



Dell Wyse Datacenter for VMware Horizon DaaS Reference Architecture

A Reference Architecture for the design, configuration and implementation of a VMware Horizon DaaS environment.

Dell Wyse Solutions Engineering
November 2014

Revisions

Date	Description
June 2014	Initial release v.1.0
August 2014	Added on-premise and hybrid architecture models v.1.1
November 2014	13G updates v.1.2



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2014 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. QLogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

Revisions.....	2
1 Introduction.....	7
1.1 Purpose of this document.....	7
1.2 Scope.....	7
1.3 New in this release.....	7
2 Reference Architecture overview.....	8
2.1 Introduction.....	8
2.2 Physical architecture overview.....	8
2.2.1 Horizon DaaS Service Provider hosted model.....	9
2.2.2 Horizon DaaS hybrid on-premise compute hosted model.....	10
2.2.3 Horizon DaaS fully on-premise model.....	11
2.3 Network layers.....	12
2.4 Compute layers.....	13
2.5 Storage layers.....	14
2.6 Other services and appliances.....	15
3 Architecture components.....	16
3.1 Component overview.....	16
3.2 Shared storage – Rack.....	17
3.2.1 Shared Storage – Rack – 1 Tenant (600 Enhanced users example).....	17
3.2.2 Shared storage – Rack – 1 to 200 Tenants (3000 Enhanced users example).....	18
4 Hardware components.....	21
4.1 Networking.....	21
4.1.1 Force10 S55 (ToR switch).....	21
4.1.2 Force10 S60 (1Gb ToR switch).....	22
4.1.3 Force10 S4810 (10Gb ToR switch).....	23
4.2 Servers.....	25
4.2.1 PowerEdge R730.....	25
4.3 Storage.....	26
4.3.1 SolidFire SF-3010 Tier 1 storage (iSCSI).....	26
4.4 Wyse Cloud Clients.....	27
4.4.1 Wyse 5020-P25.....	27



4.4.2	Wyse 5012-D10DP	27
4.4.3	Wyse 7020-P45	27
4.4.4	Wyse 7250-Z50D	28
4.4.5	Wyse 7290-Z90D7	28
4.4.6	Chromebook 11	28
5	Software components	29
5.1	VMware Horizon DaaS	29
5.2	VMware vSphere 5 hypervisor	30
6	Solution architecture.....	31
6.1	Compute server infrastructure	31
6.2	Management server infrastructure.....	32
6.2.1	SQL databases	32
6.2.2	DNS.....	32
6.2.3	DNS for SQL.....	33
6.3	Scaling guidance	33
6.3.1	Management/Desktop Server Host Sizing.....	34
6.3.2	Desktop Host and Appliance scale out matrices	35
6.4	Storage architecture overview.....	37
6.4.1	Shared Tier 1 storage.....	37
6.4.2	Storage networking – SolidFire 3010 iSCSI.....	37
6.5	Virtual networking.....	38
6.5.1	Shared Tier 1 – iSCSI	38
6.6	Solution high availability	41
6.6.1	vSphere HA.....	42
6.6.2	Horizon DaaS infrastructure protection.....	43
6.6.3	Management server high availability	43
6.6.4	Windows File Services high availability.....	44
6.6.5	SQL Server high availability	44
7	Solution performance and testing.....	45
7.1	Load generation and monitoring	45
7.1.1	VMware View Planner.....	45
7.1.2	Login VSI – Login Consultants	45



7.1.3	Liquidware Labs Stratusphere UX.....	45
7.1.4	VMware vCenter.....	46
7.2	Performance analysis methodology.....	47
7.2.1	Resource utilization.....	47
7.2.2	EUE tools information.....	48
7.2.3	EUE real user information.....	48
7.2.4	Workload and profile definitions.....	48
7.2.5	Profile specifications.....	48
7.2.6	Workload specifications.....	49
7.3	Testing and validation.....	50
7.3.1	Testing process.....	50
7.4	VMware Horizon DaaS test results.....	51
A	Appendix.....	53
A.1	Microsoft operating systems and licensing.....	53
A.1.1	Server operating systems and the Services Provider License Agreement.....	53
A.1.2	Desktop-as-a-Service using Windows Remote Desktop Services.....	53
	Acknowledgements.....	54
	About the authors.....	54



1 Introduction

1.1 Purpose of this document

This purpose of this document is to describe:

- How the Dell Wyse Datacenter for VMware Horizon DaaS Reference Architecture scales from 50 to 50,000+ virtual desktop infrastructure (VDI) users.
- Solution options that encompass a combination of models including SolidFire storage options.

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via VMware Horizon DaaS.

1.2 Scope

Relative to delivering the virtual desktop environment, the objectives of this document are to:

- Define the detailed technical design for the solution.
- Define the hardware requirements to support the design.
- Define the design constraints which are relevant to the design.
- Define relevant risks, issues, assumptions and concessions – referencing existing examples, wherever possible.
- Provide a breakdown of the design into key elements so that the reader receives a modular explanation of the design.
- Provide solution scaling and component selection guidance.

1.3 New in this release

- Silver Peak WAN optimization validation - <http://dell.to/UkKAb0>
- SolidFire storage validation - <http://dell.to/1im2o0f>



2 Reference Architecture overview

2.1 Introduction

This solution leverages a core set of hardware and software components consisting of 4 primary layers:

- Networking
- Compute server
- Management server
- Storage

These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. Additionally, the solution includes an approved extended list of optional components in the same categories. These components give IT departments the flexibility to custom tailor the solution for environments with unique VDI feature, scale or performance needs. The solution stack is designed to be a cost effective starting point for IT departments looking to gradually migrate to a fully virtualized desktop environment. This approach allows you to grow the investment and commitment as needed or as your IT staff becomes more comfortable with DaaS VDI technologies.

A summary of infrastructure components may include:

- Network: VLANs
- Compute resources: Management and virtual desktop hosts
- Storage
- Service provider network services
- Service provider appliances
- Resource manager appliances
- Tenant appliances
- Tenant network appliances
- Tenant network services

Offering a Desktop-as-a-Service (DaaS) solution requires appropriately licensed solutions for each end customer and end user. The Dell Wyse Datacenter for VMware Horizon DaaS Reference Architecture is focused on the use of Microsoft Windows Server as a desktop operating system that has been licensed under the Microsoft Services Provider License Agreement (SPLA). Please review the appendix of this document for additional information and resources related to the Microsoft SPLA. This Reference Architecture does not preclude the use of desktop operating systems such as Windows 7 or Windows 8.

2.2 Physical architecture overview

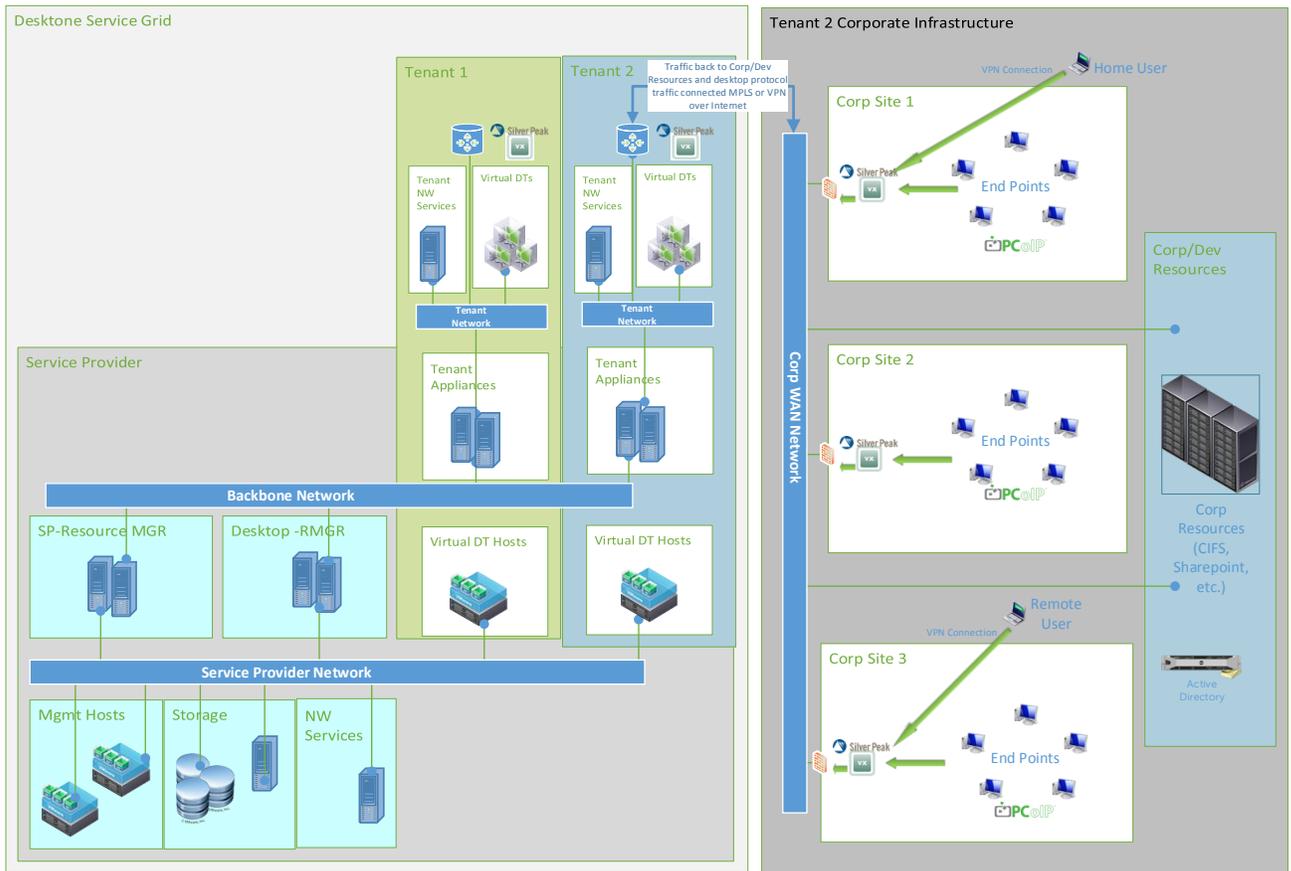
The core Dell Wyse Datacenter architecture is built on a Shared Tier 1 model. Tier 1 in the Dell Wyse Datacenter context defines from which disk source the VDI sessions execute. Tier 2 storage is present in the solution architectures and, while having a reduced performance requirement, is utilized for user profile/data and Management virtual machine (VM) execution. Management VM execution occurs using



Tier 2 storage for all solution models. Tier 2 storage can be in the form of an in-tenant infrastructure VM supporting CIFS or run on same platform as Tier1 storage utilizing components such as deduplication. Dell Wyse Datacenter is a 100% virtualized solution architecture.

2.2.1 Horizon DaaS Service Provider hosted model

The following illustration shows an overview of the VMware Horizon DaaS Architecture when fully hosted in the Service Provider's cloud. It includes WAN optimization from Silver Peak but other similar functioning technology can be used. The Corporate end of the DaaS deployment is an example of how it might be architected but a variety of connectivity options exist in this realm. In this model, the Service Provider's customers consume the desktops entirely in the cloud at the Service Provider's data center. This model is good for regular desktop use cases or where increased security does not apply.

