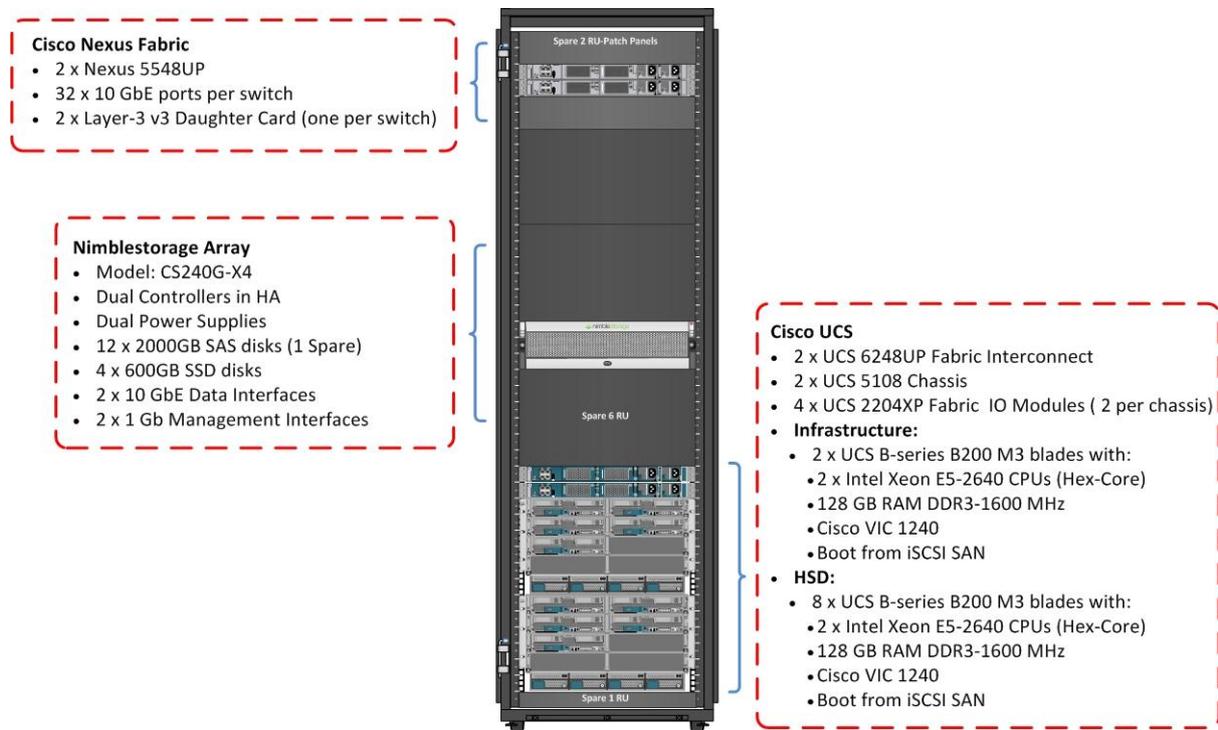


Figure 5. Physical Component View HSD

## HVD Component Overview:

Figure 6. defines the Cisco and Nimble Storage array hardware components required to provide the 1,000 Hosted Virtual Desktop [pod](#) delivered by Citrix XenDesktop.

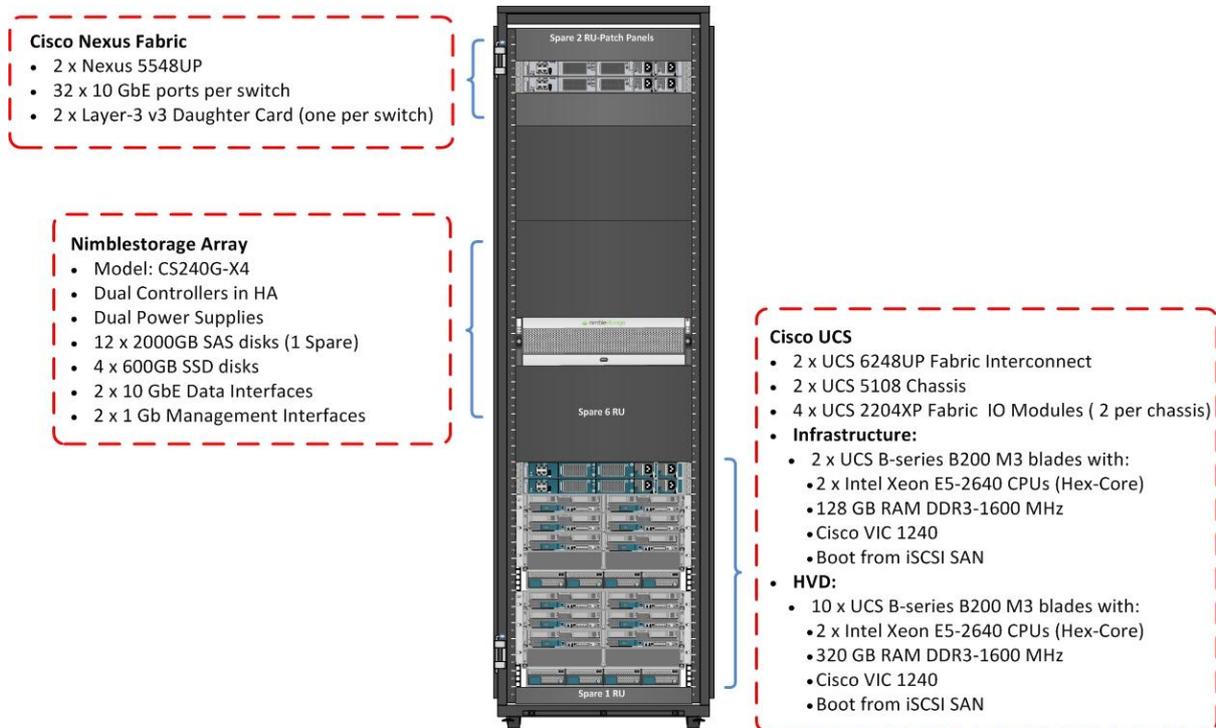


Figure 6. Physical Component View HVD

## Component Overview

Resource	Components – Patches/Revisions
Compute	<ul style="list-style-type: none"> <li>• Cisco UCS Manager 2.1(3a)</li> <li>• Cisco UCS 5108 Blade Server Chassis</li> <li>• Cisco UCS B200 M3 B-Series Blades - (B200M3.2.0.3.0.051620121210)</li> <li>• Dual Intel 2.50 GHz E5-2640 Xeon Processors</li> <li>• Cisco Virtual Interface Card 1240</li> <li>• Cisco UCS 6248UP Series Fabric Interconnects – 5.0(3)N2</li> <li>• Cisco UCS 2204XP Series Fabric Extender</li> </ul> <p><b>HSD:</b></p> <ul style="list-style-type: none"> <li>• 128GB RAM for Infrastructure hosts</li> <li>• 128GB RAM for HSD hosts</li> </ul> <p><b>HVD:</b></p> <ul style="list-style-type: none"> <li>• 128GB RAM for Infrastructure hosts</li> <li>• 320GB RAM for HVD hosts</li> </ul>
Storage	<ul style="list-style-type: none"> <li>• Nimble Storage array</li> <li>• Software version: 1.4.7.0-45626-opt (current version at the time of testing)</li> <li>• Head Shelf, Model: CS240G-X4</li> <li>• Dual HA Controllers</li> <li>• iSCSI for all data paths</li> </ul> <p><b>HSD:</b></p> <ul style="list-style-type: none"> <li>• Internal Disks:</li> <li>• 12 x 2000GB NL-SAS drives</li> <li>• 4 x 600GB SSD drives</li> </ul> <p><b>HVD:</b></p> <ul style="list-style-type: none"> <li>• Internal Disks:</li> <li>• 12 x 2000GB NL-SAS drives</li> <li>• 4 x 600GB SSD drives</li> </ul>
Network	<ul style="list-style-type: none"> <li>• Cisco Nexus 5548UP Series Switch - System version: 5.1(3)N2(1)</li> </ul>
Remote Access & Server Load Balancing	<ul style="list-style-type: none"> <li>• Citrix NetScaler SDX 11500 appliances.</li> <li>• Virtual instances configured in HA.</li> </ul>

Table 1. Hardware Components

## 4.2 Physical Component Design HSD

### BoM Cisco UCS Compute Hardware for HSD

Part Number	Description	Quantity
N20-Z0001	CISCO Unified Computing System	1
N20-C6508	CISCO UCS 5108 Blade Svr AC Chassis/0 PSU/8 fans/0 fabric extender	2
UCSB-B200-M3	CISCO UCS B200 M3 Blade Server w/o CPU, memory, HDD, mLOM/mezz	10
UCS-MR-1X162RY-A	CISCO 16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	16
UCS-MR-1X162RY-A	CISCO 16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	64
UCS-CPU-E5-2640	CISCO 2.50 GHz E5-2640/95W 6C/15MB Cache/DDR3 1333MHz	20
UCSB-MLOM-40G-01	CISCO VIC for UCS blade servers capable of up to 40GbE	10
N20-BBLKD	CISCO UCS 2.5 inch HDD blanking panel	20
UCSB-HS-01-EP	CISCO Heat Sink for UCS B200 M3 server	20
UCS-IOM-2208XP	CISCO UCS 2204XP I/O Module (4 External, 16 Internal 10Gb Ports)	4
N20-PAC5-2500W	CISCO 2500W AC power supply unit for UCS 5108	8
CAB-AC-16A-AUS	CISCO Power Cord, 250VAC, 16A, Australia C19	8
N20-FAN5	CISCO Fan module for UCS 5108	16
N01-UAC1	CISCO Single phase AC power module for UCS 5108	2
N20-CAK	CISCO Access. kit for 5108 Blade Chassis incl Rail kit, KVM dongle	2
N20-FW010	CISCO UCS 5108 Blade Server Chassis FW package	2
UCS-FI-6248UP	CISCO UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC	2
UCS-FI-DL2	CISCO UCS 6248 Layer 2 Daughter Card	2
UCS-BLKE-6200	CISCO UCS 6200 Series Expansion Module Blank	4
UCS-FAN-6248UP	CISCO UCS 6248UP Fan Module	8
UCS-ACC-6248UP	CISCO UCS 6248UP Chassis Accessory Kit	4
N10-MGT010	CISCO UCS Manager v2.0	2
CAB-9K10A-AU	CISCO Power Cord, 250VAC 10A 3112 Plug, Australia	4
UCS-PSU-6248UP-AC	CISCO UCS 6248UP Power Supply/100-240VAC	4
SFP-H10GB-CU5M	CISCO 10GBASE-CU SFP+ Cable 5 Meter	14

Table 2. Cisco UCS Compute Hardware

### **BoM Cisco Nexus Network Hardware for HSD**

Part Number	Description	Quantity
N5K-C5548UP-FA	CISCO Nexus 5548 UP Chassis, 32 10GbE Ports, 2 PS, 2 Fans	2
CAB-9K10A-AU	CISCO Power Cord, 250VAC 10A 3112 Plug, Australia	4
N55-D160L3-V2	CISCO Nexus 5548 Layer 3 Daughter Card, Version 2	2
N5KUK9-513N1.1	CISCO Nexus 5000 Base OS Software Rel 5.1(3)N1(1)	2
N55-PAC-750W	CISCO Nexus 5500 PS, 750W, Front to Back Airflow(Port-Side Outlet)	4
N5548P-FAN	CISCO Nexus 5548P and 5548UP Fan Module, Front to Back Airflow	4
N5548-ACC-KIT	CISCO Nexus 5548 Chassis Accessory Kit	2
N55-M-BLNK	CISCO Nexus 5500 Module Blank Cover	2
N55-BAS1K9	CISCO Layer 3 Base License for Nexus 5500 Platform	2
SFP-10G-SR=	CISCO 10GBASE-SR SFP Module	4

Table 3. Cisco Nexus Switch Hardware

### **BoM Nimble Storage Array Hardware for HSD**

Part Number	Description	Quantity
CS240G-X4	CS240G-X4 Storage Array w/10GbE 24TB Raw, 16-33TB Usable, 2.4TB Flash Cache, 2x10GigE + 2x1GigE, High Perf Ctlr	1

Table 4. Nimble Storage Array Hardware

## **4.3 Hardware Support and Maintenance for HSD**

### **Cisco UCS and Nexus for HSD**

Part Number	Description	Quantity
CON-SNT-2C6508	CISCO UC SUPPORT 8X5XNBD 5108 Blade Server Chassis	2
CON-SNT-B200M3	CISCO UC SUPPORT 8X5XNBD UCS B200 M3 Blade Server	10
CON-SNT-FI6248UP	CISCO UC SUPPORT 8X5XNBD UCS 6248UP 1RU Fabric Interconnect/2PSU/2	2
CON-SNT-C5548UP	CISCO SUPPORT 8X5XNBD Nexus 5548UP	2

Table 5. Cisco UCS and Nexus Maintenance HSD

**Nimble Storage for HSD**

Part Number	Description	Quantity
SLA-CS240-4HR-1YR	4 Hour Serv/Softw Support for 240; 24x7, 1 Yr, Not available in all areas *	1

Table 6. Nimble Maintenance for HSD

#### 4.4 Physical Component Design HVD

##### BoM Cisco UCS Compute Hardware for HVD

Part Number	Description	Quantity
N20-Z0001	CISCO Unified Computing System	1
N20-C6508	CISCO UCS 5108 Blade Svr AC Chassis/0 PSU/8 fans/0 fabric extender	2
UCSB-B200-M3	CISCO UCS B200 M3 Blade Server w/o CPU, memory, HDD, mLOM/mezz	12
UCS-MR-1X162RY-A	CISCO 16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	16
UCS-MR-1X162RY-A	CISCO 16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	200
UCS-CPU-E5-2640	CISCO 2.50 GHz E5-2640/95W 6C/15MB Cache/DDR3 1333MHz	24
UCSB-MLOM-40G-01	CISCO VIC for UCS blade servers capable of up to 40GbE	12
N20-BBLKD	CISCO UCS 2.5 inch HDD blanking panel	24
UCSB-HS-01-EP	CISCO Heat Sink for UCS B200 M3 server	24
UCS-IOM-2208XP	CISCO UCS 2204XP I/O Module (4 External, 16 Internal 10Gb Ports)	4
N20-PAC5-2500W	CISCO 2500W AC power supply unit for UCS 5108	8
CAB-AC-16A-AUS	CISCO Power Cord, 250VAC, 16A, Australia C19	8
N20-FAN5	CISCO Fan module for UCS 5108	16
N01-UAC1	CISCO Single phase AC power module for UCS 5108	2
N20-CAK	CISCO Access. kit for 5108 Blade Chassis incl Rail kit, KVM dongle	2
N20-FW010	CISCO UCS 5108 Blade Server Chassis FW package	2
UCS-FI-6248UP	CISCO UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC	2
UCS-FI-DL2	CISCO UCS 6248 Layer 2 Daughter Card	2
UCS-BLKE-6200	CISCO UCS 6200 Series Expansion Module Blank	4
UCS-FAN-6248UP	CISCO UCS 6248UP Fan Module	8
UCS-ACC-6248UP	CISCO UCS 6248UP Chassis Accessory Kit	4
N10-MGT010	CISCO UCS Manager v2.0	2
CAB-9K10A-AU	CISCO Power Cord, 250VAC 10A 3112 Plug, Australia	4
UCS-PSU-6248UP-AC	CISCO UCS 6248UP Power Supply/100-240VAC	4
SFP-H10GB-CU5M	CISCO 10GBASE-CU SFP+ Cable 5 Meter	14

Table 7. Cisco UCS Compute Hardware

### **BoM Cisco Nexus Network Hardware for HVD**

Part Number	Description	Quantity
N5K-C5548UP-FA	CISCO Nexus 5548 UP Chassis, 32 10GbE Ports, 2 PS, 2 Fans	2
CAB-9K10A-AU	CISCO Power Cord, 250VAC 10A 3112 Plug, Australia	4
N55-D160L3-V2	CISCO Nexus 5548 Layer 3 Daughter Card, Version 2	2
N5KUK9-513N1.1	CISCO Nexus 5000 Base OS Software Rel 5.1(3)N1(1)	2
N55-PAC-750W	CISCO Nexus 5500 PS, 750W, Front to Back Airflow(Port-Side Outlet)	4
N5548P-FAN	CISCO Nexus 5548P and 5548UP Fan Module, Front to Back Airflow	4
N5548-ACC-KIT	CISCO Nexus 5548 Chassis Accessory Kit	2
N55-M-BLNK	CISCO Nexus 5500 Module Blank Cover	2
N55-BAS1K9	CISCO Layer 3 Base License for Nexus 5500 Platform	2
SFP-10G-SR=	CISCO 10GBASE-SR SFP Module	4

Table 8. Cisco Nexus Switch Hardware

### **BoM Nimble Storage Array Hardware for HVD**

Part Number	Description	Quantity
CS240G-X4	CS240G-X4 Storage Array w/10GbE 24TB Raw, 16-33TB Usable, 2.4TB Flash Cache, 2x10GigE + 2x1GigE, High Perf Ctlr	1

Table 9. Nimble Storage Array Hardware

## **4.5 Hardware Support and Maintenance for HVD**

### **Cisco UCS and Nexus for HVD**

Part Number	Description	Quantity
CON-SNT-2C6508	CISCO UC SUPPORT 8X5XNBD 5108 Blade Server Chassis	2
CON-SNT-B200M3	CISCO UC SUPPORT 8X5XNBD UCS B200 M3 Blade Server	12
CON-SNT-FI6248UP	CISCO UC SUPPORT 8X5XNBD UCS 6248UP 1RU Fabric Interconnect/2PSU/2	2
CON-SNT-C5548UP	CISCO SUPPORT 8X5XNBD Nexus 5548UP	2

Table 10. Cisco UCS and Nexus Maintenance HVD

**Nimble Storage for HVD**

Part Number	Description	Quantity
SLA-CS240-4HR-1YR	4 Hour Serv/Softw Support for 240; 24x7, 1 Yr, Not available in all areas *	1

Table 11. Nimble Maintenance for HVD

## 5. High-Level Design

### 5.1 Network / Cisco Nexus

#### Overview

Nexus A and Nexus B identify the pair of Cisco Nexus 5548UP switches that will be deployed as part of the solution forming the network switching components of the architecture

Figure 7. illustrates the high-level connectivity for the individual components:

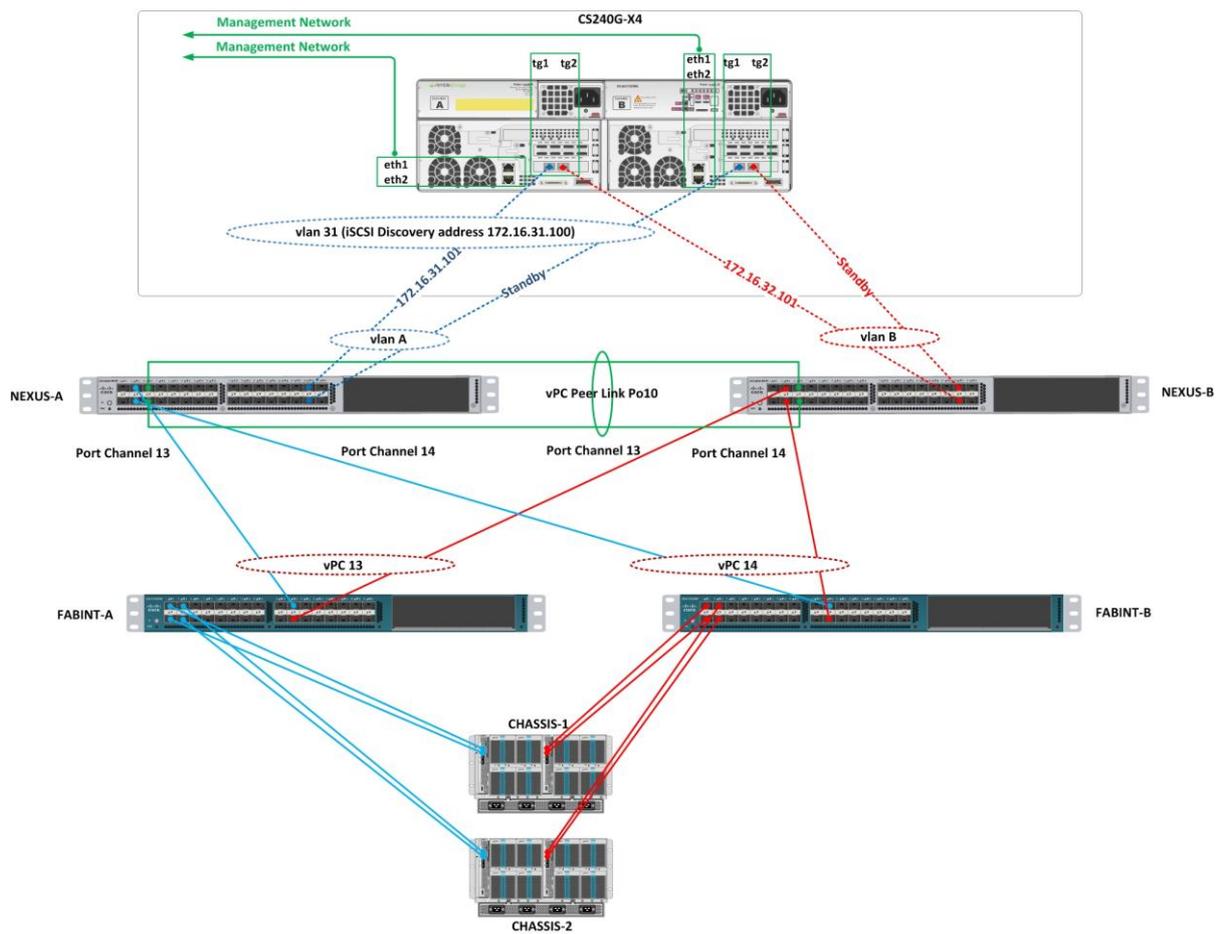


Figure 7. Network Component Connectivity

## Key Decisions

Decision Point	Description / Decision
Nexus Switch Firmware	<ul style="list-style-type: none"> <li>Cisco Nexus 5548UP Series Switch - System version: 5.1(3)N2(1)</li> </ul>
Layer 3 Routing	<ul style="list-style-type: none"> <li>Optional<sup>3</sup></li> </ul>
Number of Switches	Two: <ul style="list-style-type: none"> <li>Nexus 5548UP A</li> <li>Nexus 5548UP B</li> </ul>

Table 12. Cisco Nexus Key Decisions

## Design

### HSD Solution vlan requirements:

Vlan Name	Vlan ID	Description
Hostmgmt_vlan	vlan 20	Host Management vlan.
Infraserver_vlan	vlan 25	Infrastructure server vlan.
iscsi_vlan_A	vlan 31	iSCSI storage for Fabric A vlan.
iscsi_vlan_B	vlan 32	iSCSI storage for Fabric B vlan.
hyperv-Live-Migration	vlan 33	Hyper-V VM live migration vlan.
hsd_vlan-1	vlan 80	HSD worker vlan.
vPC-Native-vlan	vlan 2	Native vlan for vPC untagged packets.
VM-Cluster-HB-vlan	vlan 34	VM File Server Cluster Heartbeat <sup>4</sup>

Table 13. Cisco Nexus/UCS vlan requirements for HSD

### HVD Solution vlan requirements:

Vlan Name	Vlan ID	Description
Hostmgmt_vlan	vlan 20	Host Management vlan.
Infraserver_vlan	vlan 25	Infrastructure server vlan.
iscsi_vlan_A	vlan 31	iSCSI storage for Fabric A vlan.
iscsi_vlan_B	vlan 32	iSCSI storage for Fabric B vlan.
hyperv-Live-Migration	vlan 33	Hyper-V VM live migration vlan.
hvd_vlan-1	vlan 40	HVD vlan 1
hvd_vlan-2	vlan 42	HVD vlan 2
hvd_vlan-3	vlan 44	HVD vlan 3

<sup>3</sup> Layer 3 routing on the Cisco Nexus switch provided by N55-BAS1K9 - Cisco Layer 3 Base License for Nexus 5500 Platform. This can either be terminated on the Nexus 5548UP level or utilise existing network infrastructure.

<sup>4</sup> Required if using the virtual file server cluster options, refer to the section [SMB File Services](#).

Vlan Name	Vlan ID	Description
hvd_vlan-4	vlan 46	HVD vlan 4
vPC-Native-vlan	vlan 2	Native vlan for vPC untagged packets.
VM-Cluster-HB-vlan	vlan 34	VM File Server Cluster Heartbeat5

Table 14. Cisco Nexus/UCS vlan requirements for HVD

At a high-level, the pair of Nexus switches will provide Layer 2 redundancy using virtual port channel configurations between the switch pair and the Fabric Interconnects (vPCs). Layer 3 routing is expected to be carried out by the customer's existing aggregation or core switching layer infrastructure.

Optionally, Layer 3 can be configured on the Nexus 5548UP switch pair<sup>6</sup> using Hot Standby Router Protocol (HSRP) to add Layer 3 redundancy capability. Connectivity to the Nimble Storage array is by individual switch ports on each Nexus switch with network redundancy and failover being provided at the Nimble Storage array level in conjunction with the Nexus switch pair and the Microsoft Windows Server 2012 native Multipath I/O driver.

<sup>5</sup> Required if using the virtual file server cluster options, refer to the section [SMB File Services](#).

<sup>6</sup> Layer 3 routing on the Cisco Nexus switch provided by N55-BAS1K9 - Cisco Layer 3 Base License for Nexus 5500 Platform.

## 5.2 Cisco UCS

### Overview

Cisco UCS comprises a number of physical and logical entities, managed by Cisco UCS Manager. Cisco UCS provides the next-generation data centre platform that unifies computing, networking, storage access, and virtualisation resources into a single unified system.

### Key Decisions

Decision Point	Description / Decision
Service Profiles	Allows the servers to be stateless, with logical entities provided as part of a profile that applies: identity connectivity HBA, NIC firmware and other assignments based on templates that can be assigned to a Server hardware profile.
Network	The Cisco Virtual Interface Card 1240 will be configured within the service profile to present the following Networks to each Hyper-V host: <ul style="list-style-type: none"> <li>• Live Migration (CSV I/O Redirection Network)</li> <li>• VM Traffic (multiple vlans)</li> <li>• Host Management (Cluster Heartbeat)</li> <li>• iSCSI storage (IPSAN)</li> </ul>
Storage	The Cisco UCS B200 M3 blades will be diskless and configured to boot from iSCSI storage presented from the Nimble Storage array.
Service Profile Templates	Two Service profile templates are required: Service Profile Infrastructure Hyper-V hosts: <ul style="list-style-type: none"> <li>• HyperV_Infra_BootiSCSI</li> </ul> Service Profile <FlexCast model> Hyper-V hosts: <ul style="list-style-type: none"> <li>• HyperV_&lt;FlexCast model&gt;_BootiSCSI</li> </ul>
UUID Suffix Pool	Single UUID Suffix Pool: <ul style="list-style-type: none"> <li>• Hyper-V-Hosts</li> </ul>
MAC Address Pool	Two MAC Pools: <ul style="list-style-type: none"> <li>• Fabric-A</li> <li>• Fabric-B</li> </ul>
iSCSI Initiator IP Pools	Two IP Pools are required, HyperV-iSCSI-Initiator-Pools: <ul style="list-style-type: none"> <li>• IP Range Fabric-A</li> <li>• IP Range Fabric-B</li> </ul>
IQN Pools	Four Pools are required: For Service Profile HyperV_<FlexCast model>_BootiSCSI: <ul style="list-style-type: none"> <li>• iSCSI-&lt;FlexCast model&gt;-Fabric-A</li> <li>• iSCSI-&lt;FlexCast model&gt;-Fabric-B</li> </ul> For Service Profile HyperV_Infra_BootiSCSI: <ul style="list-style-type: none"> <li>• iSCSI-Infra-Fabric-A</li> <li>• iSCSI-Infra-Fabric-B</li> </ul>
QoS Policies	Two QoS Policies are required:

Decision Point	Description / Decision
	<ul style="list-style-type: none"> <li>• LiveMigration</li> <li>• iSCSI</li> </ul>
Boot Policies	Name: <ul style="list-style-type: none"> <li>• Boot-iSCSI</li> </ul>
vNIC Templates	Hyper-V Live Migration network: <ul style="list-style-type: none"> <li>• HV-LiveMig-Fab-A</li> <li>• HV-LiveMig-Fab-B</li> </ul> Hyper-V host management network: <ul style="list-style-type: none"> <li>• HV-MGMT-Fab-A</li> <li>• HV-MGMT-Fab-B</li> </ul> VM Data for HyperV_<FlexCast model>_BootiSCSI Service Profile: <ul style="list-style-type: none"> <li>• HV-VM-&lt;FlexCast model&gt;-Fab-A</li> <li>• HV-VM-&lt;FlexCast model&gt;-Fab-B</li> </ul> VM Data for HyperV_Infra_BootiSCSI Service Profile: <ul style="list-style-type: none"> <li>• HV-VM-INF-Fab-A</li> <li>• HV-VM-INF-Fab-B</li> </ul> iSCSI Traffic: <ul style="list-style-type: none"> <li>• HV-iSCSI-Fab-A</li> <li>• HV-iSCSI-Fab-B</li> </ul> VM Cluster Heart Beat: <ul style="list-style-type: none"> <li>• FS-VM-CHB-Fab-A</li> <li>• FS-VM-CHB-Fab-B</li> </ul>
BIOS Policies	Hyper-V_BIOS

Table 15. Cisco UCS Key Decisions

## **Design**

### **Hosted Shared Desktop:**

Two Cisco UCS 5108 Blade Server Chassis will be deployed to support 10 (2 x Infrastructure nodes and 8 x HSD nodes) Cisco UCS B200 M3 B-Series Blades that will define the Server 2012 Hyper-V hosts.

### **Hosted Shared Desktop:**

Two Cisco UCS 5108 Blade Server Chassis will be deployed to support 12 (2 x Infrastructure nodes and 10 x HVD nodes) Cisco UCS B200 M3 B-Series Blades that will define the Server 2012 Hyper-V hosts.

## Common

Cisco UCS 6248UP Series Fabric Interconnects will provide the connectivity to the Cisco UCS 2204XP Series Fabric Extender fitted to each 5108 Blade Server Chassis. Cisco UCS Manager will be used to create the Service Profiles defining the virtual and logical entities required to configure each component.

Each Hyper-V host server will be configured with multiple paths to the Nimble Storage array using iSCSI. Separate vlans from Fabric A and Fabric B using the Microsoft Windows Server 2012 native Multipath I/O driver will be used. Least queue depth load balancing method for iSCSI data traffic will be utilised as per Nimble Storage recommendation and best practice.

## 5.3 Nimble Storage

### Overview

The storage platform utilised for this solution is a Nimble Storage array with internal disk drives only (no additional expansion shelves). At a high level, the Nimble Storage array provides the following features:

- **CASL™ architecture.** Patented “Cache Accelerated Sequential Layout” CASL features include:
  - Dynamic Caching using SSDs to cache data and metadata in flash reads
  - Write-Optimized Data Layout
  - Application-Tuned Block Size
  - Universal Compression
  - Efficient, Instant Snapshots
  - Efficient Replication
  - Zero-Copy Clones
- **Array Hardware.** Features include:
  - Dual Controller Architecture
  - Dual Power Supplies
  - Capacitor backed Non Volatile Random Access Memory ensures that all writes to the array not yet committed to disk are safely protected in the event of an unexpected power outage.
  - Nimble Storage arrays utilise RAID 6 which provides dual parity for disk protection.
  - Each Nimble Storage controller shelf and expansion shelf houses a single hot spare drive.
  - Dedicated Ethernet 10GbE for data traffic
  - Dedicated Ethernet 1 Gb for management traffic

Figure 8. below provides a high level overview of the Nimble Storage architecture and describes a typical Hyper-V Host.