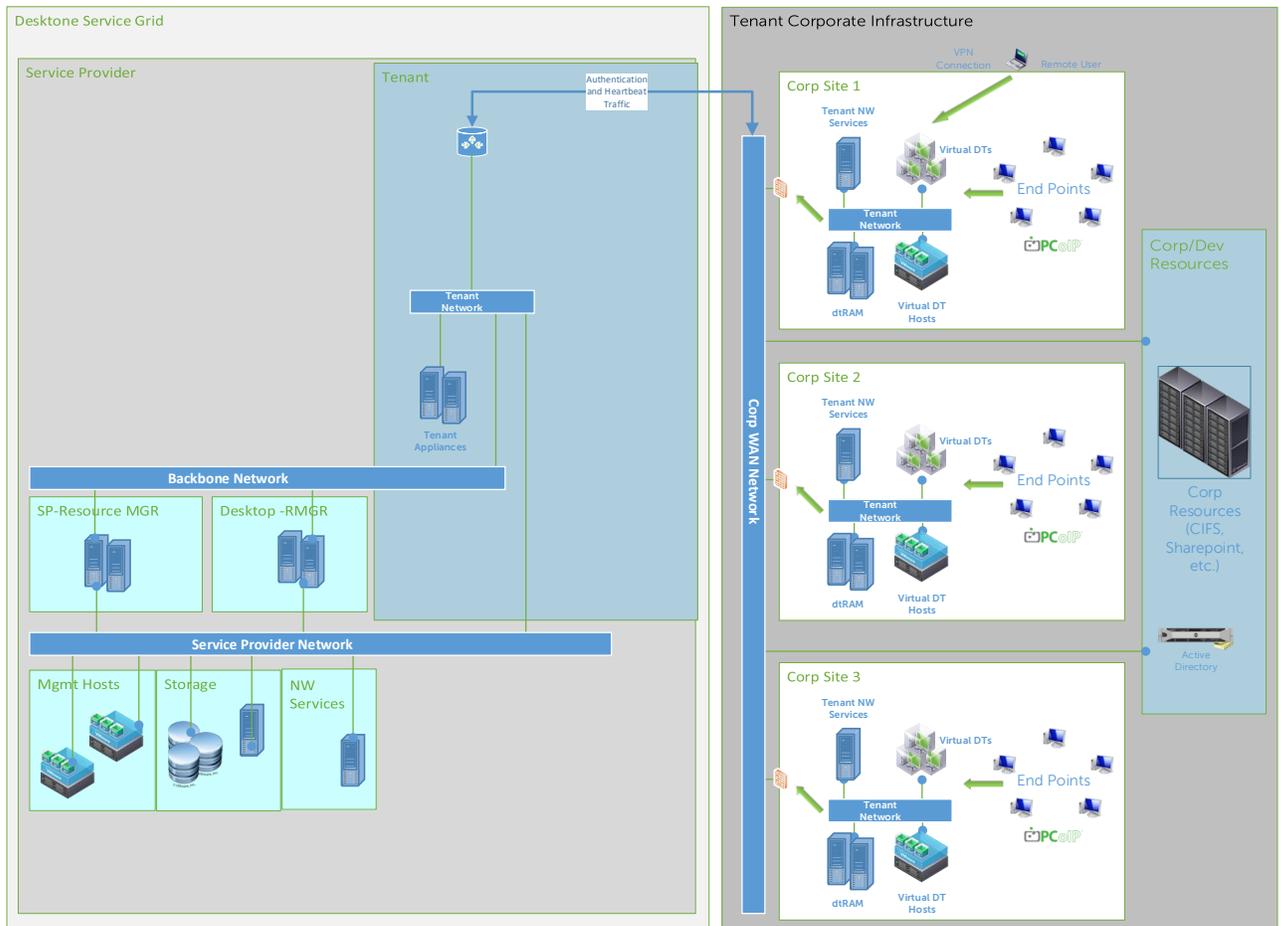


See the larger illustration in PDF format [here](#)



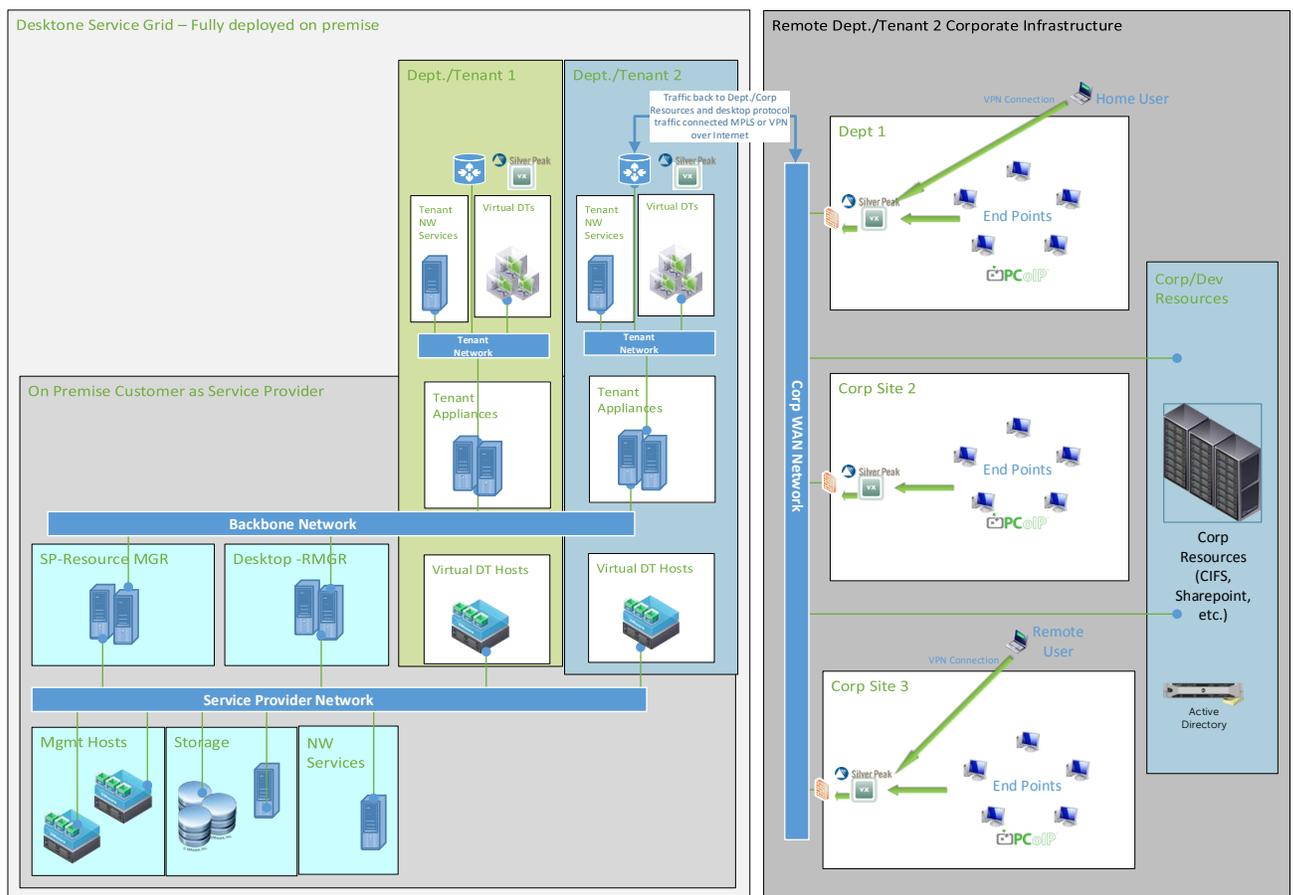
2.2.2 Horizon DaaS hybrid on-premise compute hosted model

The following illustration shows an overview of the VMware Horizon DaaS Architecture where a hybrid model is employed. All management functions are hosted in the Service Provider's cloud while compute, storage resources and some Horizon DaaS tenant infrastructure are hosted at the customer's site. In this model the Service Provider's customers consume the desktops locally on-premise at customer's infrastructure. This is a good architectural model for desktop use cases that require vSGA graphics with NVIDIA K1 or K2 cards installed in the on-premise compute servers or where increased security models are required.