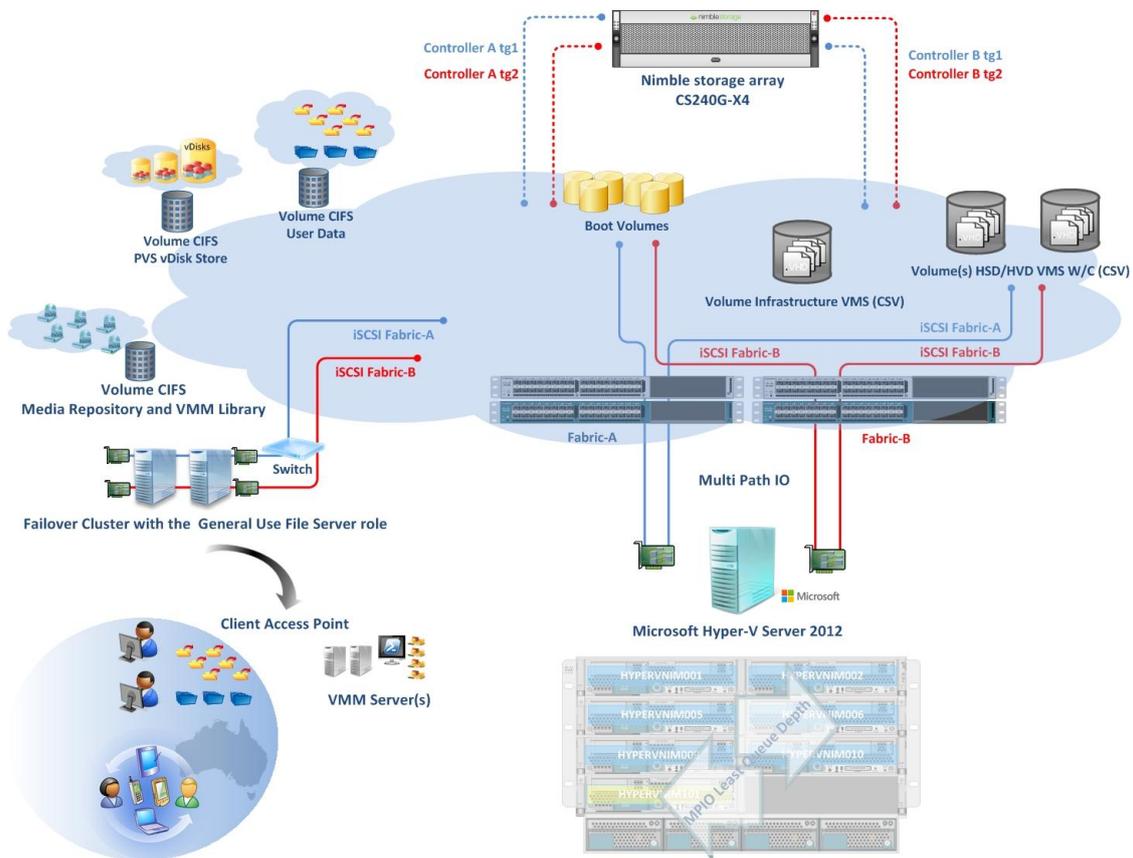


Figure 8. Nimble Storage - System Overview

### Key Decisions

Decision Point	Description / Decision
Hardware Details	<ul style="list-style-type: none"> <li>Nimble array CS240G-X4</li> </ul>
Software Version	<ul style="list-style-type: none"> <li>1.4.7.0-45626-opt</li> </ul>
iSCSI Initiator Groups	<ul style="list-style-type: none"> <li>Required, per volume</li> </ul>
Storage Types	<ul style="list-style-type: none"> <li>iSCSI Hyper-V Boot volumes</li> <li>iSCSI hyper-V CSV volumes</li> <li>iSCSI volumes for CIFS file sharing</li> </ul>
Thin Provisioning (volume reserve)	<ul style="list-style-type: none"> <li>Enabled</li> </ul>
Networks	<ul style="list-style-type: none"> <li>Dedicated 10GbE interfaces will be used for data traffic</li> <li>Dedicated 1 Gb interfaces will be used for management traffic</li> <li>Two discreet vlans will be used to separate data traffic (iSCSI Fabric A and iSCSI Fabric B) through the fabric to the</li> </ul>

Decision Point	Description / Decision
	10GbE interfaces on the array <ul style="list-style-type: none"> <li>• iSCSI discovery will be configured using the 2 data addresses</li> </ul>
MTU	<ul style="list-style-type: none"> <li>• Jumbo Frames will be enabled for both data interfaces on each controller at 9000 MTU</li> </ul>
Performance Policies	Performance Policies: <ul style="list-style-type: none"> <li>• “Hyper-V CSV” for Clustered Shared volumes               <ul style="list-style-type: none"> <li>• Compress: On</li> <li>• Cache: On</li> </ul> </li> <li>• “Default” for Boot volumes               <ul style="list-style-type: none"> <li>• Compress: On</li> <li>• Cache: On</li> </ul> </li> <li>• “Windows File Server” for CIFS volumes               <ul style="list-style-type: none"> <li>• Compress: On</li> <li>• Cache: On</li> </ul> </li> </ul>
Multipath I/O	Native Windows Server 2012 MPIO driver.
MPIO Load Balancing Method	Least Queue Depth (LQD)
SMTP Server	An SMTP server will be specified to allow the array to send email alerts.
Auto Support	Send “Auto Support” data to Nimble Storage support will be checked to allow the array to upload data to Nimble technical support. Proxy Server: <ul style="list-style-type: none"> <li>• Optional</li> </ul>
SNMP	Enabled as per customer requirements

Table 16. Nimble Storage Key Decisions

## Design

The Nimble Storage array CS240G-X4 used within this design provides a highly available, redundant controller solution within a single 3U enclosure. The Nimble Storage array is a converged storage and backup system in one, which contains sufficient internal disk capacity providing the required performance to meet the demands of the solution.

Each controller in the Nimble Storage array high availability pair is connected with dual data paths to the network, which allows the storage system to operate in the event of component failure. A single failure of a data path will not result in a controller failover. From the hypervisor host server perspective, a multipath I/O driver will be used to ensure the optimum path to the storage layer.

The Nimble Storage array only supports block based storage using iSCSI. This CVS design document discusses the use of a Microsoft Windows-based File server for the purpose of hosting SMB file shares for data such as user data, user profiles, ISO media repository and Citrix Provisioning Services vDisk image files.

The following sections contain recommended configuration parameters for the logical storage entities.

### Required Volumes for HSD:

Volume Name	Performance Policy	Volume Size	Description
Quorum-Infra01	Default	2GB	Hyper-V Infrastructure Failover Cluster disk witness
Quorum-Infra02	Default	2GB	File Server VM Failover Cluster disk witness
Quorum-hsd01	Default	2GB	HVD Failover Cluster disk witness
hypervnimxxx	Default	2,000GB	Hyper-V Boot volumes, where xxx represents the server's ordinal number.
infra-pvs01	Windows File Server	1,000GB	PVS CIFS Share for vDisk storage.
infra_iso01	Windows File Server	1,000GB	Media and ISO repository.
infra_cifs01	Windows File Server	1,000GB	UPM data and Redirected Folders. Assuming 1GB data per user.
hsd-csv01	Hyper-V CSV	2,500GB <sup>7</sup>	HSD RDS VM storage, PVS write cache drives and hypervisor.
infra-csv01	Hyper-V CSV	2,500GB	Infrastructure VM virtual disks.
<b>TOTAL</b>		<b>10,006GB ~10TB</b>	

Table 17. Required Nimble Storage volumes for HSD

<sup>7</sup> Minimum storage requirement: The total storage size is based on a 20GB persistent drive and ~16 GB Hypervisor overhead (VM Memory) per RDS server. This drive will contain the Windows Pagefile, PVS write cache and redirected logs.

Required Volumes for HVD:

Volume Name	Performance Policy	Volume Size	Description
quorum-Infra01	Default	2GB	Hyper-V Infrastructure Failover Cluster disk witness
quorum-Infra02	Default	2GB	File Server VM Failover Cluster disk witness
quorum-hvd01	Default	2GB	HVD Failover Cluster disk witness
hypervnimxxx	Default	2,000GB	Hyper-V Boot volumes, where xxx represents the server's ordinal number.
infra-pvs01	Windows File Server	1,000GB	PVS CIFS Share for vDisk storage.
infra_iso01	Windows File Server	1,000GB	Media and ISO repository.
infra_cifs01	Windows File Server	1,000GB	UPM data and Redirected Folders. Assuming 1GB data per user.
hvd-csv01	Hyper-V CSV	13,000GB <sup>8</sup>	HVD VM storage, PVS write cache drives and hypervisor.
infra-csv01	Hyper-V CSV	2,500GB	Infrastructure VM virtual disks.
<b>TOTAL</b>		<b>20,506GB ~21TB</b>	

Table 18. Required Nimble Storage volumes for HVD

Volume Parameters:

Volume Name	Volume Reserve	Volume Quota	Volume Warning	Description
hypervnimxxx	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>
infra-pvs01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>
infra_iso01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>
infra_cifs01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>
hsd-csv01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>

<sup>8</sup> Minimum storage requirement: The total storage size is based on a 10GB persistent drive and ~3 GB Hypervisor overhead (VM Memory) per VM guest. This drive will contain the Windows Pagefile, PVS write cache and redirected logs.

Volume Name	Volume Reserve	Volume Quota	Volume Warning	Description
hvd-csv01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>
infra-csv01	0%	100%	80%	<ul style="list-style-type: none"> <li>• 100% Thin Provisioned</li> <li>• Allows 100% usage of disk</li> <li>• Warn at 80% utilisation</li> </ul>

Table 19. Nimble Volume configuration

Volume Snapshot Parameters:

Volume Name	Snapshot Reserve	Snapshot Quota	Snapshot Warning	Description
hypervnimxxx	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
infra-pvs01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
infra_iso01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
infra_cifs01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
hsd-csv01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
hvd-csv01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>
infra-csv01	0%	Unlimited (selection box checked)	0%	<ul style="list-style-type: none"> <li>• No snapshot reserves</li> <li>• No snapshot quota limits</li> <li>• No warnings</li> </ul>

Table 20. Nimble Volume snapshot configuration

## Initiator Groups

Initiator Group	Volume Access	Initiator Names
<b>Virtual Desktop Hosts</b>		
ig-hypervnim001	<ul style="list-style-type: none"> <li>hypervnim001</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:001 iqn.2013-07.com.microsoft.b.hypervnim:001
ig-hypervnim002	<ul style="list-style-type: none"> <li>hypervnim002</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:002 iqn.2013-07.com.microsoft.b.hypervnim:002
ig-hypervnim003	<ul style="list-style-type: none"> <li>hypervnim003</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:003 iqn.2013-07.com.microsoft.b.hypervnim:003
ig-hypervnim004	<ul style="list-style-type: none"> <li>hypervnim004</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:004 iqn.2013-07.com.microsoft.b.hypervnim:004
ig-hypervnim005	<ul style="list-style-type: none"> <li>hypervnim005</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:005 iqn.2013-07.com.microsoft.b.hypervnim:005
ig-hypervnim006	<ul style="list-style-type: none"> <li>hypervnim006</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:006 iqn.2013-07.com.microsoft.b.hypervnim:006
ig-hypervnim007	<ul style="list-style-type: none"> <li>hypervnim007</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:007 iqn.2013-07.com.microsoft.b.hypervnim:007
ig-hypervnim008	<ul style="list-style-type: none"> <li>hypervnim008</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:008 iqn.2013-07.com.microsoft.b.hypervnim:008
ig-hypervnim009	<ul style="list-style-type: none"> <li>hypervnim009</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:009 iqn.2013-07.com.microsoft.b.hypervnim:009
ig-hypervnim010	<ul style="list-style-type: none"> <li>hypervnim010</li> <li>&lt;FlexCast model&gt;-csv&lt;xx&gt;</li> <li>Quorum-&lt;FlexCast model&gt;</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:010 iqn.2013-07.com.microsoft.b.hypervnim:010
<b>Infrastructure Hosts</b>		
ig-hypervnim101	<ul style="list-style-type: none"> <li>hypervnim101</li> <li>infra-csv01</li> <li>Quorum-Infra01</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:101 iqn.2013-07.com.microsoft.b.hypervnim:101
ig-hypervnim102	<ul style="list-style-type: none"> <li>hypervnim102</li> <li>infra-csv01</li> <li>Quorum-Infra01</li> </ul>	iqn.2013-07.com.microsoft.a.hypervnim:102 iqn.2013-07.com.microsoft.b.hypervnim:102
<b> </b>		

Initiator Group	Volume Access	Initiator Names
<u>File Server Cluster Nodes</u>		
ig-cifscluster01	<ul style="list-style-type: none"> <li>• infra_cifs01</li> <li>• infra-pvs01</li> <li>• infra_iso01</li> <li>• Quorum-Infra02</li> </ul>	iqn.2013-07.com.microsoft:<fileservername01> iqn.2013-07.com.microsoft:<fileservername02>

Table 21. Nimble Storage Initiator Groups for HSD and HVD

## 5.4 Microsoft Hyper-V Server 2012

### Overview

Microsoft Hyper-V Server® 2012 is utilised to provide the hypervisor hosting platform to the virtualised desktop and infrastructure server instances required to support the 1,000 user pod solution.

Figure 9. below depicts the physical connectivity between Cisco UCS blade chassis, Cisco 6248UP Fabric Interconnects, Cisco Nexus 5548UP switches and the Nimble Storage CS240G-X4 array:

- Converged network with a total of 4 x 10GbE server ports per Cisco UCS Chassis (2 x 10GbE connections per Fabric Extender).
- 2 x 10GbE uplink connections between the Fabric Interconnect and Nexus switch layer.
- 2 x 10GbE connections per Nimble Storage CS240G-X4 array to support iSCSI data traffic.