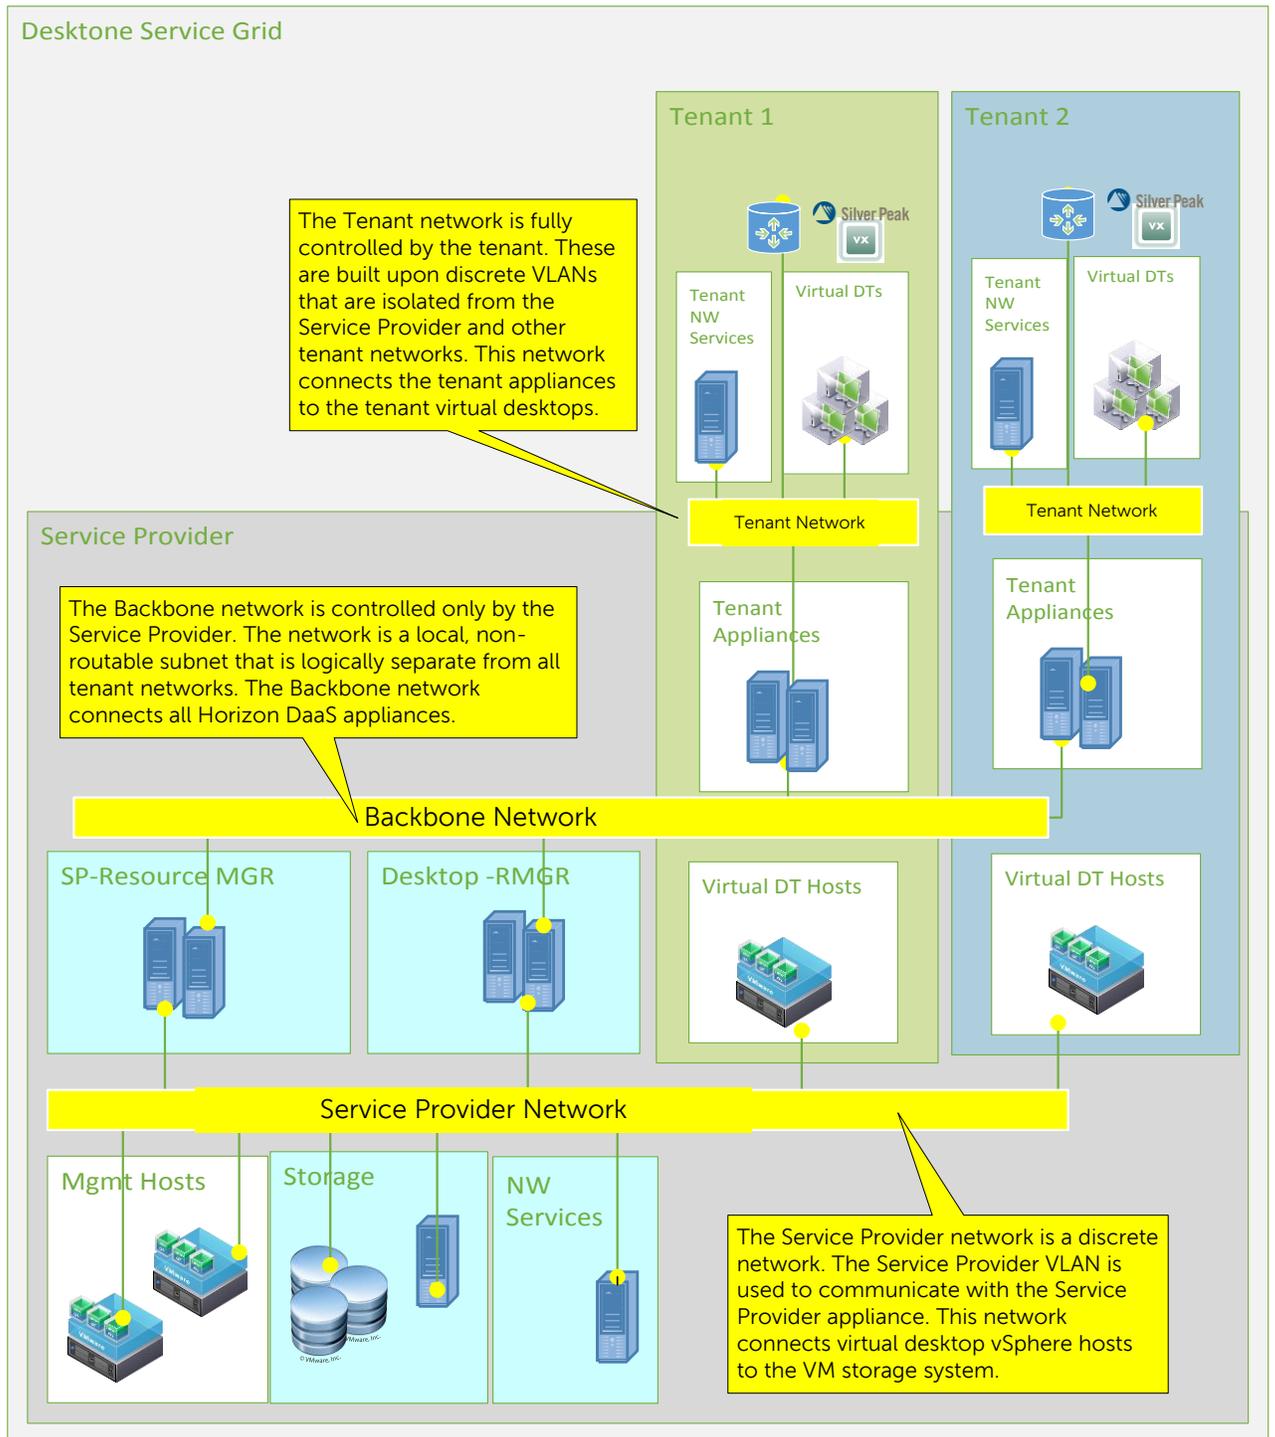


2.2.3 Horizon DaaS fully on-premise model

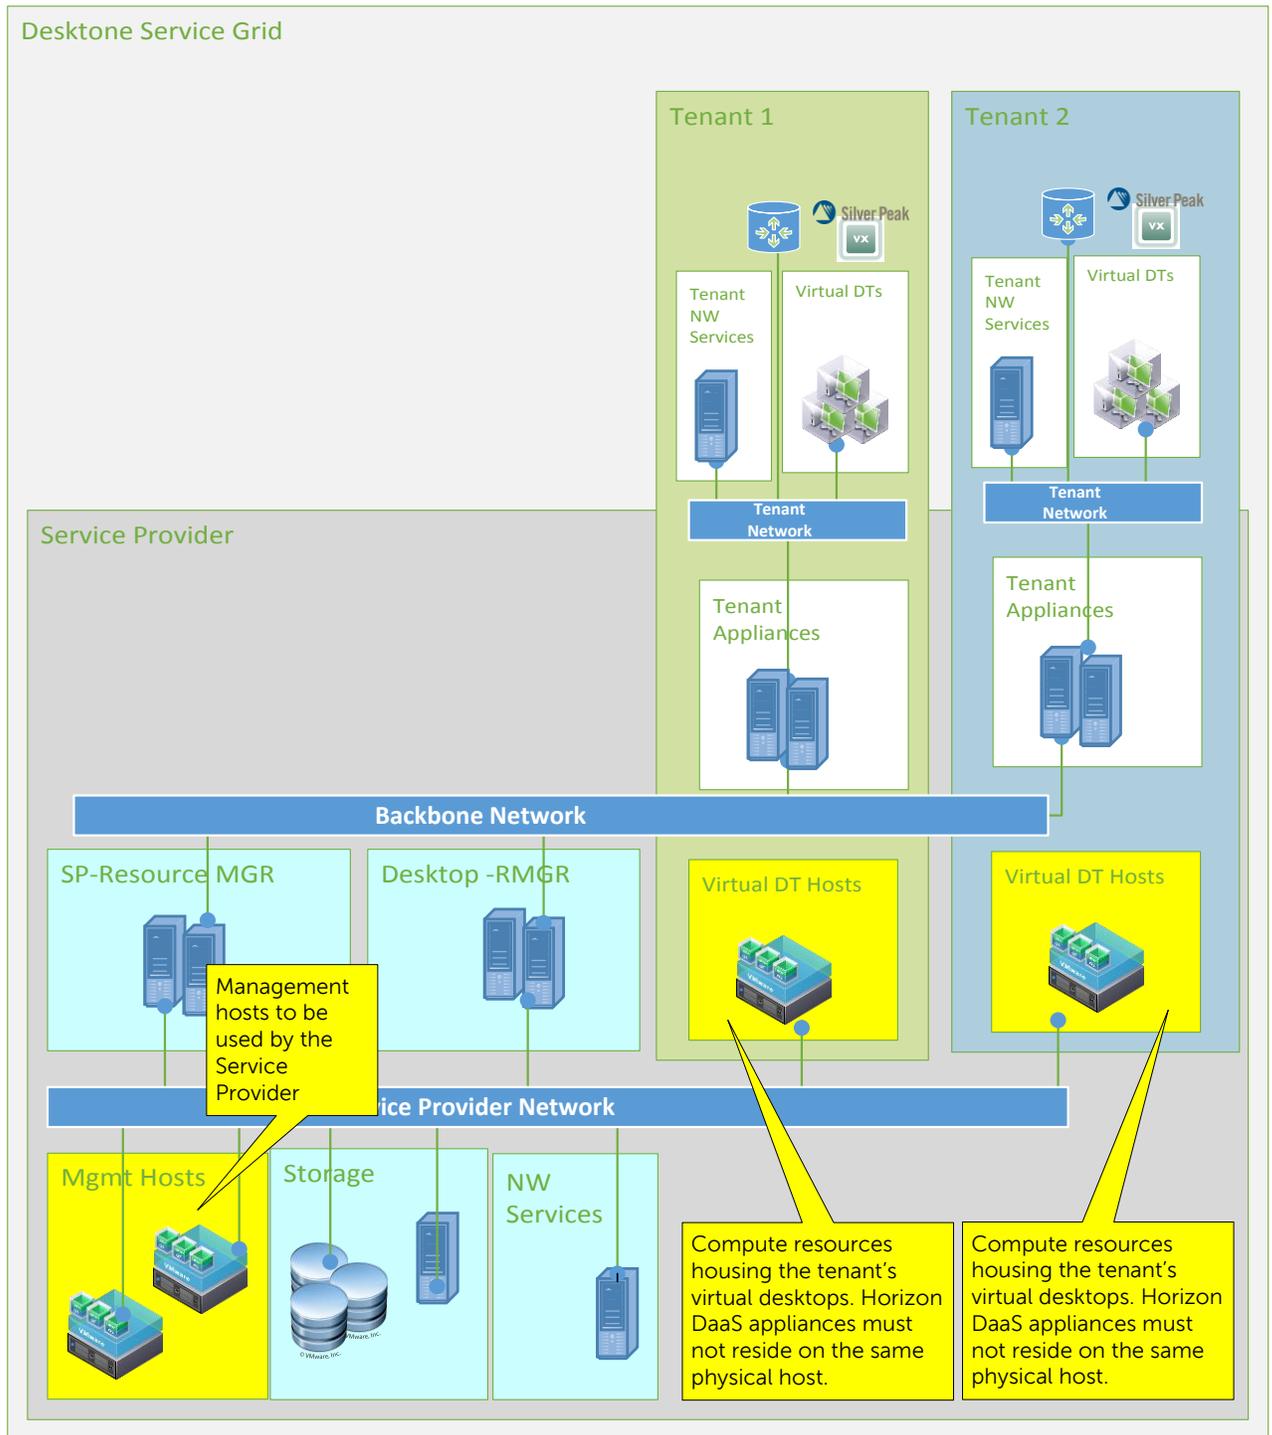
The following illustration shows an overview of the VMware Horizon DaaS Architecture where a Horizon DaaS Infrastructure is fully deployed on-premise at a customer site. All management functions and compute resources are hosted locally at the customer's site. In this model the Service Provider's customers consume the desktops locally on-premise at customer's infrastructure. Similar to the customer hosted compute model, this model works well for desktop use cases that require vSGA graphics with NVIDIA K1 or K2 cards installed in the on-premise compute servers or where the highest security models are required. In addition, some customers have "sub-tenants" (internal user groups) which they support but need to treat as separate and distinct customers even though they may be part of the same entity. Licensing structures for this model are different. The architecture is identical to the Service Provider hosted model above except that all Horizon DaaS infrastructure is installed at the customer site and not in the Service Provider's data center.



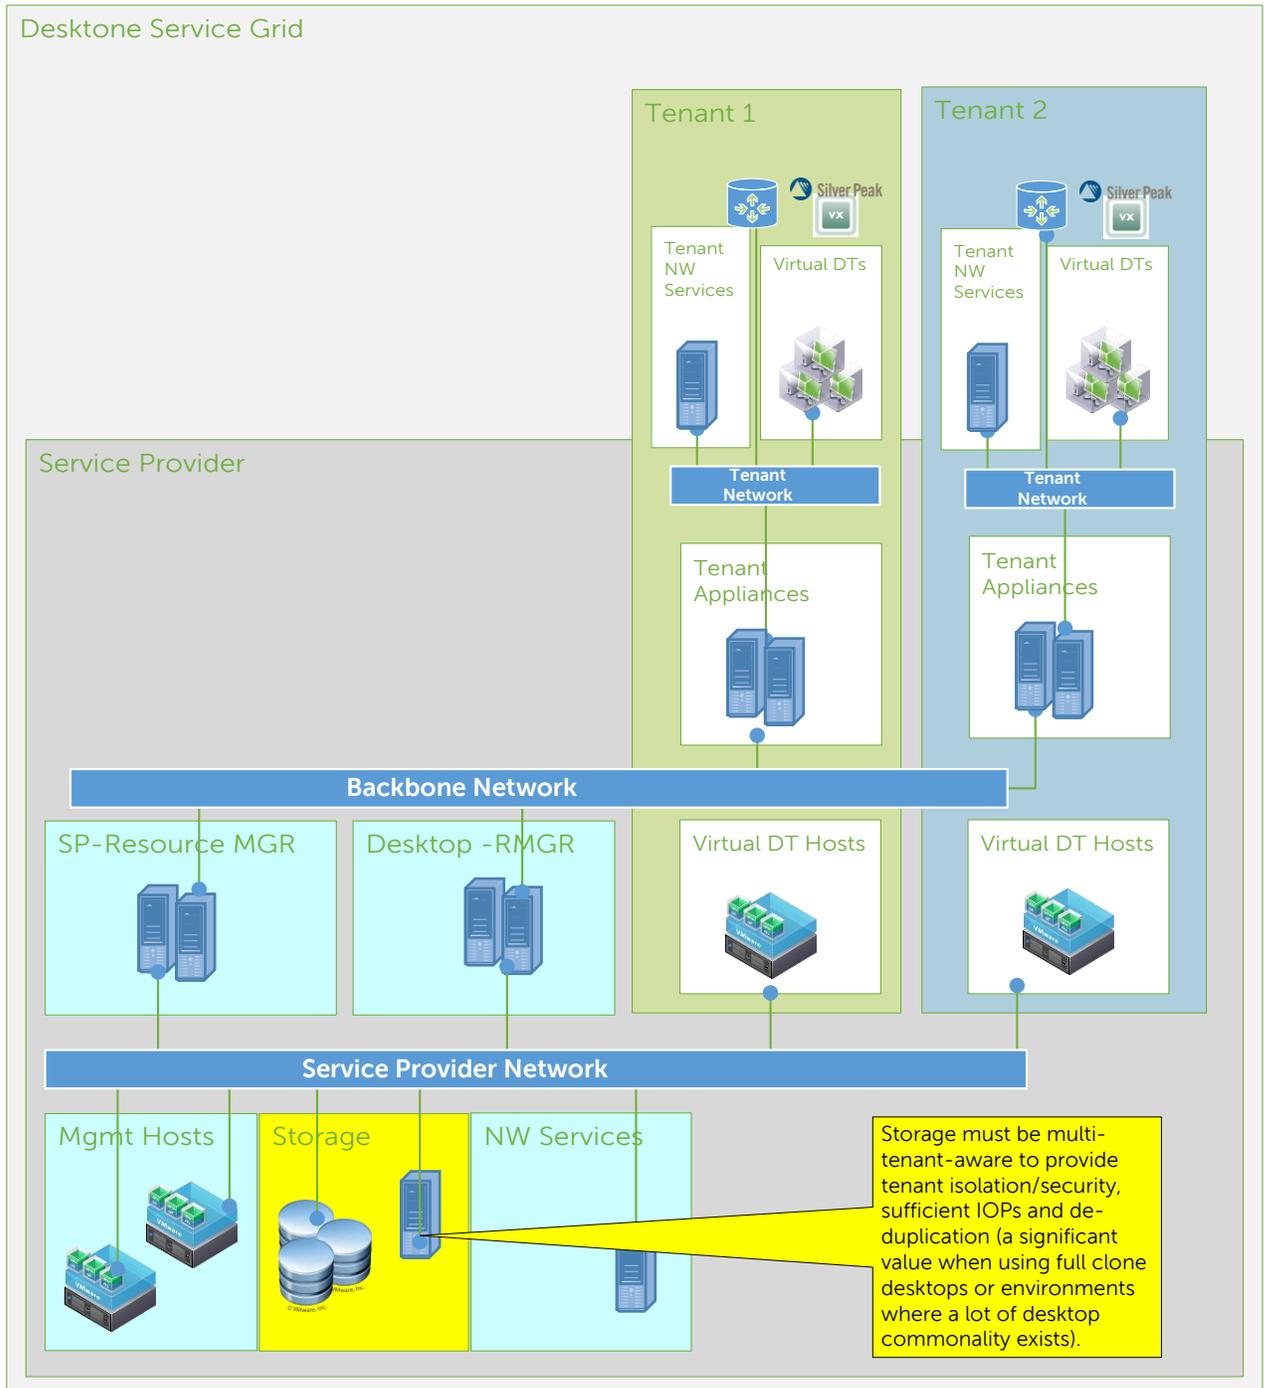
2.3 Network layers



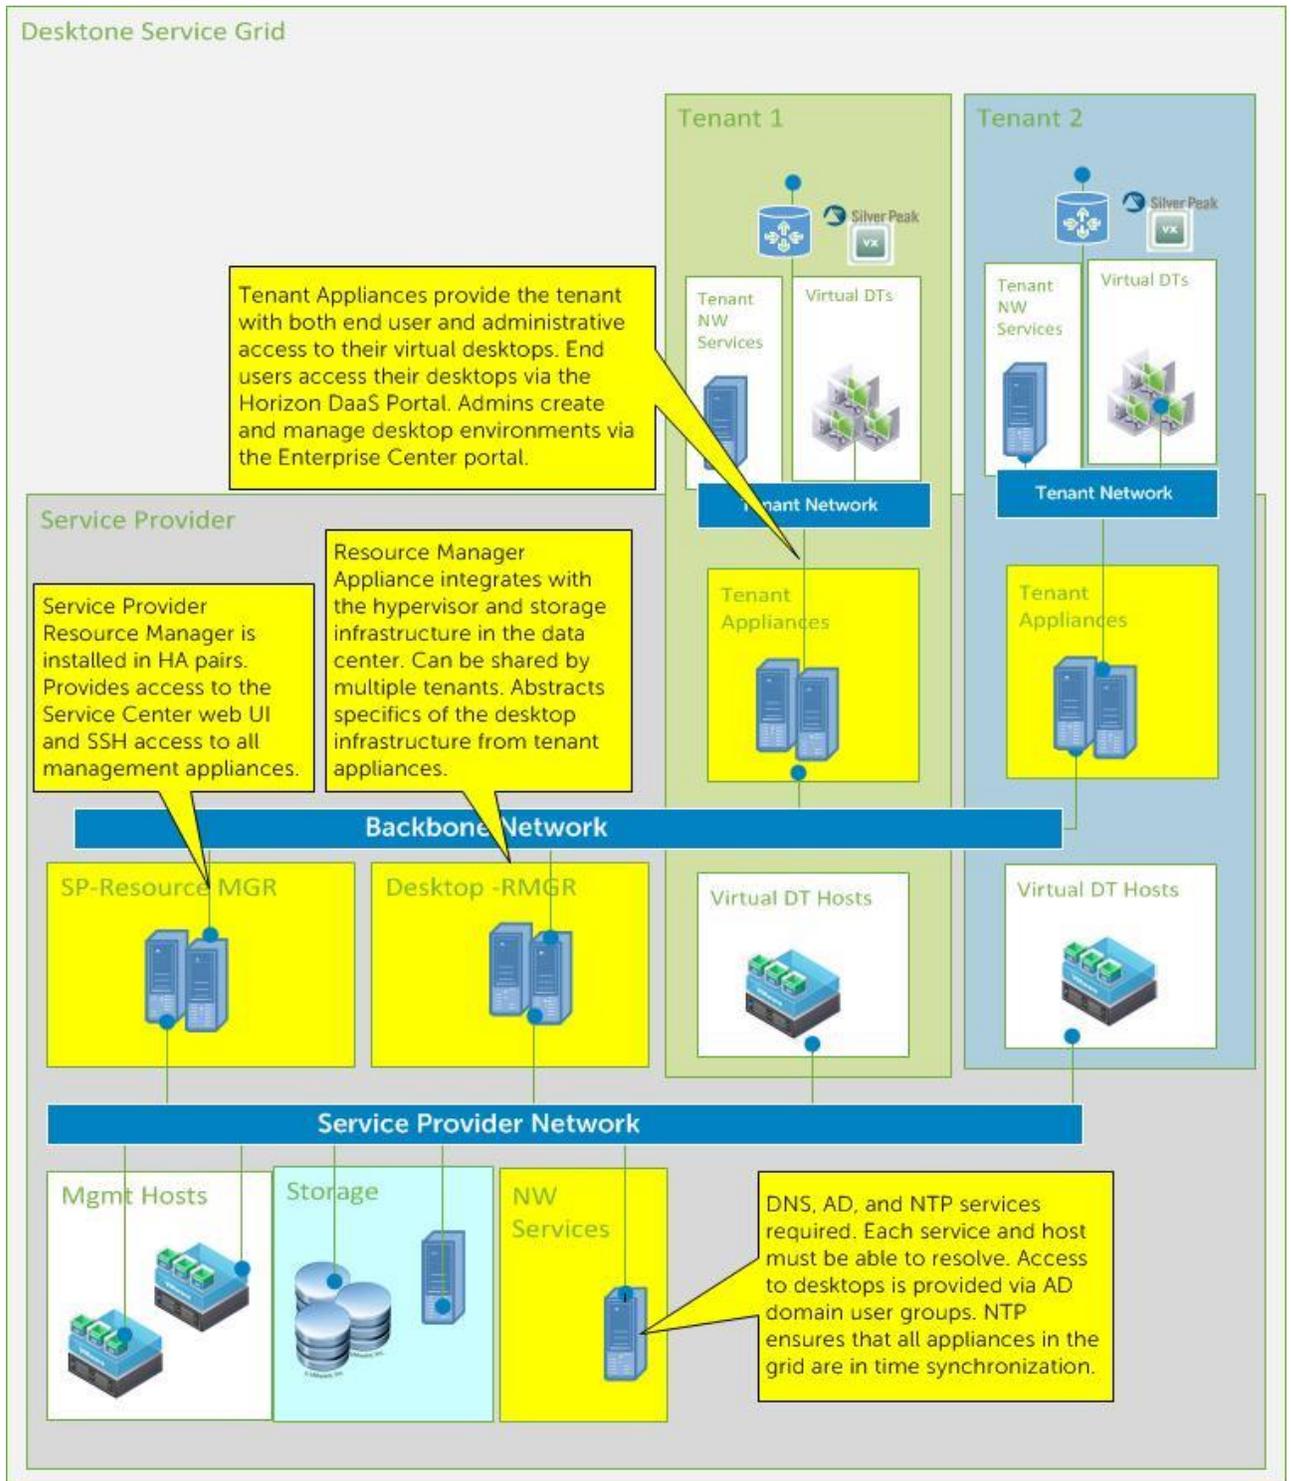
2.4 Compute layers



2.5 Storage layers



2.6 Other services and appliances



3 Architecture components

3.1 Component overview

A high performance Force10 S55 48-port switch is required to get started in the management network layer. This switch will host all solution management and/or “northbound” traffic. Another high performance Force10 S4810 switch is required to support the 10GB iSCSI network layer. This switch will host all storage IP traffic.



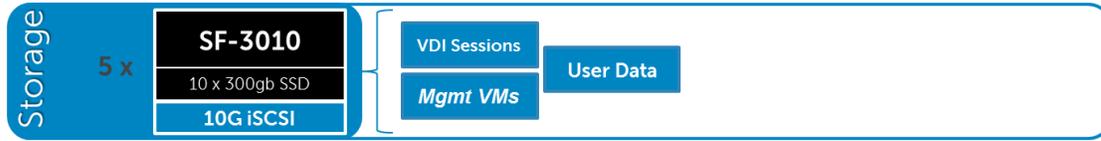
The compute layer consists of the server resources responsible for hosting the VMware Horizon DaaS desktop sessions, hosted via the VMware vSphere hypervisor.



Horizon DaaS management components are dedicated to their own layer so as to not negatively impact the desktop sessions running in the compute layer. This physical separation of resources provides clean, linear and predictable scaling without the need to reconfigure or move resources within the solution as you grow. This compute layer will support all the management VM functions needed to support the Horizon DaaS infrastructure including tenant management functions. There are two levels of separation for management compute infrastructure: Tenant management functions and service provider level management functions.



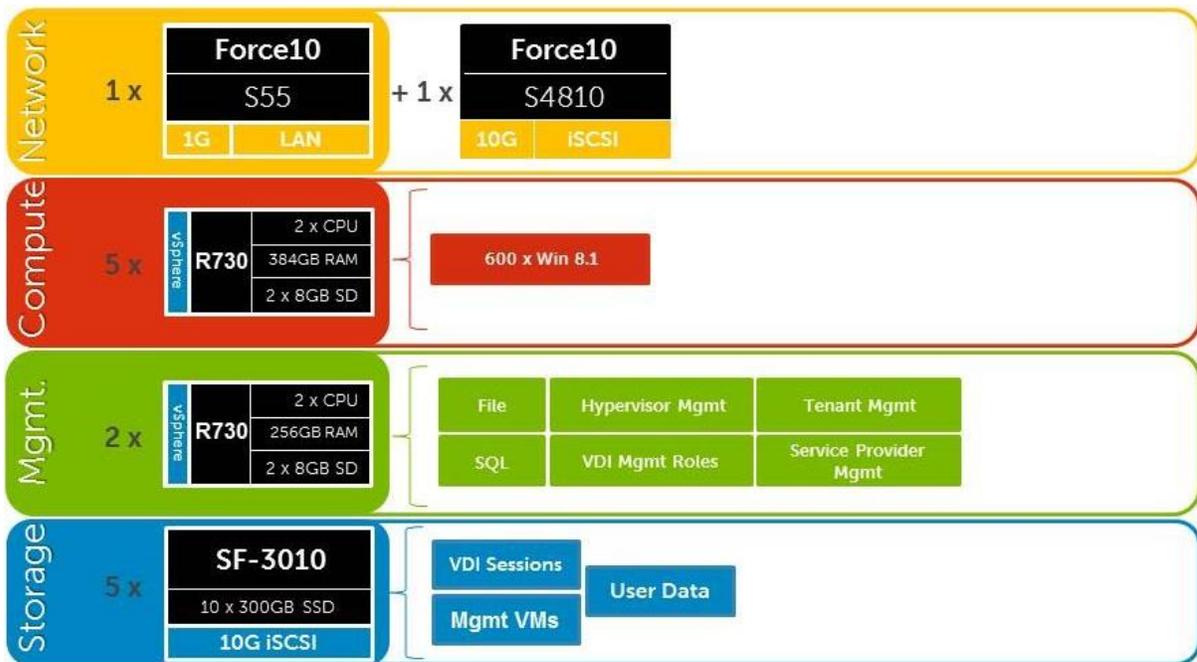
The storage layer consists of options provided by SolidFire for iSCSI to suit your Tier 1 and Tier 2 scaling and capacity needs. It consists of inline deduplication to lower storage requirements across the collective tenant environment and provides SLA/QoS levels on a per volume/tenant basis to ensure a consistent experience from tenant to tenant. A minimum configuration consisting of 5 nodes is required to support full node failure in the storage cluster.



3.2 Shared storage – Rack

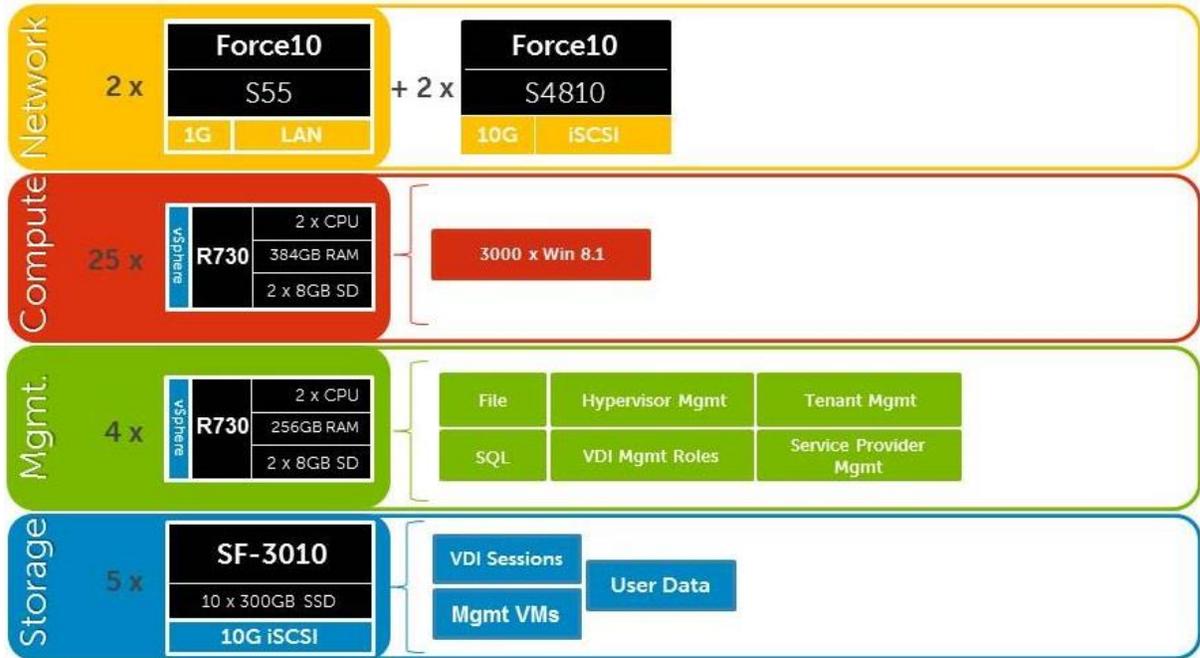
3.2.1 Shared Storage – Rack – 1 Tenant (600 Enhanced users example)

Within the VMware Horizon DaaS offering, tenants/customers are typically brought onboard with dedicated vSphere ESXi hosts either as a starting single node or a collective cluster due to Microsoft desktop OS licensing requirements. While VMware Horizon DaaS does support tenants in a shared mode on the same hardware the “desktop OS” would need to either be Windows 2008 R2 based desktops or otherwise. Tier 1 and Tier 2 can be combined on a single SF-3010 storage array cluster of nodes starting with 5 nodes to provide full node failure resiliency. The following illustration depicts a single tenant with 600 users in their collective cluster. Five SF-3010 nodes are depicted to provide the above mentioned resiliency. This configuration can provide IOPS capacity for approximately 1000 users per node or 5000 users per cluster and can scale linearly on a per node basis.