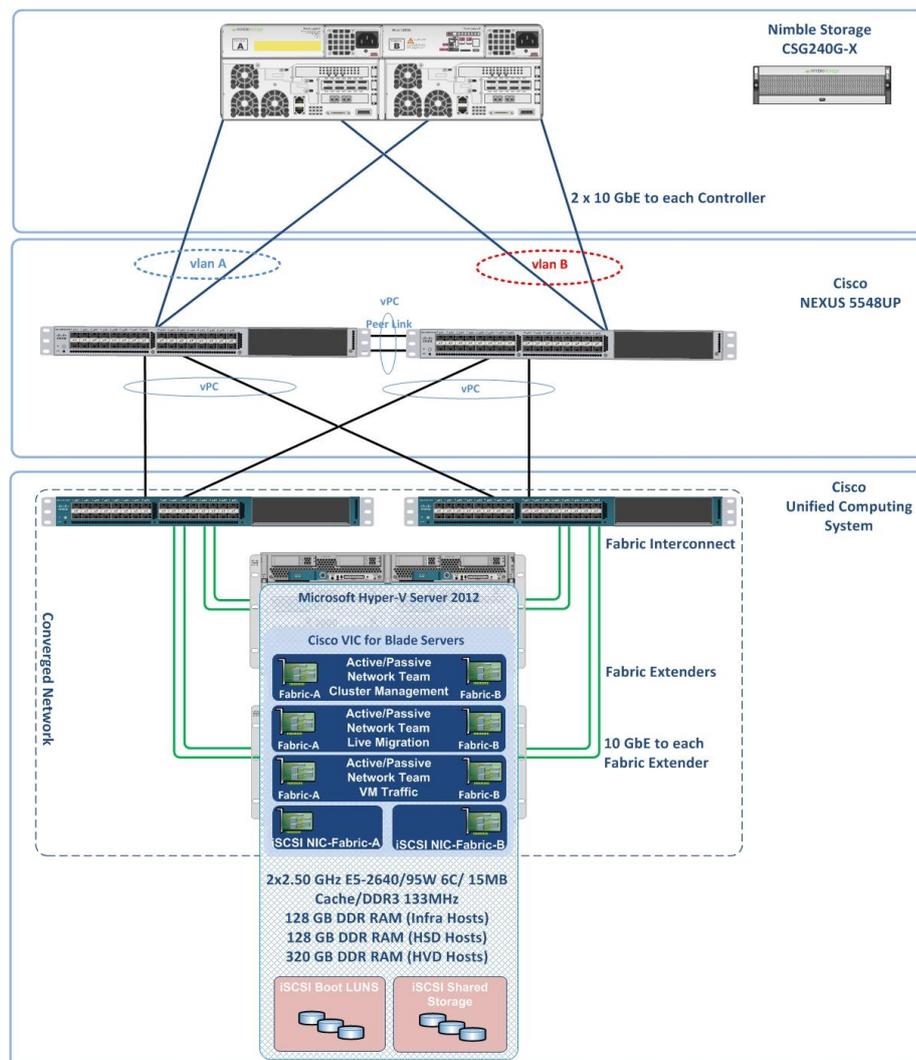


Figure 9. Hyper-V Host Configuration

## Key Decisions

Configuration	Decision
Version	Microsoft Hyper-V Server® 2012
Hardware Setting for Infrastructure nodes	2 x Cisco UCS B200 M3 blades spread across 2 x UCS 5108 chassis: <ul style="list-style-type: none"> <li>• 2 x Intel Xeon 2.50 GHz E5-2640 CPUs(12 cores, 24 with HT enabled)</li> <li>• 128GB RAM DDR3-1600-MHz</li> <li>• Cisco VIC 1240</li> <li>• Diskless blades – boot from iSCSI SAN</li> </ul>
Hardware Settings for HSD nodes	10 x Cisco UCS B200 M3 blades spread across 2 x UCS 5108 chassis: <ul style="list-style-type: none"> <li>• 2 x Intel Xeon 2.50 GHz E5-2640 CPUs(12 cores, 24 with HT enabled)</li> <li>• 128GB RAM DDR3-1600-MHz</li> <li>• Cisco VIC 1240</li> <li>• Diskless blades – boot from iSCSI SAN</li> </ul>
Hardware Settings for HVD nodes	10 x Cisco UCS B200 M3 blades spread across 2 x UCS 5108 chassis: <ul style="list-style-type: none"> <li>• 2 x Intel Xeon 2.50 GHz E5-2640 CPUs(12 cores, 24 with HT enabled)</li> <li>• 320GB RAM DDR3-1600-MHz</li> <li>• Cisco VIC 1240</li> <li>• Diskless blades – boot from iSCSI SAN</li> </ul>
Storage Settings	<ul style="list-style-type: none"> <li>• Boot from iSCSI SAN</li> <li>• Shared storage using iSCSI (CSV's)</li> </ul>
Network Settings	Cisco VIC 1240 presenting 8 x vNICs to each host: <ul style="list-style-type: none"> <li>• 1 x iSCSI Fabric-A (Boot and shared storage)</li> <li>• 1 x iSCSI Fabric-B (Boot and shared storage)</li> </ul> <p><b>Network Team:</b> (Cluster Heartbeat path 1)</p> <ul style="list-style-type: none"> <li>• Team-Host-Management (Active Passive, Switch Independent Mode)</li> <li>• Host-Management-Fabric-A</li> <li>• Host-Management-Fabric-B</li> </ul> <p><b>Network Team:</b> (Cluster Heartbeat path 2, internal cluster network)</p> <ul style="list-style-type: none"> <li>• Team-Live-Migration (Active Passive, Switch Independent Mode)</li> <li>• Live-Mig-Fabric-A</li> <li>• Live-Mig-Fabric-B</li> </ul> <p><b>Network Team:</b></p> <ul style="list-style-type: none"> <li>• Team-VM-Data (Active Passive, Switch Independent Mode)</li> <li>• VM-Data-Fabric-A (Trunk Ports)</li> <li>• VM-Data-Fabric-B (Trunk Ports)</li> </ul>
Cluster Shared Volumes	The requirement for a dedicated CSV network is considered a low priority for the solution. The Live Migration of Guest VMs is also considered a low priority; each component of the architecture is redundant and can allow at least a single component failure without loss of service. Therefore a dedicated Live Migration network was deemed unnecessary; the Live Migration network will be shared with CSV I/O redirection traffic (in the unlikely event I/O redirection occurs).

Configuration	Decision
Hyper-V Switch	VM-Switch01: <ul style="list-style-type: none"> <li>Associated Hyper-V interface: Team-VM-Data</li> </ul>
Failover clustering	<p><b>Infrastructure Hyper-V hosts (2 x hosts):</b></p> <ul style="list-style-type: none"> <li>Failover Clustering enabled <ul style="list-style-type: none"> <li>Cluster name: clust-infra001</li> <li>Node and Disk Majority</li> </ul> </li> <li>High Availability is required</li> <li>Availability Sets are required: <ul style="list-style-type: none"> <li>DHCP Role</li> <li>PVS Role</li> <li>Delivery Controller Role</li> <li>File Services Role (refer to section <a href="#">SMB File Services</a>)</li> </ul> </li> </ul> <p><b>Hosted Shared Desktop Hyper-V hosts (8 x hosts):</b></p> <ul style="list-style-type: none"> <li>Failover Clustering enabled <ul style="list-style-type: none"> <li>Cluster name: clust-hsd001</li> <li>Node Majority (adjust voting configuration)</li> </ul> </li> <li>High Availability is required</li> </ul> <p><b>Hosted Virtual Desktop Hyper-V hosts (10 x hosts):</b></p> <ul style="list-style-type: none"> <li>Failover Clustering enabled <ul style="list-style-type: none"> <li>Cluster name: clust-hvd001</li> <li>Node Majority (adjust voting configuration)</li> </ul> </li> <li>High Availability is required</li> </ul>
Scale-out Recommendation	Additional pods should be deployed to scale out HSD or HVD capacity, thus additional Hyper-V host/failover clusters will be subsequently added.
System Center 2012 - Virtual Machine Manager SP1 (VMM) - Hardware Settings	<ul style="list-style-type: none"> <li>2 x Windows Server 2012 Standard (Guest Clustering)</li> <li>4 vCPUs</li> <li>16GB RAM</li> <li>100GB disk for Operating System (C:\)</li> <li>1 vNIC for production traffic</li> </ul>
Clustered Shared Volume Cache	<ul style="list-style-type: none"> <li>Enabled 2 GB for all clusters</li> <li>Enabled for all Clustered Shared Volumes</li> </ul>

Table 22. Hyper-V Key Decisions

## Design

**Virtual Machine Manager.** System Center 2012 - Virtual Machine Manager (VMM) will be deployed as the management solution for the virtualised environment. VMM will provide the management interface to the virtualised Hyper-V environment for VM Templates, logical networks Hyper-V hosts, failover clusters and other related services.

High availability mode will be deployed using a guest based System Center VMM 2012 SP1 failover cluster for VMM redundancy.

**Chassis Assignment.** Hyper-V hosts will be configured such that hosts with even numbers have their primary network teams configured to the active NIC on Fabric-A and hosts with odd numbers have their primary network teams configured to the active NIC on Fabric-B. iSCSI traffic will be roughly distributed across both Fabrics using the “Least Queue Depth” MPIO load balancing method. This is to ensure even distribution of traffic across the Fabric and minimise impact in the event of a Fabric failure.

*Figure 10.* defines the Hyper-V hosts to physical chassis assignment which makes up the pod of a 1,000 user HSD desktop sessions (48 x RDS virtual machine instances):

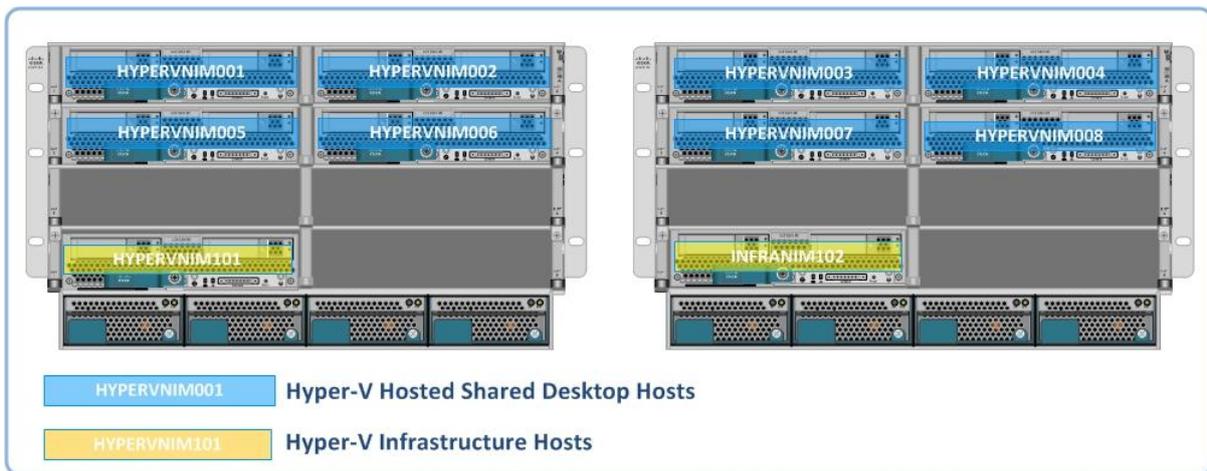


Figure 10. Hyper-V Host to Chassis Assignment for HSDs

*Figure 11.* defines the Hyper-V hosts to physical chassis assignment which makes up the pod of a 1,000 user HVD desktop sessions (1,000 x Windows 7 virtual machine instances):

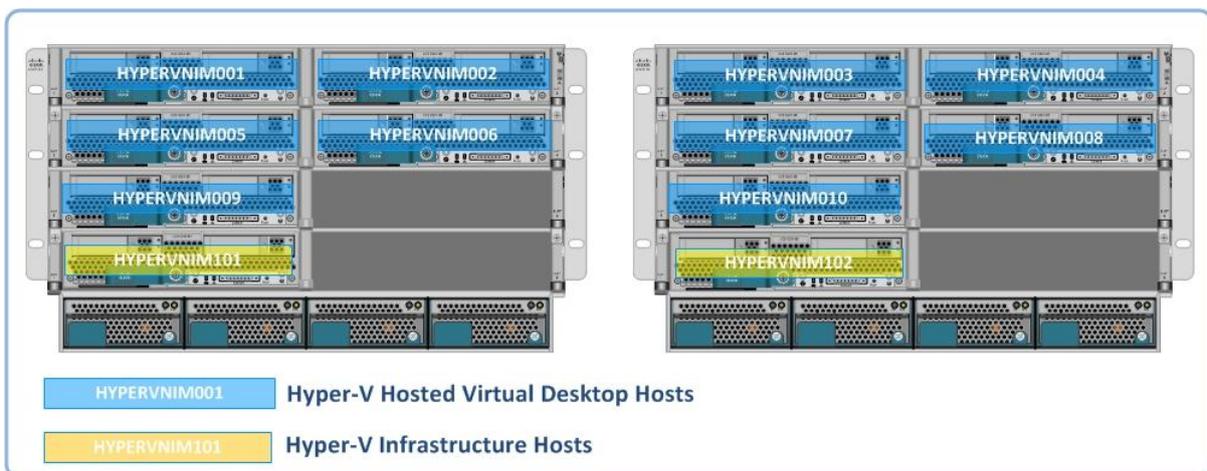


Figure 11. Hyper-V Host to Chassis Assignment for HVDs

**Storage Environment.** As the Cisco UCS B200 M3 blades are diskless, the Hyper-V hosts will be configured to boot from SAN via iSCSI. Shared storage used to host the virtual machine disk images will be mounted by the Hyper-V hosts as Clustered Shared Volumes via iSCSI over dedicated vlans.

**Network Environment.** Each Hyper-V host will utilise a Cisco UCSB-MLOM-40G-01 Virtual Interface Card. The virtual interface card (VIC) will present multiple virtual NIC's to the host, which will be mapped to I/O modules installed within the UCS Chassis. Each UCS 5108 chassis is equipped with two 2204XP I/O Modules (Fabric Extenders). These will have two connections from each I/O module to the upstream 6248UP Fabric Interconnects. The Fabric Interconnects will have upstream connections to the Nexus 5548UP switches that provide connectivity to the core switching infrastructure.

**Microsoft Failover Clusters.** Failover clusters will be deployed for the infrastructure Hyper-V hosts, HSD Hyper-V hosts or HVD Hyper-V hosts. Each failover cluster will be deployed with two separate paths for the cluster heartbeat and a shared network for CSV I/O redirection and Live Migration traffic.

Availability Sets will be used to identify virtual machines that SCVMM will keep on separate hosts for redundancy, e.g. DHCP servers, virtual file server nodes. The Infrastructure Failover cluster will utilise a "Node and Disk Majority" Quorum configuration to ensure only a single node can own 2 of the 3 quorum resources *e.g. node and disk majority*.

In the event the solution is scaled beyond the 1,000 desktops, an additional node will be added to the infrastructure cluster; at this time the disk witness will not be required.

**Active Directory Integration.** Each Hyper-V host will be logically located in a dedicated Organisational Unit with specific Group Policy Objects applied as appropriate to the Hyper-V role.

## 5.5 SMB File Services

### Overview

This Citrix Validated Solution has a dependency for Windows SMB file shares to host various storage components of the design, specifically:

- Provisioning Services vDisk store
- User Personalisation data
- Profile Management and redirected User folders
- ISO and media repository
- Microsoft VMM Library Server

Since the Nimble Storage array only supports block based storage using iSCSI, this design discusses the use of a Microsoft Windows based file server. The file server must be deployed in a high availability mode to provide a level of resiliency.