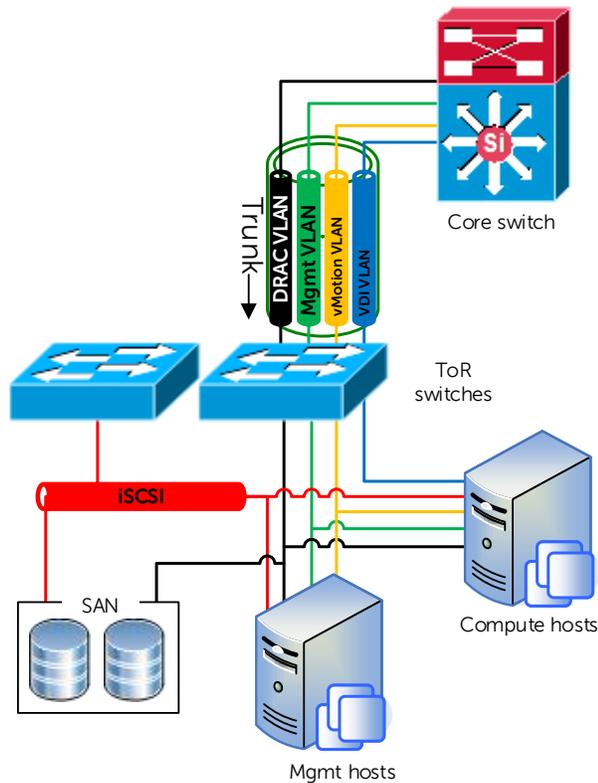
3.2.2 Shared storage – Rack – 1 to 200 Tenants (3000 Enhanced users example)

This illustration depicts a 3000 user build. The 3000 users could belong to a single tenant (as in the diagram below) or they could be spread amongst multiple tenants. In this case, discrete compute hosts would be required for each tenant following desktop scaling guidance to support the number of desktop users for that tenant.

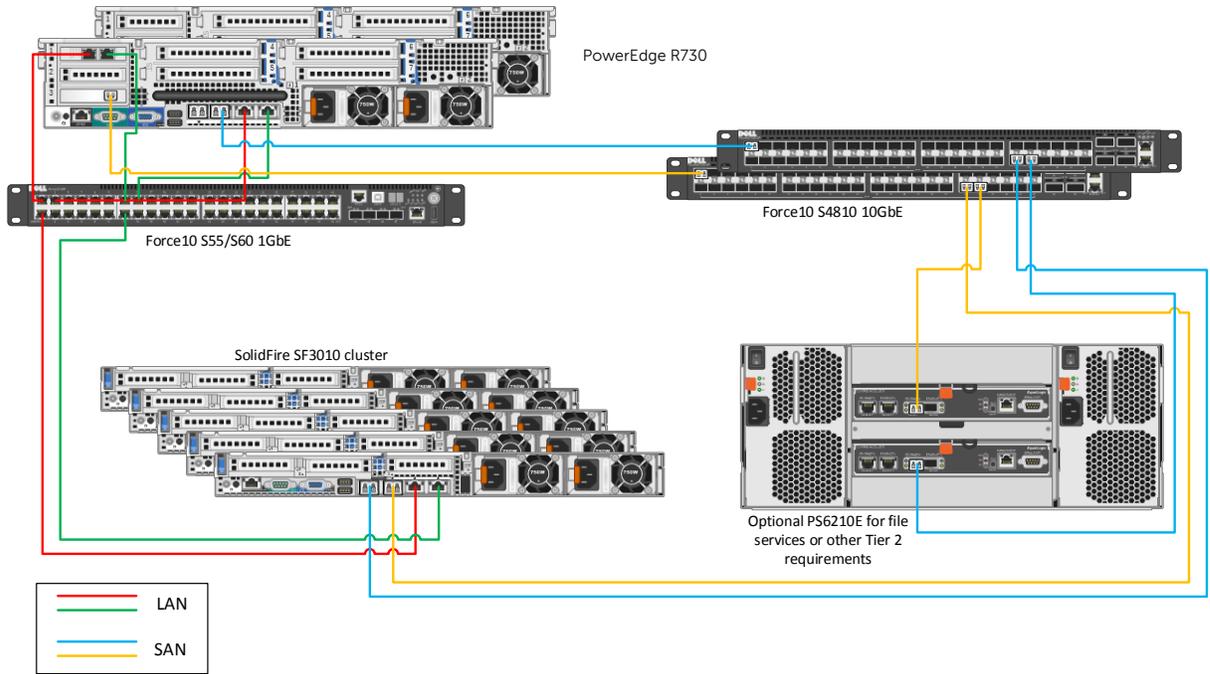


3.2.2.1 Shared storage – Rack – Network architecture (iSCSI)

In the Shared Tier 1 architecture for rack servers, both management and compute servers connect to shared storage in this model. All ToR traffic has been designed to be layer 2 (switched locally), with all layer 3 (routable) VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN traffic flow in relation to the core switch.



3.2.2.2 Shared storage – Rack – Cabling diagram (Rack – SolidFire - EQL)



4 Hardware components

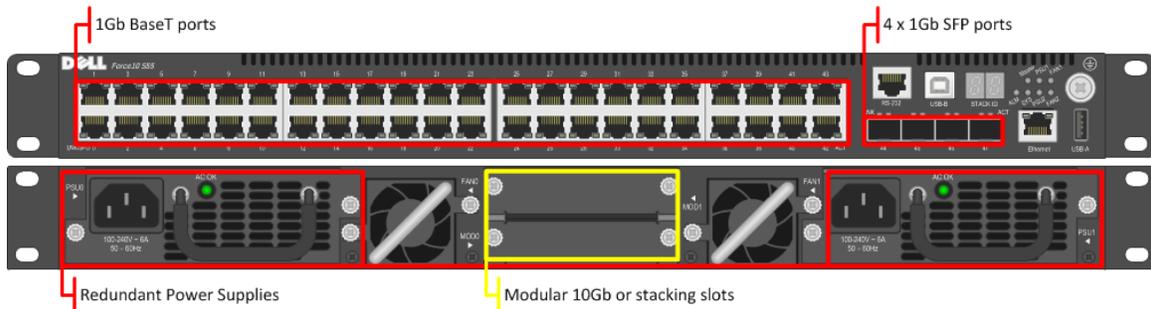
4.1 Networking

The following sections contain the core network components for all Dell Wyse Datacenter solutions. General uplink cabling guidance to consider in all cases is that Twinax is very cost effective for short 10Gb runs and for longer runs it is best to use fiber with SFPs.

4.1.1 Force10 S55 (ToR switch)

The Dell Force10 S-Series S55 1/10 GbE Top-of-Rack (ToR) switch is optimized for lowering operational costs while increasing scalability and improving manageability at the network edge. Optimized for high-performance data center applications, the S55 is recommended for solution deployments of 6000 users or less and leverages a non-blocking architecture that delivers line-rate, low-latency L2 and L3 switching to eliminate network bottlenecks. The high-density S55 design provides 48 GbE access ports with up to four modular 10 GbE uplinks in just 1 rack unit conserving valuable rack space. The S55 incorporates multiple architectural features that optimize data center network efficiency and reliability, including IO panel to PSU airflow or PSU to IO panel airflow for hot/cold aisle environments and redundant, hot-swappable power supplies and fans. A “scale-as-you-grow” ToR solution that is simple to deploy and manage, up to 8 S55 switches can be stacked to create a single logical switch by utilizing Dell Force10 stacking technology and high-speed stacking modules.

Model	Features	Options	Uses
Force10 S55	<ul style="list-style-type: none">44 x BaseT (10/100/1000) + 4 x SFP	<ul style="list-style-type: none">Redundant PSUs4 x 1Gb SFP (Cu or fiber)12 or 24Gb stacking port (up to 8 switches)2 x slots for 10Gb uplink or stacking modules	ToR switch for LAN and iSCSI in Local Tier 1 solution



Guidance:

- 10Gb uplinks to a core or distribution switch are the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports can be used.



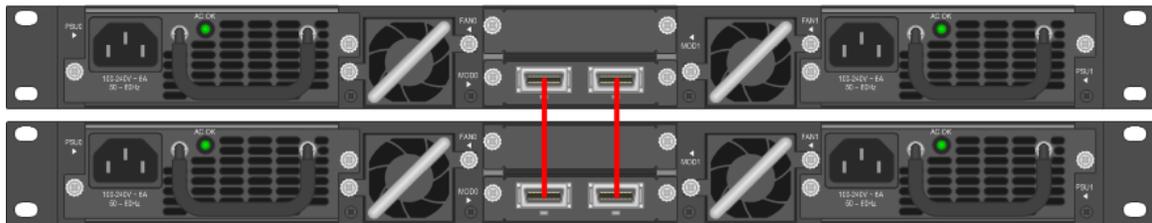
- The front 4 SFP ports can support copper cabling and can be upgraded to optical if a longer run is needed.

For more information on the S55 switch and Dell Force10 networking, please visit:

<http://www.dell.com/us/enterprise/p/force10-s55/pd>

4.1.1.1 Force10 S55 stacking

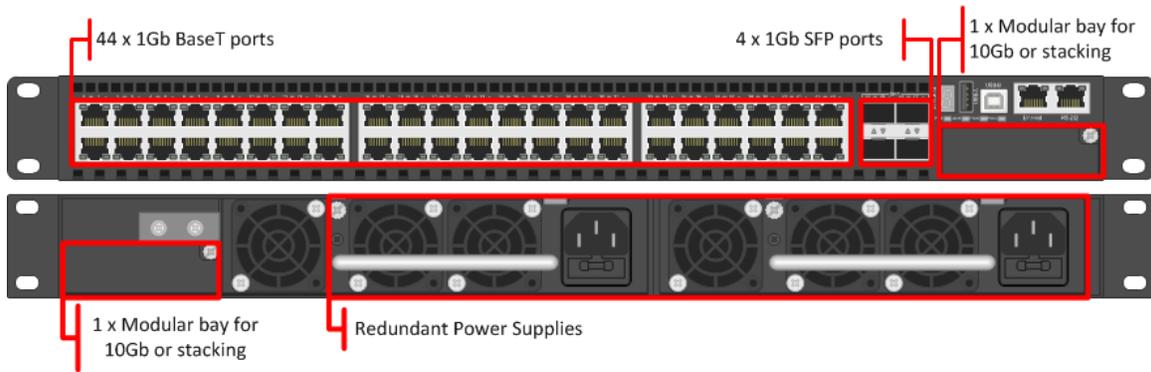
The Top of Rack switches in the Network layer can be optionally stacked with additional switches, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and connected with a stacking cable. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.



4.1.2 Force10 S60 (1Gb ToR switch)

The Dell Force10 S-Series S60 is a high-performance 1/10 GbE access switch optimized for lowering operational costs at the network edge and is recommended for solution deployments of over 6000 users. The S60 answers the key challenges related to network congestion in data center ToR and service provider aggregation deployments. As the use of large data burst applications and services continue to increase, huge spikes in network traffic that can cause network congestion and packet loss also become more common. The S60 is equipped with the industry's largest packet buffer (1.25 GB), enabling it to deliver lower application latency and maintain predictable network performance even when faced with significant spikes in network traffic. Providing 48 line-rate GbE ports and up to four optional 10 GbE uplinks in just 1 rack unit, the S60 conserves valuable rack space. Further, the S60 design delivers unmatched configuration flexibility, high-reliability and power and cooling efficiency to reduce costs.

Model	Features	Options	Uses
Force10 S60	<ul style="list-style-type: none"> • 44 x BaseT (10/100/1000) + 4 x SFP • High performance • High scalability 	<ul style="list-style-type: none"> • Redundant PSUs • 4 x 1Gb SFP (Cu or fiber) • 12 or 24Gb stacking ports (up to 12 switches) • 2 x slots for 10Gb uplink or stacking modules 	Higher scale ToR switch for LAN in Local + Shared Tier 1 and iSCSI in Local Tier 1 solution



Guidance:

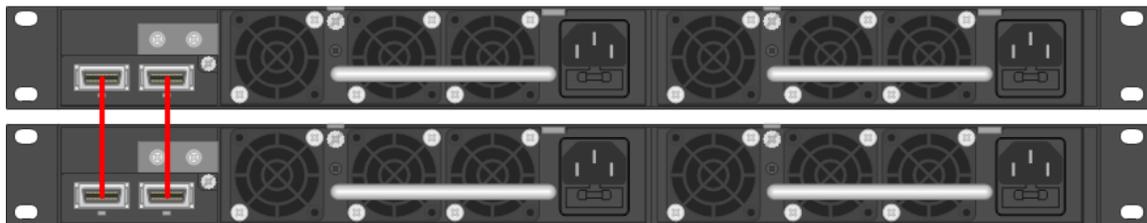
- 10Gb uplinks to a core or distribution switch is the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports can be used.
- The front 4 SFP ports can support copper cabling and can be upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S60 switch and Dell Force10 networking, please visit:

<http://www.dell.com/us/enterprise/p/force10-s60/pd>

4.1.2.1 S60 stacking

The S60 switch can be optionally stacked with 2 or more switches, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and connected with a stacking cable. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.



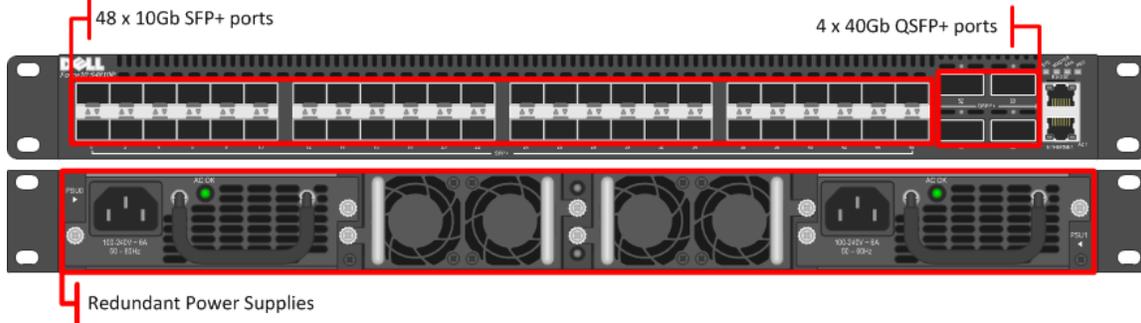
4.1.3 Force10 S4810 (10Gb ToR switch)

The Dell Force10 S-Series S4810 is an ultra-low latency 10/40 GbE ToR switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact S4810 design provides industry-leading density of 48 dual-speed 1/10 GbE (SFP+) ports as well as four 40 GbE QSFP+ uplinks to conserve valuable rack space and simplify the migration to 40 Gbps in the data center core (Each 40 GbE QSFP+ uplink can



support four 10 GbE ports with a breakout cable). Priority-based Flow Control (PFC), Data Center Bridge Exchange (DCBX), Enhance Transmission Selection (ETS), coupled with ultra-low latency and line rate throughput, make the S4810 ideally suited for iSCSI storage, FCoE Transit & DCB environments.

Model	Features	Options	Uses
Force10 S4810	<ul style="list-style-type: none"> 48 x SFP+ (1Gb/10Gb) + 4 x QSFP+ (40Gb) Redundant PSUs 	<ul style="list-style-type: none"> Single-mode/multi-mode optics, Twinax, QSFP+ breakout cables Stack up to 6 switches with SFP or QSFP (2 with VLT) 	ToR switch for iSCSI in Shared Tier 1 solution



Guidance:

- The 40Gb QSFP+ ports can be split into 4 x 10Gb ports using breakout cables for stand-alone units, if necessary. This is not supported in stacked configurations.
- 10Gb or 40Gb uplinks to a core or distribution switch is the preferred design choice.
- The front 4 SFP ports can support copper cabling and can be upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S4810 switch and Dell Force10 networking, please visit: <http://www.dell.com/us/enterprise/p/force10-s4810/pd>

4.1.3.1 S4810 stacking

The S4810 switch can be optionally stacked up to 6 switches or configured to use Virtual Link Trunking (VLT) up to 2 switches. Stacking is supported on either SFP or QSFP ports as long as that port is configured for stacking. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.



4.2 Servers

4.2.1 PowerEdge R730



The rack server platform for the Dell Wyse Datacenter solution is the best-in-class Dell PowerEdge R730. This dual socket CPU platform runs the fastest Intel Xeon E5-2600 v3 family of processors, can host up to 768GB RAM and supports up to 16 2.5" SAS disks. The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor. For more information, please visit:

<http://www.dell.com/us/business/p/poweredge-r730/pd>

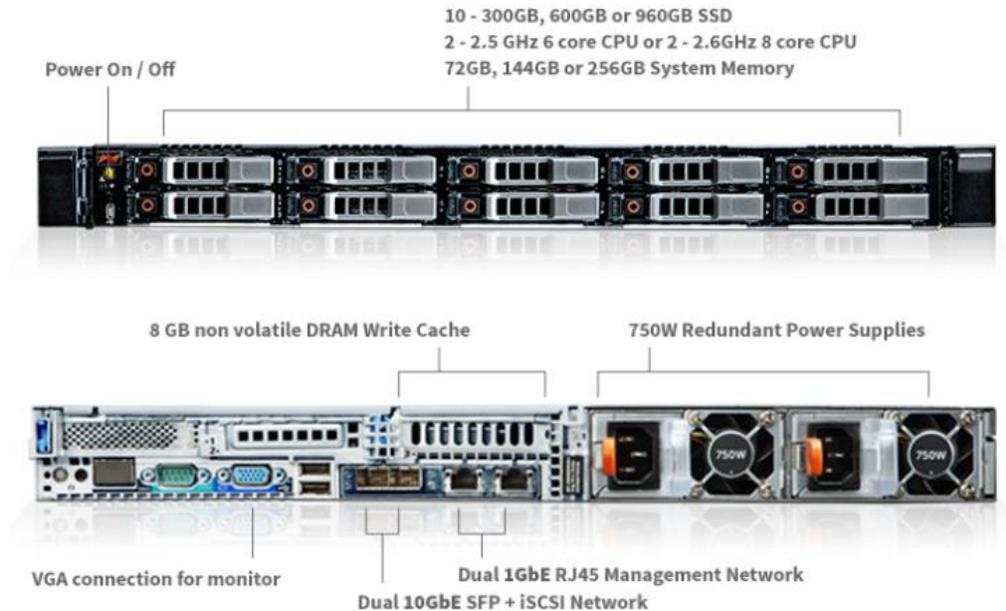
4.3 Storage

4.3.1 SolidFire SF-3010 Tier 1 storage (iSCSI)

Utilizing high-speed, low-latency solid-state disk (SSD) technology in 1U nodes, the SF 3010 10GbE iSCSI array is a virtualized scale-out solution that delivers enhanced storage performance and reliability that is easy to manage and scale for future needs.

For more information please visit: <http://www.solidfire.com/storage-system/product-specifications>

Model	Features	Options	Uses
SF3010	<ul style="list-style-type: none"> 10-drive array (300 GB SSD) 48TB effective capacity (calculated with data high availability and SolidFire's data reduction algorithms) 2 x 2.5 GHz 6 core CPU 72GB Shared System Memory/Read Cache 8GB non-volatile DRAM Write Cache 2 x 10GbE SFP+ iSCSI 2 x 1 GbE RJ45 mgmt 2 Hot-plug redundant 750W PS 1 RU 17.2 kg (38 lbs.) 	<ul style="list-style-type: none"> SF3010 – 300 GB drives SF6010 – 600 GB drives SF9010 – 960 GB drives (higher end processors, etc.) 	Tier 1 array (10Gb – iSCSI)



4.4 Wyse Cloud Clients

The following Wyse Cloud Clients have been tested and are the recommended choices for this solution.

4.4.1 Wyse 5020-P25



Uncompromising computing with the benefits of secure, centralized management. The Wyse 5020-P25 PCoIP zero client for VMware View is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. Smaller than a typical notebook, this dedicated zero client is designed specifically for use with VMware Horizon View. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Wyse 5020-P25 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

4.4.2 Wyse 5012-D10DP



The Wyse 5012-D10DP is a high-performance and secure ThinOS 8 thin client that is absolutely virus and malware immune. Combining the performance of a dual core AMD G-Series APU with an integrated graphics engine and ThinOS, the D10DP offers exceptional thin client PCoIP processing performance for VMware Horizon View environments that handles demanding multimedia apps with ease and delivers brilliant graphics. Powerful, compact and extremely energy efficient, the Wyse 5012-D10DP is a great VDI end point for organizations that need high-end performance but face potential budget limitations.

4.4.3 Wyse 7020-P45



Uncompromising computing with the benefits of secure, centralized management. The Wyse 7020-P45 PCoIP zero client is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. About the size of a notebook, this dedicated zero client designed specifically for VMware Horizon View. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of display performance available over 4 HD displays in a compact, energy-efficient form factor. The Wyse 7020-P45 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

4.4.4 Wyse 7250-Z50D



Designed for power users, the Wyse 7250-Z50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the Wyse 7250-Z50D combines Dell Wyse-enhanced SUSE Linux Enterprise with dual core AMD 1.65 GHz processor and a revolutionary unified engine for an unprecedented user experience. The Wyse 7250-Z50D eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modelling.

4.4.5 Wyse 7290-Z90D7



This is super high performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a dual core AMD processor and a revolutionary unified engine that eliminates performance constraints, the Wyse 7290-Z90D7 achieves incredible speed and power for the most demanding embedded windows applications, rich graphics and HD video. With touch screen capable displays, the Wyse 7290-Z90D7 adds the ease of an intuitive multi touch user experience and is an ideal thin client for the most demanding virtual desktop workload applications.

4.4.6 Chromebook 11



With its slim design and high performance, the Chromebook 11 features a 4th Generation Intel Celeron 2955U processor, 11.6-inch screen, up to 10-hours of battery life and 16GB embedded Solid State Drive which allow it to boot in seconds. The Chromebook 11 is available in two models with either 2GB or 4GB of internal DDR3 RAM. This provides options for the education ecosystem, allowing students, teachers and administrators to access, create and collaborate throughout the day at a price point that makes widespread student computing initiatives affordable. The Chromebook 11 features an 11.6-inch, edge-to-edge glass screen that produces exceptional viewing clarity at a maximum resolution of 1366x768 and is powered by Intel HD Graphics. The high-performing display coupled with a front-facing 720p webcam creates exciting opportunities for collaborative learning. The Chromebook 11 is less than one inch in height and starts at 2.9lbs, making it highly portable. With two USB 3.0 ports, Bluetooth 4.0 and an HDMI port, end users have endless possibilities for collaborating, creating, consuming and displaying content. With battery life of up to 10-hours, the Chromebook is capable of powering end users throughout the day. Also included is a fully compliant HTML5 browser, which makes the Chromebook 11 an excellent choice as an endpoint to connect to an HTML5/BLAST Horizon DaaS VDI desktop.



5 Software components

5.1 VMware Horizon DaaS

VMware Horizon DaaS adds cloud scale to VDI and changes the way desktops are deployed and managed through the marriage of VDI and supporting services to deliver Desktops as a Service.

- By delivering virtual desktops in the cloud, Horizon DaaS users can rapidly provision desktops to users on any device, anywhere, without the upfront costs and complexity of traditional desktop virtualization
- Deliver reliable desktops that users can access from anywhere
- With Horizon DaaS, desktop computing becomes what it should be: a cost-effective service that can be implemented and scaled quickly up or down, delivering immediate financial and service-level improvements.
- Horizon DaaS provides the only Desktop as a Service platform that offers a scalable utility grid

Horizon DaaS is easy to use, and allows for rapid virtual desktop deployment to meet your business needs. The Horizon DaaS Cloud turns virtual desktops into a predictable, easy to budget monthly cost, as opposed to the capital expenditures of a physical PC refresh or an onsite enterprise VDI.

For the complete set of details, please see the Horizon DaaS resources page at <http://www.vmware.com/products/daas>

The core Horizon DaaS components include:

Tenant Appliances

- Provides the tenant with both end user and administrative access to their virtual desktops.
- End users can access and manage their individual virtual desktops via the Desktop Portal.
- Administrators have the ability to create and manage their virtual desktops via the Enterprise Center.