Nimble Storage provides further guidelines and recommendations as described in this document, <http://info.nimblestorage.com/bpg-windows-file-sharing.html>

This section discusses design requirements and integration points for a Microsoft Server 2012 Failover Cluster running the General Use file server role. *Figure 12.* below describes the high level architecture:

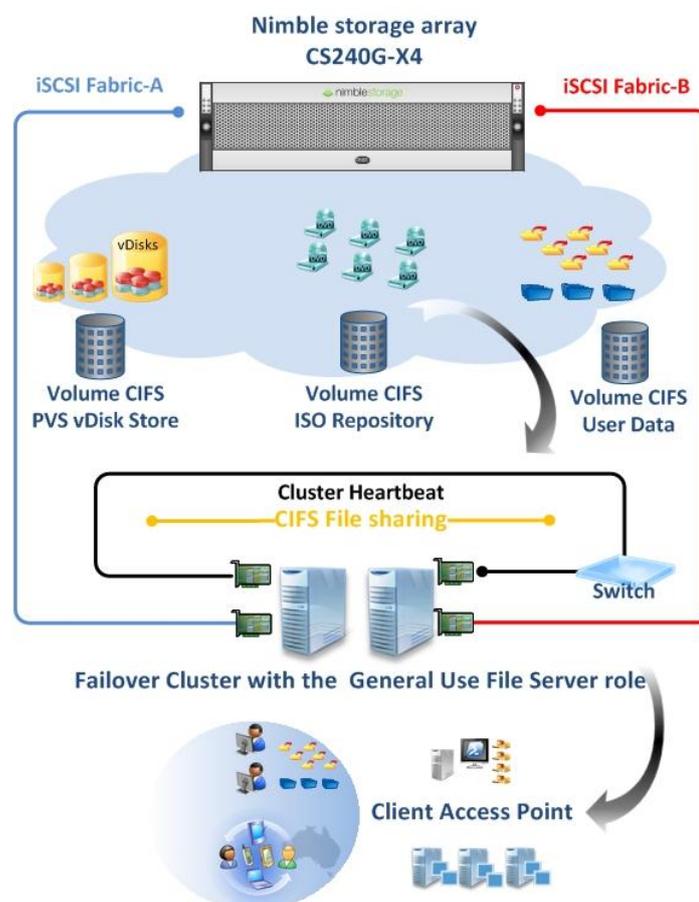


Figure 12. Windows Server 2012 File Server Architecture

## Key Decisions

Configuration	Decision
Highly available File Server solution	2-node Failover Cluster running the General Use file server role with a quorum witness disk.
Storage	iSCSI shared storage.
Volumes	<p><a href="#">Nimble Volumes:</a></p> <ul style="list-style-type: none"> <li>• PVS vDisk Store</li> <li>• Profile Management and redirected User folders</li> <li>• ISO media repository</li> <li>• Disk Quorum</li> </ul>
Client Access Name	<a href="#">\\Infra-cifs01</a>
SMB Shares	<p>PVS-vDisk Store:</p> <ul style="list-style-type: none"> <li>• "<a href="#">\\infra-cifs01\PVS-Store</a>" <ul style="list-style-type: none"> <li>• Path=F:\PVS-Store</li> </ul> </li> </ul> <p>Profile Management and redirected User folders</p> <ul style="list-style-type: none"> <li>• "<a href="#">\\infra-cifs01\UPM</a>" <ul style="list-style-type: none"> <li>• Path=E:\UPM</li> </ul> </li> <li>• "<a href="#">\\infra-cifs01\UserData</a>" <ul style="list-style-type: none"> <li>• Path=E:\UserData</li> </ul> </li> </ul> <p>ISO media repository:</p> <ul style="list-style-type: none"> <li>• "<a href="#">\\infra-cifs01\ISO</a>" <ul style="list-style-type: none"> <li>• Path=G:\ISO</li> </ul> </li> </ul>
Cluster Networks	<p>Cluster Management:</p> <ul style="list-style-type: none"> <li>• Management</li> <li>• Client Access</li> <li>• Cluster Communications (Path 1)</li> </ul> <p>Cluster Communications:</p> <ul style="list-style-type: none"> <li>• Cluster Communications (Path 2)</li> </ul> <p>iSCSI Fabric-A:</p> <ul style="list-style-type: none"> <li>• iSCSI data traffic fabric A (Path 1)</li> </ul> <p>iSCSI Fabric-B:</p> <ul style="list-style-type: none"> <li>• iSCSI data traffic fabric B (Path 2)</li> </ul>
Hardware Settings	<p>There are two options on how the file server should be deployed - either utilising physical or virtualized instances:</p> <p><b>Option1:</b></p> <p>The <a href="#">Hyper-V infrastructure Failover Cluster</a> may be built with the full version of Microsoft Server 2012 Datacenter Edition and the General Use file server role configured alongside the Hyper-V role.</p> <p><b>Option2:</b></p> <p>2 x virtualised Microsoft Server 2012 VM guest cluster nodes within a Hyper-V Failover Cluster running the General Use file server role:</p>

Configuration	Decision
	<ul style="list-style-type: none"> <li>• Hyper-V VM guest</li> <li>• Windows Server 2012 Standard edition</li> <li>• 4 vCPUs</li> <li>• 16GB RAM</li> <li>• 100GB disk for Operating System (C:\)</li> <li>• 1 vNIC for cluster management (Host teamed network)</li> <li>• 1 vNIC for cluster communications (Host teamed network)</li> <li>• 1 vNIC for iSCSI Fabric-A traffic</li> <li>• 1 vNIC for iSCSI Fabric-B traffic</li> </ul> <p>Hypervisor additional requirements to support the virtual cluster (virtual machines instances):  UCS Service profile update: <i>"HyperV_Infra_BootiSCSI"</i>  Four additional vNICs:</p> <ol style="list-style-type: none"> <li>1. HV-VM-iSCSI-A</li> <li>2. HV-VM-iSCSI-B</li> </ol> <p>New host Team <i>Team-VM-ClusterHB</i>, Interface members:</p> <ol style="list-style-type: none"> <li>1. VM-Cluster-A</li> <li>2. VM-Cluster-B</li> </ol>
Failover Clustering	<ul style="list-style-type: none"> <li>• Failover Clustering enabled <ul style="list-style-type: none"> <li>• Node and Disk Majority</li> </ul> </li> </ul>

Table 23. Windows Server 2012 SMB File Services Key Decisions

## Design

This high level design discusses two options for the deployment of a highly available General Use file server cluster presenting SMB shares for different file data sharing purposes.

### Option 1:

- The Hyper-V Failover Cluster "clust-infra001" hosting the infrastructure virtual machines may be configured with Microsoft Server 2012 Datacenter Edition thereby allowing the Failover cluster role "General Use file server" to be deployed. HA will be maintained as per the cluster configuration.
- Each node in the Infrastructure cluster will be granted access to the volumes used for presenting the SMB 3.0 shares. Nimble Storage initiator groups will be used to manage access to the volumes.

### Option 2:

- The Hyper-V Failover Cluster "clust-infra001" hosting the infrastructure virtual machines will host two additional virtual machines configured as a 2-node guest Failover cluster hosting the "General Use file server" role. The operating system deployed to the VMs will be Microsoft Server 2012 Standard Edition. Availability sets will be configured to ensure each of the nodes remain on separate physical hosts for redundancy.

- Each node in the file server cluster will be granted access to the volumes used for presenting the SMB shares. Nimble Storage initiator groups will be used to manage access to the volumes.
- To support the requirements of the virtualised Failover cluster (running as guest VMs) the underlying Hyper-V hosts will require additional Cisco UCS service profile and Hyper-V configuration. These changes are described at a high level in the following tables.

Required Cisco UCS Configuration Changes to support Option 2:

Decision Point	Description / Decision
Service Profile Templates	Service profile template to be amended: Service Profile Infrastructure Hyper-V hosts: <ul style="list-style-type: none"> <li>• HyperV_Infra_BootiSCSI</li> </ul>
Networks	The Cisco Virtual Interface Card 1240 will be configured within the service profile to present the following Networks to each Hyper-V host: <ul style="list-style-type: none"> <li>• Live Migration (CSV I/O Redirection Network)</li> <li>• VM Traffic (multiple vlans)</li> <li>• Host Management (Cluster Heartbeat)</li> <li>• iSCSI storage (IPSAN)</li> <li>• iSCSI for virtual machine storage</li> <li>• Dedicated cluster network for virtual machine clusters</li> </ul>
vNICs	Hyper-V Live Migration network: <ul style="list-style-type: none"> <li>• HV-LiveMig-Fab-A</li> <li>• HV-LiveMig-Fab-B</li> </ul> Hyper-V host management network: <ul style="list-style-type: none"> <li>• HV-MGMT-Fab-A</li> <li>• HV-MGMT-Fab-B</li> </ul> VM Data for HyperV_Infra_BootiSCSI Service Profile: <ul style="list-style-type: none"> <li>• HV-VM-INF-Fab-A</li> <li>• HV-VM-INF-Fab-B</li> </ul> iSCSI Traffic: <ul style="list-style-type: none"> <li>• HV-VM-iSCSI-A</li> <li>• HV-VM-iSCSI-B</li> <li>• HV-iSCSI-Fab-A</li> <li>• HV-iSCSI-Fab-B</li> </ul> VM Cluster Heart Beat <ul style="list-style-type: none"> <li>• FS-VM-CHB-Fab-A</li> <li>• FS-VM-CHB-Fab-B</li> </ul>
vlans	Additional vlan required for the virtual machine cluster heartbeat: <ul style="list-style-type: none"> <li>• vlan ID: 34</li> </ul>

Table 24. Cisco UCS configuration changes to support the guest Failover Cluster

Required Hyper-V Configuration Changes to support Option 2:

Decision Point	Description / Decision
Network Team	New host Team: <i>Team-VM-Cluster</i> , Interface members: <ol style="list-style-type: none"> <li>1. VM-Cluster-A</li> <li>2. VM-Cluster-B</li> </ol>
Hyper-V Switch	New: VM-Switch-ClusterHB <ul style="list-style-type: none"> <li>• Associated Hyper-V interface: Team-VM-Cluster</li> <li>• New vlan ID 34</li> </ul> New VM-Switch-iSCSI-Fabric-A <ul style="list-style-type: none"> <li>• Associated Hyper-V interface: HV-VM-iSCSI-A</li> <li>• Native vlan</li> </ul> New VM-Switch-iSCSI-Fabric-B <ul style="list-style-type: none"> <li>• Associated Hyper-V interface: HV-VM-iSCSI-B</li> <li>• Native vlan</li> </ul>

Table 25. Hyper-V Configuration Changes to support the guest Failover Cluster

## 5.6 Citrix Provisioning Services

### Overview

The Citrix Provisioning Services (PVS) environment is designed as a single farm and one initial Site. A single Site is used to host three Provisioning servers for the proposed workloads for up to two Hosted Shared Desktop pods (up to 2,000 users) or a single Hosted Virtual Desktop pod (up to 1,000 users).

- **DHCP on Windows Server® 2012.** The Citrix Validated Solution uses the DHCP failover feature providing the ability to have two DHCP servers serve IP addresses and option configuration to the same subnet or scope, providing uninterrupted availability of DHCP service to clients. The two DHCP servers will be configured to replicate lease information between themselves, allowing one server to assume responsibility for servicing of clients for the entire subnet when the other server is unavailable, without using split scopes.

Figure 13. below describes the high-level components showing a HSD and a HVD pod in relation to the Citrix Provisioning Services farm and DHCP configuration:

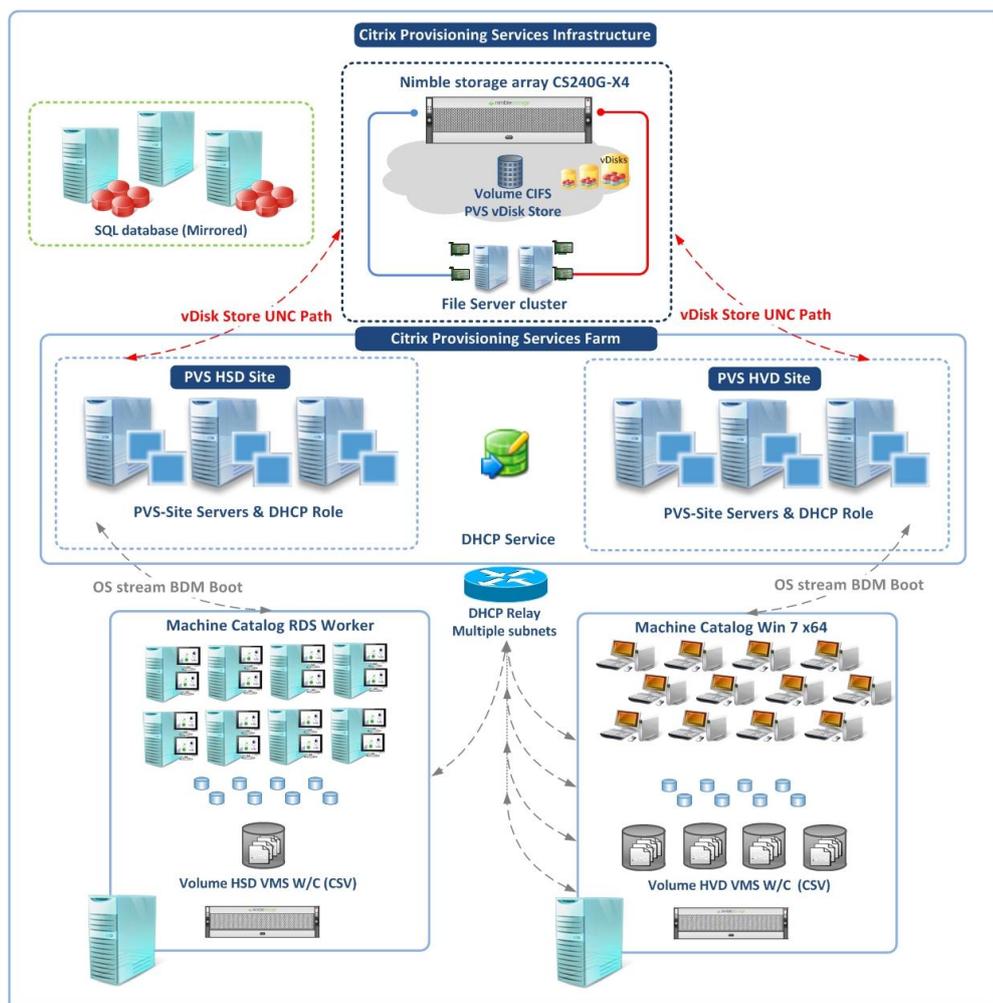


Figure 13. Citrix Provisioning Services Farm and related infrastructure

## Key Decisions: PVS

Configuration	Decision
Version	Citrix Provisioning Services 7.1
Servers	3 x Provisioning Services servers will be deployed, two are required to maintain high availability at all times. The third allows for maintenance of a single server while maintaining high availability.
Boot Services	Boot Device Manager (BDM). BDM vhd Disk attached to each target device. The Boot Devices Manager Utility provides an optional method for providing IP and boot information to target devices; as an alternative to using the PXE, and TFTP methods.
Hardware Settings	<ul style="list-style-type: none"> <li>3 x virtualised PVS servers</li> <li>Hyper-V VM guest</li> <li>Windows Server 2012 Standard Edition</li> <li>4 vCPUs</li> <li>16GB RAM<sup>9</sup> (allows for caching of ~ 4 vDisk images)</li> <li>100GB disk for Operating System (C:\)</li> <li>1 vNIC for production traffic</li> </ul>
Storage Settings	PVS vDisk Store hosted on a Windows Server 2012 file share associated with a volume presented by the Nimble Storage array.
Network Settings	<p>PVS Server:</p> <ul style="list-style-type: none"> <li>1 vNIC for production traffic (Synthetic)</li> </ul> <p>PVS target devices will be multi homed with 2 x vNICs as follows:</p> <ul style="list-style-type: none"> <li>1 vNIC for production traffic (Synthetic)</li> <li>1 vNIC for streaming traffic (Emulated) The PXE Boot option is required to support the BDM drivers during the boot phase only. Once the Synthetic NIC is operational within the VM the PVS software will automatically switch streaming traffic to this NIC</li> </ul>
PVS Write Cache Settings HSD	<p>Local disk on the target device; a 20GB (D: drive) persistent virtual disk will be associated to each target device.</p> <p>Sizing guideline based on the HSD and application workload tested<sup>10</sup>:</p> <ul style="list-style-type: none"> <li>Write cache size after 24 hours of testing is ~2GB x 7 days of uptime = 14 GB + redirected logs = ~15GB with ~25% spare storage capacity</li> </ul>
PVS Write Cache Settings HVD	<p>Local disk on the target device; a 10GB (D: drive) persistent virtual disk will be associated to each target device:</p> <p>Sizing Guideline based on the HVD and application workload tested<sup>11</sup>:</p> <p>Write cache size after 24 hours ~ 512 MB x 4 days of uptime = ~2 GB + logs 1 GB + 5 GB Page file =10 GB with ~25% spare capacity</p>
PVS Farm	Farm Name and details: Refer to the Appendix: <a href="#">DECISION POINT</a>

<sup>9</sup> Recommended minimum memory requirement for Citrix PVS servers. Caters for up to 4 x vDisk images

<sup>10</sup> Sizing guidelines are based on the application set tested as part of the scalability testing conducted within the CVS labs. This value is a guideline and the actual metrics may differ depending on unique customer applications and requirements.

<sup>11</sup> Sizing guidelines are based on the application set tested as part of the scalability testing conducted within the lab. This value is a baseline and the actual metrics may differ depending on unique customer applications and requirements.

Configuration	Decision
Configuration	
Database and Service account information	<p>Mirrored database running on Microsoft SQL Server 2008 R2</p> <p>Refer to the Appendix: <a href="#">DECISION POINT</a></p> <ul style="list-style-type: none"> <li>• Database information</li> <li>• Service Account information <ul style="list-style-type: none"> <li>• Additional permissions: SELF "Write Public Information" Required for automatic updating of the Service Principle Name (SPN)</li> </ul> </li> <li>• Failover over partner information</li> </ul>
PVS vDisk Store	<p>A single store is shared by the 3 x PVS servers within the PVS Site. The vDisk store will be configured to utilise the Windows SMB file share.</p> <p>Path:</p> <ul style="list-style-type: none"> <li>• "<a href="#">\\infra-cifs01\PVS-Store</a>"</li> </ul>
Device Collections	<p>A single device collection created for each XenDesktop Catalog.</p>
vDisk Properties	<ul style="list-style-type: none"> <li>• Access Mode - Standard image</li> <li>• Cache Type - Cache on Target device hard drive (VM vhd file residing on iSCSI shared storage)</li> <li>• Enable Active Directory machine account password management – Enabled</li> <li>• Microsoft Volume Licensing: Refer to the Appendix: <a href="#">DECISION POINT</a></li> </ul>
Hypervisor integration	<ul style="list-style-type: none"> <li>• VMM 2012 spk1 console installed on all PVS servers</li> </ul>

Table 26. Citrix Provisioning Services Key Decisions

### **Key Decisions: DHCP**

Configuration	Decision
Version, Edition	Windows Server® 2012 DHCP Role enabled
Servers	Two Windows Server® 2012 PVS servers will be deployed with the DHCP role enabled
(IPv4 Options) Failover	Failover Enabled

Table 27. DHCP Scope Key Decisions

## Design

### PVS Design

- **PVS Farm.** The Citrix Provisioning Services (PVS) environment is designed as a single farm and one initial Site for the appropriate FlexCast model. A single Site is used to host three Provisioning servers for each FlexCast workload, two servers are required to maintain high availability at all times. The third allows for maintenance of a single server while still maintaining high availability. A windows SMB file share will be used for the storage of vDisks.
- **PVS Target Device Network.** The Hyper-V Legacy Network adapter will be used in conjunction with the Synthetic NIC. The Legacy NIC is required to allow the Boot Device Manager Disk option ROM to stay resident in memory during the PXE boot phase. The Legacy NIC will be configured on the same subnet as the Synthetic (production) NIC leveraging automatic changeover of the Legacy NIC to the Synthetic NIC during the boot process.
- **IP Addressing.** All target devices will receive IP addressing requirements from DHCP both the Synthetic and the Legacy NIC require IP Address information for the same subnet to support the NIC change over feature. Once the Synthetic NIC is operational within the VM the PVS software will automatically switch streaming traffic to the Production NIC (Synthetic).
- **PVS Farm Database.** The Farm database will be hosted on a Microsoft SQL 2008 R2 platform using synchronous database mirroring.

### DHCP Design

- **DHCP.** Two PVS servers will host Microsoft DHCP Services for the IP addressing requirements. DHCP Relay will be configured on the Cisco Nexus 5548UP switches, allowing client DHCP discover packets to be forwarded to their respective DHCP servers. DHCP scopes will be deployed as highly available in load balanced mode, using the capabilities of the DHCP Role.

### Active Directory Design

- **Active Directory Integration.** Each server will be logically located in a dedicated Organisational Unit with specific Group Policy Objects applied as appropriate to the PVS and DHCP server role.

## 5.7 Citrix XenDesktop

### Overview

This validated solution specifically defines two FlexCast delivery models, from a XenDesktop perspective each desktop will belong to a catalog specifically for that FlexCast delivery type.

Both HSD and HVD desktops will contain pre-installed core applications (Tier-1 applications) delivered to the user from within that desktop. The Hosted Shared Desktop will be configured with Themes such that the desktop has a Windows 7 “look and feel”.

Figure 14. below identifies the high-level components of the XenDesktop Site describing a HSD and a HVD Catalog:

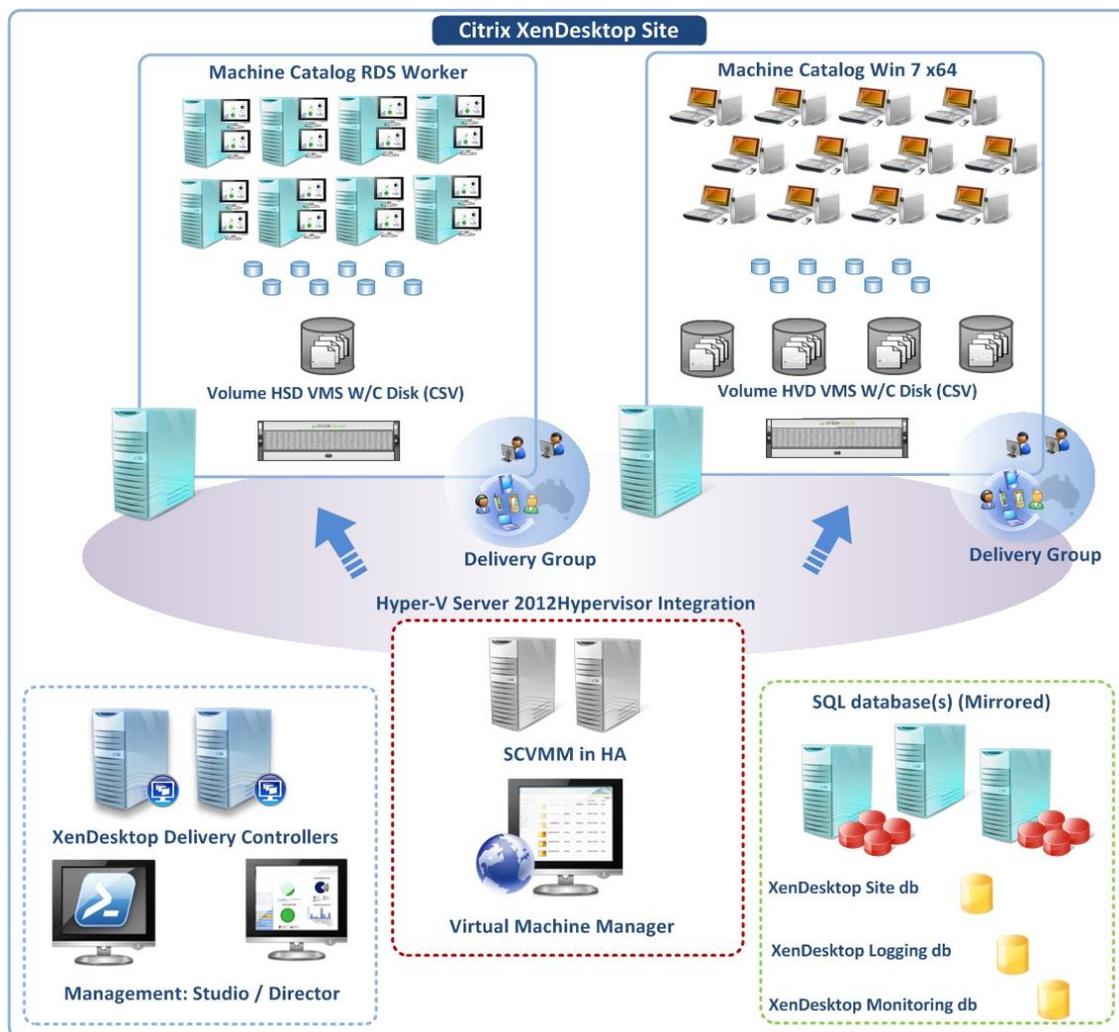


Figure 14. Citrix XenDesktop Site and related infrastructure

## Key Decisions

Configuration	Decision
Version, Edition	Citrix XenDesktop 7.1
Hardware Settings Delivery Controllers	Virtualised XenDesktop Delivery Controller servers : <ul style="list-style-type: none"> <li>• Hyper-V VM guest</li> <li>• Windows Server 2012 Standard Edition</li> <li>• 4 x vCPUs</li> <li>• 8GB RAM</li> <li>• 100GB disk for Operating System (C:\)</li> <li>• 1 x vNIC (Production Traffic)</li> </ul>
Catalog HSD	<ul style="list-style-type: none"> <li>• Windows Server OS</li> <li>• Desktop Images managed by PVS</li> </ul>
Catalog HVD	<ul style="list-style-type: none"> <li>• Windows Desktop OS</li> <li>• Virtual Machines</li> <li>• Provisioning Services</li> <li>• Random Pooled</li> </ul>
Databases	Mirrored database(s) running on Microsoft SQL Server 2008 R2: Refer to the Appendix: <a href="#">DECISION POINT</a> <ul style="list-style-type: none"> <li>• Database information <ul style="list-style-type: none"> <li>• Site database</li> <li>• Configuration Logging database:</li> <li>• Monitoring database:</li> </ul> </li> <li>• Service Account information</li> <li>• Failover over partner information</li> </ul>
Site database	Please refer to the following articles for full details: <a href="http://support.citrix.com/article/CTX139508">http://support.citrix.com/article/CTX139508</a>
Monitoring database	Database retention period 90 days = default for Platinum license Please refer to the following articles for full details: <a href="http://support.citrix.com/article/CTX139508">http://support.citrix.com/article/CTX139508</a>
Logging database	Database retention is manual only no retention policy in place Please refer to the following articles for full details: <a href="http://support.citrix.com/article/CTX139508">http://support.citrix.com/article/CTX139508</a>
Microsoft Remote Desktop Services licensing	Refer to the Appendix: <a href="#">DECISION POINT</a> <ul style="list-style-type: none"> <li>• Microsoft RDS Licensing is required for Hosted Shared Desktop types and will be based on the customer's Microsoft licensing type and model</li> </ul>
Datacentre(s)	The prescribed deployment is for a single data centre.
Citrix Policies	Citrix Policy Application: <ul style="list-style-type: none"> <li>• Applied using Active Directory Group Policy</li> </ul>
Host Connections	1 per vlan / storage Clustered Shared Volume <ul style="list-style-type: none"> <li>• For HSD: <ul style="list-style-type: none"> <li>• Cluster HSD <i>Name</i></li> <li>• vlan hsd_1</li> </ul> </li> </ul>

Configuration	Decision
	<ul style="list-style-type: none"> <li>• Shared: CSVolume1</li> <li>• For HVD: (1 per vlan) <ul style="list-style-type: none"> <li>• Cluster HVD <i>Name</i></li> <li>• vlan hvd_1 vlan, hvd_2, vlan hvd_3, vlan hvd_4</li> <li>• Shared: CSVolume1</li> </ul> </li> </ul>
Hypervisor integration	<ul style="list-style-type: none"> <li>• VMM 2012 spk1 console installed on all Delivery Controller servers</li> </ul>

Table 28. Citrix XenDesktop Key Decisions

## Design

**XenDesktop Site.** The XenDesktop site / farm consist of two virtualised Desktop Delivery Controllers. A host connection will be defined that establishes a connection to the VMM server and the Hyper-V Failover Cluster using a specified service account. Storage connections will be defined to the appropriate Cluster, VLAN and Clustered Shared Volumes.

Each FlexCast desktop type will be configured as a single catalog. HVD desktops will be configured as pooled and randomly assigned.

**Desktop Presentation.** StoreFront will be utilised for the presentation of desktops to end users. StoreFront servers that provide the required application and desktop presentation will be load balanced with Citrix NetScaler.

**Desktop Director and EdgeSight.** Citrix EdgeSight is now integrated into a single console within Desktop Director its feature set is enabled based on Citrix Licensing. The monitoring database used by EdgeSight will be separated from the site and logging database to allow appropriate management and scalability of the database. Historical data retention is available for 90 days by default with platinum licensing. Administrators can select specific views delegating permissions concisely for helpdesk staff allowing easy troubleshooting and faster resolution of problems. Citrix EdgeSight provides two key components:

- **Performance Management.** EdgeSight provides the historical retention with reporting capabilities while Director provides real time views.
- **Network Analysis.** NetScaler HDX Insight is the network component of EdgeSight, providing network analysis information for LAN, WAN and mobile users. Please refer to the [Citrix NetScaler](#) section for more details.

**Active Directory Integration.** Each Machine object will be logical located in a dedicated Organisational Unit with specific Group Policy Objects applied as appropriate to the their role.

## 5.8 Virtual Desktop VM Guest Workloads

### Overview

Each of the virtual desktops represent true-to-production configuration consisting of applications that are pre-installed as part of the virtual desktop “Disk image”. A number of configuration settings will be applied directly to the vDisk image using Active Directory Group Policies ensuring optimal performance. Antivirus will be included with specific configurations as documented within this section

Each of the virtual desktops, Windows 7 or an RDS workload will be deployed using Citrix Provisioning Services standard vDisk mode (read-only, many to one). A number of configuration settings will be applied directly to the each gold image using Active Directory Group Policies, ensuring optimal performance and consistent application.

Aside from Applications a number of components were included in the gold image that may influence scalability:

- Antivirus with specific configurations as documented within this section.  
<http://support.citrix.com/article/CTX127030>
- Themes, Windows 7 look and feel for HSD workloads.

### Key Decisions

#### Hosted Shared Desktop Workload

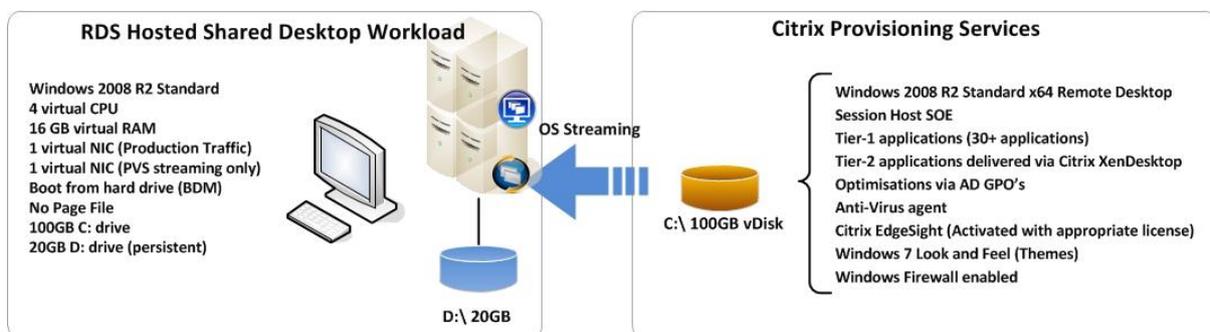


Figure 15. Hosted Shared Desktop Configuration

#### HSD Virtual Machine Specifications:

Based on the system testing carried out the following table describes the most optimal configuration for HSD on Windows Sever 2008 R2 RDS workload for user/session density:

# of VMs per host	#RAM	#vCPU	HSD Sessions per VM	Total # of HSD Sessions per Host
6	16GB	4	~22	~130

Table 29. HSD on Windows Server 2008 R2 RDS VM Specification

Configuration	Decision
Virtual Machine Specifications	Specifications: <ul style="list-style-type: none"> <li>Persistent Drive 20GB (the requirement for a Pagefile further defines the size of this drive. Refer to the Appendix: <a href="#">DECISION POINT</a>)</li> <li>System Drive 100GB (PVS vDisk)</li> </ul>
Pagefile	Refer to the Appendix: <a href="#">DECISION POINT</a> <ul style="list-style-type: none"> <li>Not required based on the application set and workload carried out during validation testing.</li> </ul>
Network Interface(s)	<ul style="list-style-type: none"> <li>NIC1 - Legacy Hyper-V NIC for boot traffic (BDM)</li> <li>NIC2 - Synthetic NIC for streaming and production traffic</li> </ul>
Memory	16GB
vCPU	4 vCPU

Table 30. HSD on Windows Server 2008 R2 RDS VM Specification

### Hosted Virtual Desktop Workload

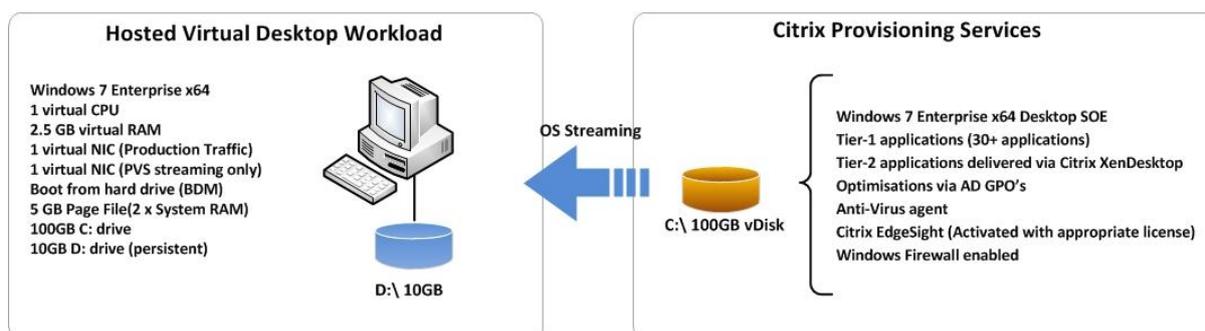


Figure 16. Hosted Virtual Desktop Configuration

### HVD Virtual Machine Specifications:

Based on the system testing carried out the following table describes the most optimal configuration for the Windows 7 workload for user/VM density:

# of VMs per host	#RAM	#vCPU
~ 108-110	2.5GB	1 <sup>12</sup>

Table 31. HVD on Windows 7 VM Specification

<sup>12</sup> Based on the [application workloads](#) described in this document a single vCPU was successfully validated and demonstrated the highest HVD VM to host density throughout testing of this Citrix Validated Solution. For completeness the same scenarios were tested using 2 vCPU in the Guest HVD VM, this configuration demonstrated improved session latency and approximately 10-15% decrease in VM to host density.