Service Provider Appliance – The first appliance installed in a data center and once bootstrapped provides the foundation to install the remainder of the Desktop application. Installed in HA pairs.

Resource Manager Appliance:

- Integrates with the hypervisor and storage infrastructure in a given data center.
- Can be shared across multiple tenants.
- Abstracts the specifics about the desktop infrastructure from the tenant appliances.

Horizon View Client – Installed on end points. It is the software for creating connections to Horizon DaaS desktops that can be run from Windows, Linux or Mac PCs; tablets, laptops, thin clients and other devices. This is the same client used to connect to Horizon View environments.



Horizon DaaS Portals – A web portal that provides Service Provider administration. Additionally, in a tenant infrastructure, Tenant Administrator and User Access portals are provided to monitor and manage virtual desktop administration and connectivity, respectively.

View Agent – Installed on all desktop VMs that are used as a source for Horizon DaaS desktops. On VMs the agent works in conjunction with the View Direct Connection Agent to provide PCoIP protocol to desktops.

View Direct Connection Agent – Installed on all desktop VMs to provide Horizon View Client to Horizon DaaS desktop PCoIP connectivity.

vCenter Server – This is a server that provides centralized management and configuration to entire virtual desktop and host infrastructure. It facilitates configuration, provision, and management services. It is installed on a Windows Server host or can be deployed using the vCenter appliance (Linux based).

5.2 VMware vSphere 5 hypervisor

VMware vSphere 5 is a virtualization platform used for building VDI and cloud infrastructures. vSphere 5 represents a migration from the ESX architecture to the ESXi architecture.

VMware vSphere 5 includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere client and the vSphere web client.

Throughout the Dell Wyse Datacenter solution, VMware best practices and prerequisites are followed (NTP, DNS, Active Directory, etc.). The vCenter 5 VM used in the solution will be a single Windows Server 2012 R2 VM (Check for current Windows Server OS compatibility at: <http://www.vmware.com/resources/compatibility>), residing on a host in the management tier. SQL Server is a core component of vCenter and will be hosted on another VM also residing in the management tier. All additional Horizon View components need to be installed in a distributed architecture assigning a single role to each VM.

For more information on VMware vSphere, visit <http://www.vmware.com/products/vsphere>



6 Solution architecture

6.1 Compute server infrastructure

In the rack server infrastructure model, VDI sessions execute on shared storage so there is no need for local disks on each server to host VMs. To provide server-level network redundancy using the fewest physical NICs possible, both the Compute and Management servers use a split QP NDC: 2 x 10Gb ports for iSCSI and 2 x 1Gb ports for LAN. Two additional DP NICs (1 x 1Gb + 1 x 10Gb) provide slot and connection redundancy for both network fabrics. All configuration options are identical except for RAM which is reduced on the Management host.

Compute Host PowerEdge R730	Management Host PowerEdge R730
2 x Intel Xeon E5-2697v3 Processor (2.6Ghz)	2 x Intel Xeon E5-2660v3 Processor (2.6Ghz)
384GB Memory (24 x 16GB RDIMMs, 2133MT/s)	256GB Memory (16 x 16GB RDIMMs, 2133MT/s)
VMware vSphere on internal 8GB Dual SD	VMware vSphere on internal 8GB Dual SD
Broadcom 57810 10Gb DP (iSCSI)	Broadcom 57810 10Gb DP (iSCSI)
Broadcom 57800 10Gb QP (LAN/iSCSI)	Broadcom 57800 10Gb QP (LAN/iSCSI)
Broadcom 5720 1Gb DP NIC (LAN)	Broadcom 5720 1Gb DP NIC (LAN)
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

In the above configurations, the R730-based solution will support the following user counts per server:

User Densities per R730 Server	
Workload	Windows 8.1
Standard	180
Enhanced	120
Professional	90



6.2 Management server infrastructure

The Management role requirements for the base solution are summarized below. Use data disks for role-specific application files and data, logs, IIS web files, etc. in the Management volume. Present SolidFire or optional Tier 2 volumes with a special purpose (called out above) in the format specified below:

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Additional Tier 1 or optional Tier 2 Volume (GB)	Sizing
VMware vCenter	2	8	1	40 + 5	100 (VMDK)	
dtRAM Appliance	1	.5	1	20	-	1 pair /DC/Tenant
Service Provider Appliance	1	2	1	20		1 pair /DC
Resource Manager Appliance	1	2	1	20		1 pair /DC/20K VMS
Tenant Appliance	1	2	1	20		1 pair /DC/Tenant/5K users
SQL Server	5	8	1	40 + 5	210 (VMDK)	
File Server	1	4	1	40 + 5	2048 (RDM)	
Total	12	26.5	4	260	2358	

6.2.1 SQL databases

The VMware databases will be hosted by a single dedicated SQL Server 2012 SP1 VM (check DB compatibility at: http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php?) in the Management layer. Use caution during database setup to ensure that SQL data, logs and TempDB are properly separated into their respective volumes. Create all databases that will be required for vCenter.

Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case databases will need to be separated into named instances. Enable auto-growth for each DB.

Follow the best practices as defined by VMware to ensure optimal database performance.

Align all disks to be used by SQL Server with a 1024K offset and then formatted with a 64K file allocation unit size (data, logs and TempDB).

6.2.2 DNS

DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various VMware software components. All hosts, VMs and consumable software components must have a presence in DNS, preferably via a dynamic and AD-integrated namespace. Microsoft best practices and organizational requirements are to be adhered to.

Consideration should also be given to future scaling and access to components that may reside on one or more servers (SQL databases, VMware services) during the initial deployment. Use CNAMEs and the round robin DNS mechanism to provide an alias to the server which is hosting the service or data source.



6.2.3 DNS for SQL

To access the SQL data sources, either directly or via ODBC, a connection to the server name\ instance name must be used. To simplify this process, as well as protect for future scaling (HA), instead of connecting to server names directly, use an alias to these connections in the form of DNS CNAMEs. In other words, instead of connecting to SQLServer1\<instance name> for every device that needs access to SQL, the preferred approach would be to connect to <CNAME>\<instance name>.

For example, the CNAME "VDISQL" is created to point to SQLServer1. If a failure scenario was to occur and SQLServer2 would need to start serving data, we would simply change the CNAME in DNS to point to SQLServer2. No infrastructure SQL client connections would need to be touched.

SQLServer1	Host (A)	10.1.1.28
SQLServer2	Host (A)	10.1.1.29
SQLVDI	Alias (CNAME)	SQLServer1.fcs.local

6.3 Scaling guidance

Each component of the solution architecture scales independently according to the desired number of supported users. Using Intel Ivy Bridge CPUs, rack and blade servers now scale equally from a compute perspective.

- The components can be scaled either horizontally (by adding additional physical and virtual servers to the server pools) or vertically (by adding virtual resources to the infrastructure)
- Eliminate bandwidth and performance bottlenecks as much as possible
- Allow future horizontal and vertical scaling with the objective of reducing the future cost of ownership of the infrastructure.

Component	Metric	Horizontal Scalability	Vertical Scalability
Virtual Desktop Host/Compute Servers	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
dtRAM	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional DT-RAM HA Paired Management VMs.
VMware vCenter	VMs per physical host and/or ESX hosts per vCenter instance	Deploy additional servers and use linked mode to optimize management	Additional vCenter Management VMs.
Component	Metric	Horizontal Scalability	Vertical Scalability
Database Services	Concurrent connections, responsiveness of reads/	Migrate databases to a dedicated SQL server and increase the	Additional RAM and CPU for the management nodes



	writes	number of management nodes	
File Services	Concurrent connections, responsiveness of reads/writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to the optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes
Resource and Tenant Appliances	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional HA Paired Management VMs for Resource or Tenant Appliances

6.3.1 Management/Desktop Server Host Sizing

Sizing the management infrastructure is a multi-dimensional exercise that involves different sizing sub elements. These elements are:

1. Host sizing for Horizon DaaS management hosts. These hosts can support multiple tenant management functions along with service provider level functions but to a limit and as more tenants and/or size of environment scales out you will need to add management appliances and as the capacity of the management servers peaks out, a pair of management hosts will need to be added (to support HA pairs of appliances). This is described below.
2. Appliance sizing scale out. Resource Manager and Tenant Appliance have scale out guidance according to the number of VMs/Desktops they support.
3. vCenter/vSphere scale out to support clusters.
4. Desktop Server compute host in terms of number of desktops can be supported on a Server.

6.3.1.1 Horizon DaaS Management Host sizing considerations

Three components need consideration in order to properly size a Horizon DaaS Management Host:

- **CPU** – Each core supports 10 tenants. For example, ten cores are required to support 100 tenants.
- **Memory** – Each tenant requires 2.5 GB of RAM on each Horizon DaaS management host (2 GB for each of the tenant appliances and 0.5 GB for each of the dtRAM appliances). For example, 250 GB of RAM on each host is required to support 100 tenants.
- **Storage** – Each Tenant requires 56 GB of storage, 28 GB allocated to each Horizon DaaS management host (20 GB for each of the tenant appliances and 8 GB for each of the dtRAM appliances). For example, a pair of management hosts that can scale to 100 tenants requires 2800 GB (2.8 TB) of storage each (5.6 TB total).



6.3.1.2 Horizon DaaS Appliances scale out considerations

Appliance	Memory/Disk Space	Sizing
Service Provider Appliance	2GB/20GB	1 pair / dc
Resource Manager Appliance	2GB/20GB	1 pair / dc / 20,000 VMs
Tenant Appliance	2GB/20GB	1 pair / dc / tenant / 5,000 users
dtRAM Appliance	512MB/8GB	1 pair / dc / tenant

6.3.2 Desktop Host and Appliance scale out matrices

The following tables indicate the desktop OS, management server and appliances density scale out. Note that the Physical Mgmt Server assumes 1 tenant. As described in section 6.3 above and expanded upon in section 6.3's subsections, more tenants or desktops across the architecture will trigger more appliances and thus more pairs of Physical mgmt. servers as do regular desktop compute servers for the number of desktops in a tenant.

6.3.2.1 Windows 7 – vSphere

Rack, Win7, vSphere							
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt. Servers	Physical Host Servers	M1000e Blade Chassis	View Conn. Servers	Virtual vCenter Server
185	130	115	1	1	1	1	1
500	390	345	1	3	1	1	1
1000	780	690	2	6	1	1	1
2000	1430	1265	2	11	1	1	1
3000	2210	1955	2	17	2	2	1
4000	2860	2530	3	22	2	2	1
5000	3640	3220	3	28	2	3	1
6000	4290	3795	4	33	3	3	1
7000	4940	4370	4	38	3	4	1
8000	5720	5060	4	44	3	4	1
9000	6370	5635	4	49	4	5	1
10,000	7150	6325	4	55	4	5	1

Note: All values based on R720 density testing.



6.3.2.2 Windows 8.1 – vSphere

Rack, Win8.1, vSphere							
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt. Servers	Physical Host Servers	M1000e Blade Chassis	View Conn. Servers	Virtual vCenter Server
180*	120*	90*	1	1	1	1	1
500	420	372	1	4	1	1	1
1000	735	651	2	7	1	1	1
2000	1470	1302	2	14	1	1	1
3000	2100	1860	2	20	2	2	1
4000	2835	2511	3	27	2	2	1
5000	3570	3162	3	34	3	3	1
6000	4200	3720	4	40	3	3	1
7000	4935	4371	4	47	4	4	1
8000	5670	5022	4	54	4	4	1
9000	6300	5580	4	60	4	5	1
10,000	7035	6231	4	67	5	5	1

(*) Values based on R730 density testing. All others based on R720 density testing.

6.3.2.3 Windows 2008R2 – vSphere

Rack, Win2008R2, vSphere							
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt. Servers	Physical Host Servers	M1000e Blade Chassis	View Conn. Servers	Virtual vCenter Server
213	150	132	1	1	1	1	1
500	450	396	1	3	1	1	1
1000	750	660	2	5	1	1	1
2000	1500	1320	2	10	1	1	1
3000	2250	1980	2	15	2	2	1
4000	2850	2508	3	19	2	2	1
5000	3600	3168	3	24	2	3	1
6000	4350	3828	4	29	3	3	1
7000	4950	4356	4	33	3	4	1
8000	5700	5016	4	38	3	4	1
9000	6450	5676	4	43	3	5	1
10000	7050	6204	4	47	4	5	1

Note: All values based on R720 density testing.