Configuration	Decision
Virtual Machine Specifications	Specifications: <ul style="list-style-type: none"> <li>• Persistent Drive 10GB</li> <li>• System Drive 100GB (PVS vDisk)</li> </ul>
Pagefile	<ul style="list-style-type: none"> <li>• 5 GB (2x the size of the assigned memory)</li> </ul>
Network Interface(s)	<ul style="list-style-type: none"> <li>• NIC1 - Legacy Hyper-V NIC for boot traffic (BDM)</li> <li>• NIC2 - Synthetic NIC for streaming and production traffic</li> </ul>
Memory	2.5GB
vCPU	1 vCPU <sup>13</sup>

Table 32. HVD on Windows 7 VM Specification

<sup>13</sup> Based on the [application workloads](#) described in this document a single vCPU was successfully validated and demonstrated the highest HVD VM to host density throughout testing of this Citrix Validated Solution. For completeness the same scenarios were tested using 2 vCPU in the Guest HVD VM, this configuration demonstrated improved session latency and approximately 10-15% decrease in VM to host density.

**Application Set.** The testing utilised application sets representative of enterprise-level SOE applications. These applications will be embedded as part of the “gold image”. The following table represents the application set that form the desktop workload profile:

HSD Application Set:

Configuration	Decision
HSD Operating System	<ul style="list-style-type: none"> <li>• Microsoft Windows Server 2008 R2 Standard Edition with Service Pack 1</li> <li>• Hyper-V Integration Services 6.2.9200.16433</li> </ul>
Citrix Applications	<ul style="list-style-type: none"> <li>• Citrix Virtual Delivery Agent 7.1.0.4033</li> <li>• Citrix Profile Management v5.1</li> <li>• Citrix Provisioning Services Target Device x64 7.1.0.4022</li> <li>• Citrix ShareFile Desktop Widget v2.22</li> <li>• Citrix Receiver v14.1.0.0</li> </ul>
Productivity Applications	<ul style="list-style-type: none"> <li>• Microsoft Excel Professional 2010 x86</li> <li>• Microsoft Outlook Professional 2010 x86</li> <li>• Microsoft PowerPoint Professional 2010 x86</li> <li>• Microsoft Word Professional 2010 x86</li> </ul>
Baseline Applications	<ul style="list-style-type: none"> <li>• Adobe Acrobat Reader v9.1</li> <li>• Adobe Flash Player v11.7.700.202</li> <li>• Adobe Shockwave Player v11.6.636</li> <li>• Adobe AIR v3.7.0.1860</li> <li>• Apple QuickTime v7.72.80.56</li> <li>• Bullzip PDF Printer v7.2.0.1304<sup>14</sup></li> <li>• Cisco WebEx Recorder/Player v 3.23.2516</li> <li>• Google Chrome v31.0.1650.57</li> <li>• Java 6 Update 21 v6.0.210</li> <li>• Kid-Key-Lock v1.2.1<sup>15</sup></li> <li>• Mozilla Firefox v14.0.1</li> <li>• Microsoft .NET Framework 4 Client Profile v4.0.30319</li> <li>• Microsoft Internet Explorer 9</li> <li>• Microsoft System Center EndPoint Protection 2012</li> <li>• Microsoft Silverlight v5.1.20913.0</li> <li>• Microsoft Windows Firewall</li> <li>• Microsoft Windows Media Player v12.x</li> <li>• Skype v5.10.116</li> <li>• WinZip v16.5.10095</li> </ul>

Table 33.HSD Pre-defined Application Set

<sup>14</sup> Application required by Login VSI for scalability testing.

<sup>15</sup> Application required by Login VSI for scalability testing.

HVD Application Set:

Configuration	Decision
HVD Operating System	<ul style="list-style-type: none"> <li>• Microsoft Windows 7 Professional Service Pack 1 x64</li> <li>• Hyper-V Integration Services 6.2.9200.16433</li> </ul>
Citrix Applications	<ul style="list-style-type: none"> <li>• Citrix Virtual Delivery Agent 7.1.0.4033</li> <li>• Citrix Profile Management v5.1</li> <li>• Citrix Provisioning Services Target Device x64 7.1.0.4022</li> <li>• Citrix ShareFile Desktop Widget v2.22</li> <li>• Citrix Receiver v14.1.0.0</li> </ul>
Productivity Applications	<ul style="list-style-type: none"> <li>• Microsoft Excel Professional 2010 x86</li> <li>• Microsoft Outlook Professional 2010 x86</li> <li>• Microsoft PowerPoint Professional 2010 x86</li> <li>• Microsoft Word Professional 2010 x86</li> </ul>
Baseline Applications	<ul style="list-style-type: none"> <li>• Adobe Acrobat Reader v9.1</li> <li>• Adobe Flash Player v11.7.700.202</li> <li>• Adobe Shockwave Player v11.6.636</li> <li>• Adobe AIR v3.7.0.1860</li> <li>• Apple QuickTime v7.72.80.56</li> <li>• Bullzip PDF Printer v7.2.0.1304<sup>16</sup></li> <li>• Cisco WebEx Recorder/Player v 3.23.2516</li> <li>• Google Chrome v31.0.1650.57</li> <li>• Java 6 Update 21 v6.0.210</li> <li>• Kid-Key-Lock v1.2.1<sup>17</sup></li> <li>• Mozilla Firefox v14.0.1</li> <li>• Microsoft .NET Framework 4 Client Profile v4.0.30319</li> <li>• Microsoft Internet Explorer 9</li> <li>• Microsoft System Center EndPoint Protection 2012</li> <li>• Microsoft Silverlight v5.1.20913.0</li> <li>• Microsoft Windows Firewall</li> <li>• Microsoft Windows Media Player v12.x</li> <li>• Skype v5.10.116</li> <li>• WinZip v16.5.10095</li> </ul>

Table 34. HVD Pre-defined Application Set

<sup>16</sup> Application required by Login VSI for scalability testing.

<sup>17</sup> Application required by Login VSI for scalability testing.

## Design

The virtual workloads are deployed using Citrix Provisioning Services. Citrix Provisioning Services utilises a read only virtual disk (vDisk) referred to as standard mode (read-only mode used during Production mode). This vDisk can also be switched to private mode (writable mode used under Maintenance mode) when updates are required to the base image. Each time updates are applied to this image in Maintenance mode, the image must be generalised to ensure it is ready to be deployed in its optimal form to many target devices.

Standard mode images are unique in that they are restored to the original state at each reboot, deleting any newly written or modified data. In this scenario, certain processes are no longer efficient and optimisation of this image is required.

Optimisations and configurations can be applied at several levels:

- **Workload Configuration Gold Image.** Changes are made directly to the gold image. These changes are considered inappropriate to be applied using GPOs or are required settings prior to generalising the image. The image must be generalised whilst it is in writable mode (Private or Maintenance mode). Once the image has been generalised it is immediately shutdown and reverted to a read-only mode (Production or Test mode) and is ready for many to one (many target devices to one vDisk image) deployment.
- **Workload Configuration GPO.** These changes are applied via Active Directory GPO and are considered baseline configurations required in almost all instances. Typical use cases for this GPO are Event log redirection, Citrix Profile Management configuration and target device optimisations. In addition this GPO may have Loopback processing enabled allowing user based settings to be applied to the virtual desktop Organisation Unit level.
- **User Optimisations GPO.** This Active Directory GPO contains optimisations for the user operations within the virtual desktop environment. User configurations cannot typically be deployed as part of the image and are independent. Typical use cases for this GPO are folder redirection and user specific optimisations.

## 5.9 Citrix StoreFront

### Overview

Citrix Storefront provides a unified application and desktop aggregation point, including Windows, web, SaaS and mobile applications. Users can access their resources through a standard Web browser using Citrix Receiver.

### Key Decisions

Configuration	Decision
Version, Edition	StoreFront Version 2.1.0.17
Hardware Settings	2 x StoreFront servers in High Availability: <ul style="list-style-type: none"><li>• Hyper-V VM guest</li><li>• Windows Server 2012 Standard Edition</li><li>• 2 vCPUs</li><li>• 4GB RAM</li><li>• 100GB disk for Operating System (C:\)</li><li>• 1 vNIC</li></ul>
Security	A server certificate will be installed to secure authentication traffic: <ul style="list-style-type: none"><li>• https will be required for all web sites, ensuring that user's credentials are encrypted as they traverse the network</li></ul>
Load Balancing	Citrix NetScaler will be deployed to perform server load balancing and health checking of the StoreFront web services.

Table 35. Citrix StoreFront Key Decisions

### Design

Citrix StoreFront servers will be load balanced using Citrix NetScaler SDX 11500 appliances with virtual instances configured in high availability mode (HA). Citrix specific service monitors will be utilised to monitor the health of the StoreFront services to ensure intelligent load balancing decisions are performed on the service. Please refer to the section [Citrix NetScaler SDX](#) for more details

**Active Directory Integration.** Each server will be logically located in a dedicated Organisational Unit with specific Group Policy Objects applied as appropriate to the web server role.

## 5.10 Citrix License Server

### Overview

The Citrix License server is a required server component which provides license service requirements to the Citrix products included in this document.

### Key Decisions

Configuration	Decision
Version, Edition	Citrix License Service version 11.11.1.
Hardware Settings	1 x virtualised License Server: <ul style="list-style-type: none"><li>• Hyper-V VM guest:</li><li>• Windows Server 2012 Standard Edition</li><li>• 2 vCPUs</li><li>• 4GB RAM</li><li>• 100GB disk for Operating System (C:\)</li><li>• 1 vNIC</li></ul>

Table 36. Citrix License Server Key Decisions

### Design

**Redundancy.** Redundancy is built into the Citrix License service via the built-in 30 day grace period. Service redundancy can be further facilitated by the underlying hypervisor; therefore a single server is recommended.

**Active Directory Integration.** The License server will be logically located in a dedicated Organisational Unit with specific Group Policy Objects applied as appropriate to the web server role.

## 5.11 Citrix NetScaler SDX

### Overview

This section provides a high-level description of the proposed Citrix NetScaler SDX functionality and Access Gateway features (Secure Remote Access functionality) required.

Figure 17. depicts the proposed Citrix NetScaler SDX logical architecture for a single data centre integrated to a XenDesktop site managing both Hosted Shared and Hosted Virtual desktop pods:

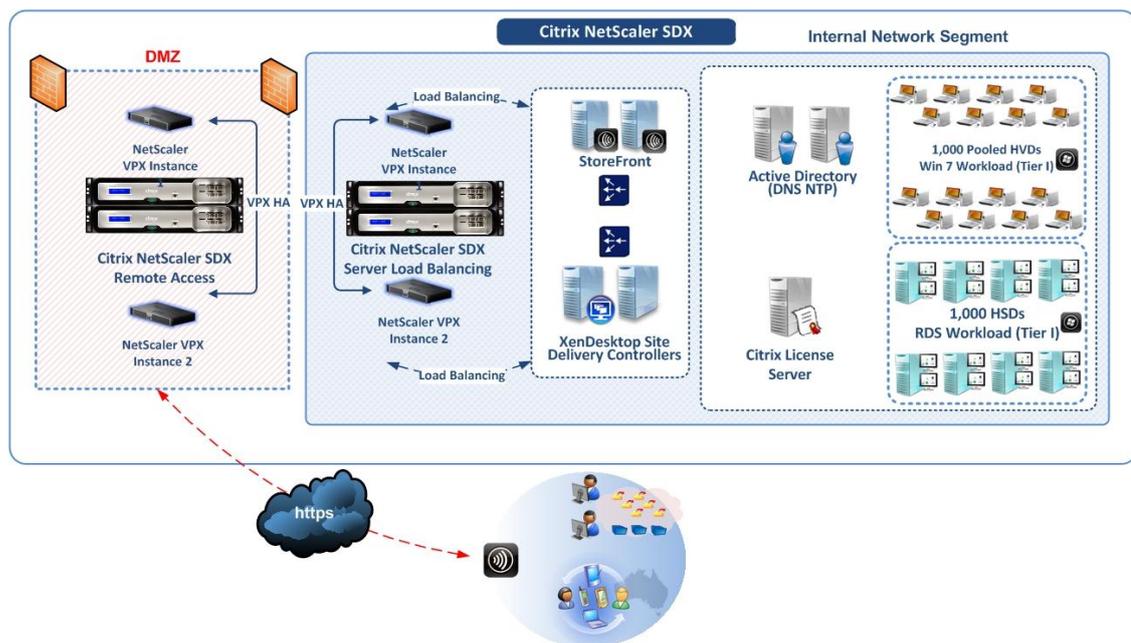


Figure 17. Citrix NetScaler SDX High Level Overview

### Key Decisions

Item	Decision
Appliance Type	Citrix NetScaler SDX 11500
NetScaler Configuration	<p>Four Citrix NetScaler SDX appliances are required for the following:</p> <p>Appliance 1 &amp; 2:</p> <ul style="list-style-type: none"> <li>2 x appliances within the secured DMZ providing remote access capability. Separate VPX instances will be created and configured in high availability between physical appliances.</li> </ul> <p>Appliance 3 &amp; 4:</p> <ul style="list-style-type: none"> <li>2 x appliances within the internal network segment providing load balancing capabilities. Separate VPX instances will be created and configured in high availability between physical appliances to support load balancing of: <ul style="list-style-type: none"> <li>Citrix StoreFront</li> <li>Citrix XenDesktop Delivery Controllers XML Brokers</li> </ul> </li> </ul>

Item	Decision
Citrix Access Gateway	A single Access scenario will be created for Citrix Receiver
Server Load Balancing	Load Balancing of the following Citrix services: <ul style="list-style-type: none"> <li>• Citrix StoreFront</li> <li>• Citrix XenDesktop XML Broker</li> </ul>
Citrix NetScaler Insight Center	Recommended: please refer to the design section <a href="#">below</a> for further details
Deployment	Single data centre.
Global Server Load Balancing (GSLB)	GSLB directs DNS requests to the best performing GSLB site in a distributed Internet environment. GSLB enables distribution of traffic across multiple sites / data centres, and ensures that applications or desktops are consistently accessible. When a client sends a DNS request, the system determines the best performing site and returns its IP address to the client. <ul style="list-style-type: none"> <li>• <a href="#">DECISION POINT</a></li> </ul>

Table 37. Citrix NetScaler SDX Key Decisions

## Design

Two pairs of Citrix NetScaler appliances will be deployed in the appropriate network security zones and network segments. Each physical appliance will be configured with a single instance (initially) to support high availability between the physical appliances.

- **External facing (DMZ Network).** The two NetScaler SDX appliances will be configured such that each virtual NetScaler instance will be in two arm mode. A single SSL VPN virtual server will be created to support a single access scenario providing access to the virtual desktops using a standard web browser and Citrix Receiver.
- **Internal facing (Internal Network).** The two NetScaler SDX appliances will be configured such that each virtual NetScaler instance will provide load balancing capabilities to internal Web sites. Load balancing will be provided for:
  - Citrix StoreFront servers / Sites
  - Citrix XML Brokers
- **Citrix NetScaler Insight Center.** Although not validated within this design Citrix NetScaler Insight Center should be considered as part of the deployment. NetScaler Insight Center. Is deployed as a virtual appliance that collects and provides detailed information regarding web traffic and virtual desktop traffic passing through a NetScaler appliance. There are two components of Insight Center as described below:
  - **HDX Insight.** Allows monitoring of ICA traffic passing through the NetScaler virtual servers defined on the appliance. HDX Insight provides the ability to monitor Citrix XenApp and Citrix XenDesktop environments, monitoring users and performance of hosted applications and desktops. HDX Insight integrated into XenDesktop Director provides Network analysis and advanced monitoring features.
  - **Web Insight.** Allows monitoring of HTTP traffic (web-application traffic) passing through load balancing and content switching virtual servers defined on the NetScaler appliance.

## 5.12 User Profile Management Solution

### Overview

Profile management is enabled through a Windows service that provides a mechanism for capturing and managing user personalisation settings within the virtual desktop environment. Citrix Profile Management is installed by default during the installation of the Virtual Desktop agent.

### Key Decisions

Configuration	Decision
Version, Edition	Citrix User Profile Management version 5.1
Profile Storage Location	Windows SMB Share - " <a href="#">\\infra-cifs01\UPM</a> "
Folder redirection	Applied using Group Policy: (minimum requirements): <ul style="list-style-type: none"><li>• Application Data</li><li>• Documents</li></ul> Redirected folder location: <ul style="list-style-type: none"><li>• Windows SMB file share<ul style="list-style-type: none"><li>• "<a href="#">\\infra-cifs01\UserData</a>"</li></ul></li></ul> Refer to the Appendix: <a href="#">DECISION POINT</a>
Configuration	Profile Management configurations will be applied using Active Directory GPOs.

Table 38. Citrix Profile Management Key Decisions

### Design

Citrix Profile Management coupled with standard Microsoft Windows Folder Redirection using Active Directory GPOs will be deployed. Storage presented via a Windows SMB file share provided by the Nimble Storage array will host the User Profiles and User Redirected Folders. All Profile Management configurations will be deployed using Active Directory GPOs.

## 5.13 Active Directory

### Overview

This validated solution has a requirement to use Microsoft Active Directory Domain Services and as such, it is an assumption that such an environment already exists within the customer's environment. The decisions discussed below describe requirements from the existing Active Directory in the form of Organisational Units and Group Policy Objects.

Supplementary requirements must also be met, to ensure sufficient capacity from authenticating Domain Controllers and can handle any additional requirements or load placed on the system by adding further Users, Groups, machine Objects and policy processing load.

[DECISION POINT](#)

### Key Decisions

Configuration	Decision
Group Policy Application	<p>Recommended:<sup>18</sup></p> <ul style="list-style-type: none"><li>• Each infrastructure server role will have a minimum security baseline applied (MSB) via GPO</li><li>• All RDS workloads will have a minimum security baseline applied (MSB) via GPO</li><li>• Windows 7 workloads will have a minimum security baseline applied (MSB) via GPO</li><li>• RDS workloads will have a Machine GPO applied specific to their application delivery requirements. This GPO will have Loopback mode enabled to apply user based settings at the RDS workload OU level</li><li>• Windows 7 workloads will have a Machine GPO applied specific to their application delivery requirements. This GPO may have Loopback mode enabled to apply user based settings at the machine workload OU level</li><li>• User based policies may be applied at the machine level using the loopback mode</li><li>• Infrastructure servers such as Hyper-V host will be deployed in relevant OUs and MSBs applied appropriate to their role.</li></ul>

Table 39. Active Directory Key Decisions

### Design

The recommended Group Policy and Organisational Unit strategy applied to this validated solution is based on deploying Group Policy Objects in a functional approach, e.g. settings are applied based on service, security or other functional role criteria. This ensures that security settings targeted for specific role services such as IIS, SQL etc. receive only their relevant configurations.

It is anticipated that the final design will be customer dependant and based on other factors such as role based administration and other typical elements outside the scope of this document. Refer to the Appendix: [DECISION POINT](#)

---

<sup>18</sup> Reference to Minimum Security Baselines in the form of GPOs will be the customer's responsibility. GPOs described in this document in all cases will be integrated into the customer Active Directory environment.



Figure 18. Organisational Units and GPO Application

## 5.14 Database Platform

### Overview

Citrix XenDesktop, Citrix Provisioning Services and Virtual Machine Manager require databases to store configuration metadata and statistical information. A highly available database platform utilising Microsoft SQL Server is required as the platform of choice.

The following tables describe minimum requirements of the database platform.

### Key Decisions

Configuration	Decision
Version, Edition	<p>Microsoft SQL Server 2008 R2 Standard edition (Used at the time of testing)</p> <p>Please refer to the following article for a list of supported database platforms:  <a href="http://support.citrix.com/servlet/KbServlet/download/18493-102-706969/Database%20Chart.pdf">http://support.citrix.com/servlet/KbServlet/download/18493-102-706969/Database%20Chart.pdf</a></p>
Databases	<p><b>XenDesktop 7.1 databases:</b></p> <ul style="list-style-type: none"> <li>• Mirrored: Synchronous mirroring with Witness node</li> </ul> <p>Please refer to the following articles for details on database sizing:  <a href="http://support.citrix.com/article/CTX139508">http://support.citrix.com/article/CTX139508</a></p> <p>Please refer to the following article for database fault tolerance:  <a href="http://support.citrix.com/proddocs/topic/xendesktop-71/cds-plan-high-avail-rho.html">http://support.citrix.com/proddocs/topic/xendesktop-71/cds-plan-high-avail-rho.html</a></p> <p><b>Provisioning Services:</b></p> <ul style="list-style-type: none"> <li>• Mirrored: Synchronous mirroring with Witness node</li> </ul> <p>Please refer to the following article for further details:  <ul style="list-style-type: none"> <li>• <a href="http://support.citrix.com/proddocs/topic/provisioning-60/pvs-install-task1-plan-6-0.html">http://support.citrix.com/proddocs/topic/provisioning-60/pvs-install-task1-plan-6-0.html</a></li> </ul> </p> <p><b>Microsoft VMM:</b></p> <p>Please refer to the following article for further details:  <ul style="list-style-type: none"> <li>• <a href="http://technet.microsoft.com/en-us/library/gg610574.aspx">http://technet.microsoft.com/en-us/library/gg610574.aspx</a></li> <li>• <a href="http://technet.microsoft.com/en-us/sqlserver/gg490638.aspx">http://technet.microsoft.com/en-us/sqlserver/gg490638.aspx</a></li> </ul> </p>

Table 40. Microsoft SQL Database Key Decisions

### Design Considerations

This document provides design guidelines for the actual databases used in this Citrix Validated Solution, however does not attempt to provide design guidelines for Microsoft SQL Server. The design and implementation for a highly available Microsoft SQL Server platform is required although considered out of scope for this high level design document.

## Appendix A. Decision Points

This section defines the elements which need further discussions with the Customer as these may be customer-specific.

DECISION POINT	Description
Naming Convention	<ul style="list-style-type: none"> <li>Component nomenclature will need to be defined by the customer during the Analysis phase of the project</li> </ul>
Database Information	<ul style="list-style-type: none"> <li>Microsoft SQL Version</li> <li>Server name</li> <li>Instance name</li> <li>Port</li> <li>Database name</li> <li>Resource Capacity (CPU Memory Storage)</li> </ul>
CTX Licensing	<ul style="list-style-type: none"> <li>License server name</li> </ul>
Microsoft Volume Licensing	<p>Microsoft licensing of the target devices is a requirement for the solution and will be based on the customer's existing Microsoft licensing agreement. The appropriate licensing option must be selected based on Microsoft KMS or MAK volume licenses for PVS target devices.</p> <ul style="list-style-type: none"> <li>Note: The vDisk license mode must be set before target devices can be activated</li> </ul>
Microsoft RDS Licensing (Terminal Server CALS)	<p>At least two Microsoft RDS License servers should be defined when using RDS workloads within the customer environment including the mode of operation:</p> <ul style="list-style-type: none"> <li>per user</li> <li>per device</li> </ul> <p>Once defined these configuration items will be deployed via the Active Directory GPO.</p>
Windows Pagefile	<p>The final applications used and workload usage patterns required by the customer will influence the decision for the requirements and sizing of the Windows Pagefile. Further customer validation will be required. Dependant on the sizing of the Pagefile and its associated storage footprint, the write cache drive may require additional storage considerations.</p>
User Logon	<p>Further analysis may be required for customers with aggressive user logon time frames to their desktops. In this scenario additional resources may be required. This may impact Citrix StoreFront, host Density or other related infrastructure.</p>
Active Directory Domain services	<p>The Active Directory Forest and domain will need to be discussed with the Customer to ensure sufficient capacity exists to support any additional authentication requirements the proposed solution may impose.</p> <p>Group Policy is likely to be deployed to suit the requirements of the customer. Assuming the existing deployment meets best practices, the GPOs described within this Citrix Validated Solution can be integrated into the customer environment or configurations may be added directly to existing GPOs. Further analysis is required. Reference to Minimum Security Baselines in the form of GPOs will be the customer's</p>

DECISION POINT	Description
	responsibility. GPOs described in this document in all cases must be integrated into the customer Active Directory environment.
User Personalisation	<p>User Profile Management will need to be further defined to meet customer expectations and application specific requirements. This includes folder redirection using GPO objects. Currently this document only describes minimal requirements, that were used for testing and validation purposes</p> <p>Please refer to the following link for further details:  <a href="http://support.citrix.com/article/CTX134081">http://support.citrix.com/article/CTX134081</a></p>

Table 41. Decision Points

## Appendix B. Server Inventory

This section defines the inventory of servers (physical and virtual) required to deliver the 1,000 pod virtual desktop solution. The following tables describe the requirements for the:

- 1,000 user Hosted Shared Desktop Pod
- 1,000 user Hosted Virtual Desktop Pod

Note: if deploying two full pods (2 x 1,000 users) or part thereof some infrastructure components may be shared between pods (e.g. Citrix StoreFront servers or Citrix Delivery Controllers). This is likely to reduce the requirement for separate infrastructure hosts for both pods and requires further consideration.

### Hosted Shared Desktops

Qty	Server role	Type	CPU	RAM	Disk	NIC
<b>Physical Servers</b>						
2	Hyper-V Host (Infrastructure)	Physical – B200-M3	2 x Hex-Core	128GB	SAN Boot – 150GB	VIC1240
8	Hyper-V Host (HSD)	Physical – B200-M3	2 x Hex-Core	128GB	SAN Boot – 150GB	VIC1240
<b>Virtual Servers</b>						
2	Citrix Desktop Delivery Controllers servers	VM	4 vCPU	8GB	100GB	1 vNIC
2	Citrix StoreFront server	VM	2 vCPU	4GB	100GB	1 vNIC
3	Citrix Provisioning servers	VM	4 vCPU	16GB	100GB	1 vNICs
1	Citrix License server	VM	2 vCPU	4GB	100GB	1 vNIC
48	RDS Workload (HSD) servers	VM	4 vCPU	16GB	100GB (PVS) 20GB (W/C)	2 vNICs
2	Virtual Machine Manager	VM	4 vCPU	16GB	100GB	1 vNIC
<b>Virtual Servers (Failover Cluster for General Use file shares)</b>						
2	File Server Cluster Nodes	VM	4 vCPU	16GB	100GB	4 vNICs

Table 42. Server Inventory for a HSD pod of 1,000 desktops/sessions

## Hosted Virtual Desktops

Qty	Server role	Type	CPU	RAM	Disk	NIC
<b>Physical Servers</b>						
2	Hyper-V Host (Infrastructure)	Physical – B200-M3	2 x Hex-Core	128GB	SAN Boot – 150GB	VIC1240
10	Hyper-V Host (HSD)	Physical – B200-M3	2 x Hex-Core	320GB	SAN Boot – 150GB	VIC1240
<b>Virtual Servers</b>						
2	Citrix Desktop Delivery Controllers servers	VM	4 vCPU	8GB	100GB	1 vNIC
2	Citrix StoreFront server	VM	2 vCPU	4GB	100GB	1 vNIC
3	Citrix Provisioning servers	VM	4 vCPU	16GB	100GB	1 vNICs
1	Citrix License server	VM	2 vCPU	4GB	100GB	1 vNIC
1,000	Windows 7 workload (HVD)	VM	1 vCPU <sup>19</sup>	2.5GB	100GB (PVS) 10GB (W/C)	2 vNICs
2	Virtual Machine Manager	VM	4 vCPU	16GB	100GB	1 vNIC
<b>Virtual Servers (Failover Cluster for General Use file shares)</b>						
2	File Server Cluster Nodes	VM	4 vCPU	16GB	100GB	4 vNICs

Table 43. Server Inventory for a HVD pod of 1,000 desktops/sessions

<sup>19</sup> Based on the [application workloads](#) described in this document a single vCPU was successfully validated and demonstrated the highest HVD VM to host density throughout testing of this Citrix Validated Solution. For completeness the same scenarios were tested using 2 vCPU in the Guest HVD VM, this configuration demonstrated improved session latency and approximately 10-15% decrease in VM to host density.

The copyright in this report and all other works of authorship and all developments made, conceived, created, discovered, invented or reduced to practice in the performance of work during this engagement are and shall remain the sole and absolute property of Citrix, subject to a worldwide, non-exclusive license to you for your internal distribution and use as intended hereunder. No license to Citrix products is granted herein. Citrix products must be licensed separately. Citrix warrants that the services have been performed in a professional and workman-like manner using generally accepted industry standards and practices. Your exclusive remedy for breach of this warranty shall be timely re-performance of the work by Citrix such that the warranty is met. THE WARRANTY ABOVE IS EXCLUSIVE AND IS IN LIEU OF ALL OTHER WARRANTIES, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE SERVICES OR PRODUCTS PROVIDED UNDER THIS AGREEMENT, THE PERFORMANCE OF MATERIALS OR PROCESSES DEVELOPED OR PROVIDED UNDER THIS AGREEMENT, OR AS TO THE RESULTS WHICH MAY BE OBTAINED THEREFROM, AND ALL IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR AGAINST INFRINGEMENT. Citrix' liability to you with respect to any services rendered shall be limited to the amount actually paid by you. IN NO EVENT SHALL EITHER PARTY BE LIABLE TO THE OTHER PARTY HEREUNDER FOR ANY INCIDENTAL, CONSEQUENTIAL, INDIRECT OR PUNITIVE DAMAGES (INCLUDING BUT NOT LIMITED TO LOST PROFITS) REGARDLESS OF WHETHER SUCH LIABILITY IS BASED ON BREACH OF CONTRACT, TORT, OR STRICT LIABILITY. Disputes regarding this engagement shall be governed by the internal laws of the State of Florida.



Level 3, 1 Julius Avenue

North Ryde, Sydney 2113 02-8870-0800

<http://www.citrix.com>

Copyright © 2012 Citrix Systems, Inc. All rights reserved. Citrix, the Citrix logo, Citrix ICA, Citrix MetaFrame, and other Citrix product names are trademarks of Citrix Systems, Inc. All other product names, company names, marks, logos, and symbols are trademarks of their respective owners.