6.4 Storage architecture overview

The Dell Wyse Datacenter solution has a wide variety of tier 1 and tier 2 storage options to provide maximum flexibility to suit any use case. Customers have the choice to leverage iSCSI solutions from SolidFire for high performance Tier 1 storage and optionally utilize Tier 2 storage for running management VM infrastructure and/or serving File Server user data. With SolidFire's deduplication technology, all storage needs can be served on one platform.

6.4.1 Shared Tier 1 storage

Choosing the Shared Tier 1 storage model means that the virtualization compute hosts are deployed in a diskless mode and leverage shared storage hosted on a high performance SolidFire storage array. In this model, shared storage will be leveraged for Tier 1 and used for VDI execution, write cache and can optionally service Tier 2 functions mentioned previously.

Volumes	Size (GB)	Storage Array	Purpose	File System
VDI-ImagesTenant1	Varies*	Tier 1	Storage for gold image deployment and VDI desktops in same volume sized to house all tenants desktops	VMFS
VDI-Mgmt Infrastructure	2620 GB	Tier 1	Storage for Horizon DaaS infrastructure (can/should be split into discrete volumes to support various management functions)	VMFS

* Calculated based upon number of desktops and size of virtual hard disk

6.4.2 Storage networking – SolidFire 3010 iSCSI

SolidFire implements high-speed, low-latency solid-state disk (SSD) technology 1U nodes. The SF 3010 10GbE iSCSI array is a virtualized scale-out solution that delivers enhanced storage performance and reliability that is easy to manage and scale for future needs. For more information please visit:

<http://www.solidfire.com/storage-system/product-specifications>

Some key settings (See SolidFire documentation for full configuration information) to be used as part of the iSCSI configuration:

- Specify 2 IP addresses for iSCSI on each host
- Specify NICs
- Specify jumbo frames at least 9018 MTU on the switches and on client side (end points, vSwitches) at 9000 MTU to account for jumbo frame overhead and to ensure proper forwarding throughout the network.
- Initialize iSCSI initiator
- Specify IP for the iSCSI target (SolidFire)
- Flow control enablement
- SolidFire recommendation is that switches only be configured to receive pause frames



- Port Fast enablement
- RSTP protocol enablement
- Disable unicast storm control
- Disable delayed ACKs
- Disable large receive offload on network cards

6.5 Virtual networking

6.5.1 Shared Tier 1 – iSCSI

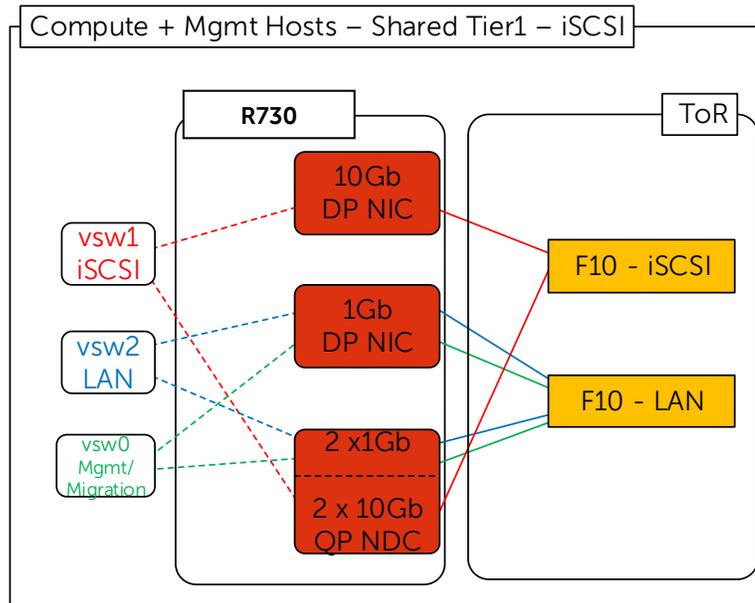
The network configuration in this model is identical to the Compute and Management hosts. Both need access to iSCSI storage since they are hosting VDI sessions from shared storage and both can leverage vMotion as a result. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor management traffic – L3 routed via core switch; VLAN will vary with different tenants
 - vMotion VLAN: Configured for vMotion traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor management traffic – L3 routed via core switch; VLAN will vary with different management services
 - vMotion VLAN: Configured for vMotion traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- An optional iDRAC VLAN can be configured for all hardware management traffic – L3 routed via core switch

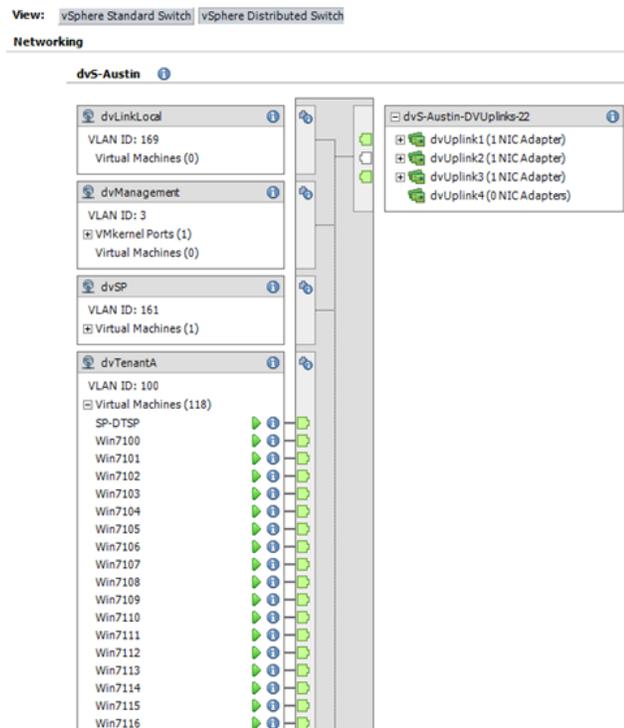
Following best practices, iSCSI and LAN traffic will be physically separated into discrete fabrics. Each Shared Tier 1 Compute and Management host will have a quad port NDC (2 x 1Gb + 2 x 10Gb SFP+), a 10Gb dual port NIC, as well as a 1Gb dual port NIC. Isolate iSCSI onto its own vSwitch with redundant ports. Connections from all 3 vSwitches should pass through both the NDC and add-in NICs per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.



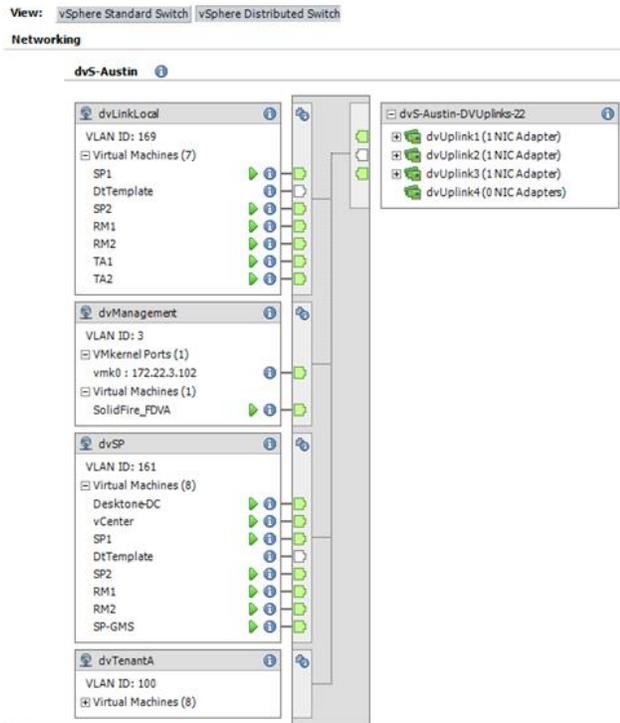
6.5.1.1 vSphere



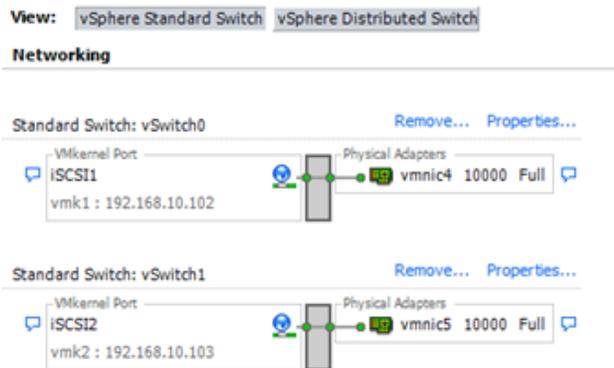
vSwitch0 carries traffic for both Management and vMotion which needs to be VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN will be L3 routable while the vMotion VLAN will be L2 non-routable.



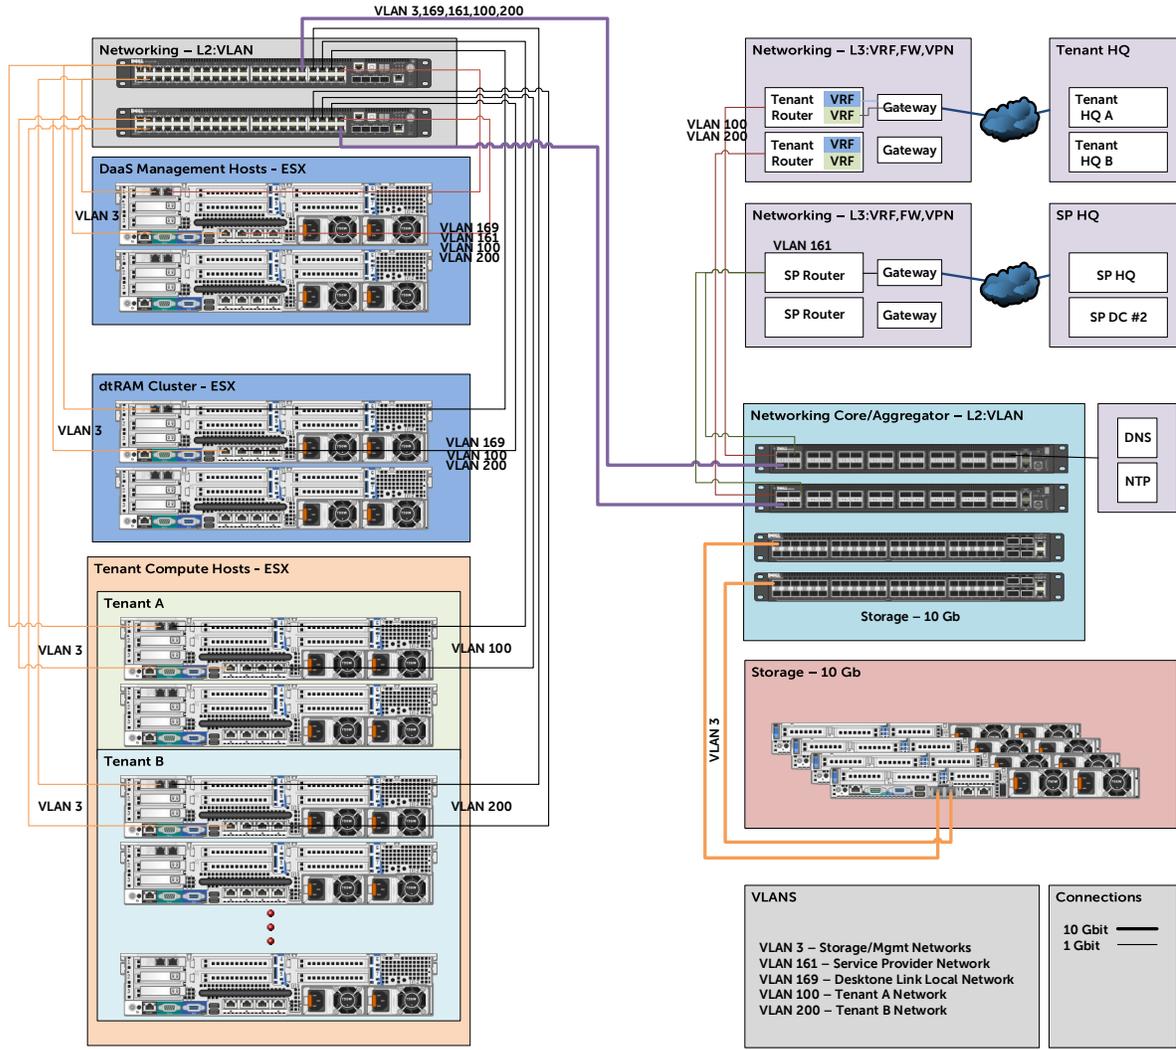
The Management server is configured identically except for the VDI Management VLAN which is fully routed but should be separated from the VDI VLAN used on the Compute host. Care should be taken to ensure that all vSwitches are assigned redundant NICs that are not from the same PCIe device.



iSCSI sample vSwitch configuration:



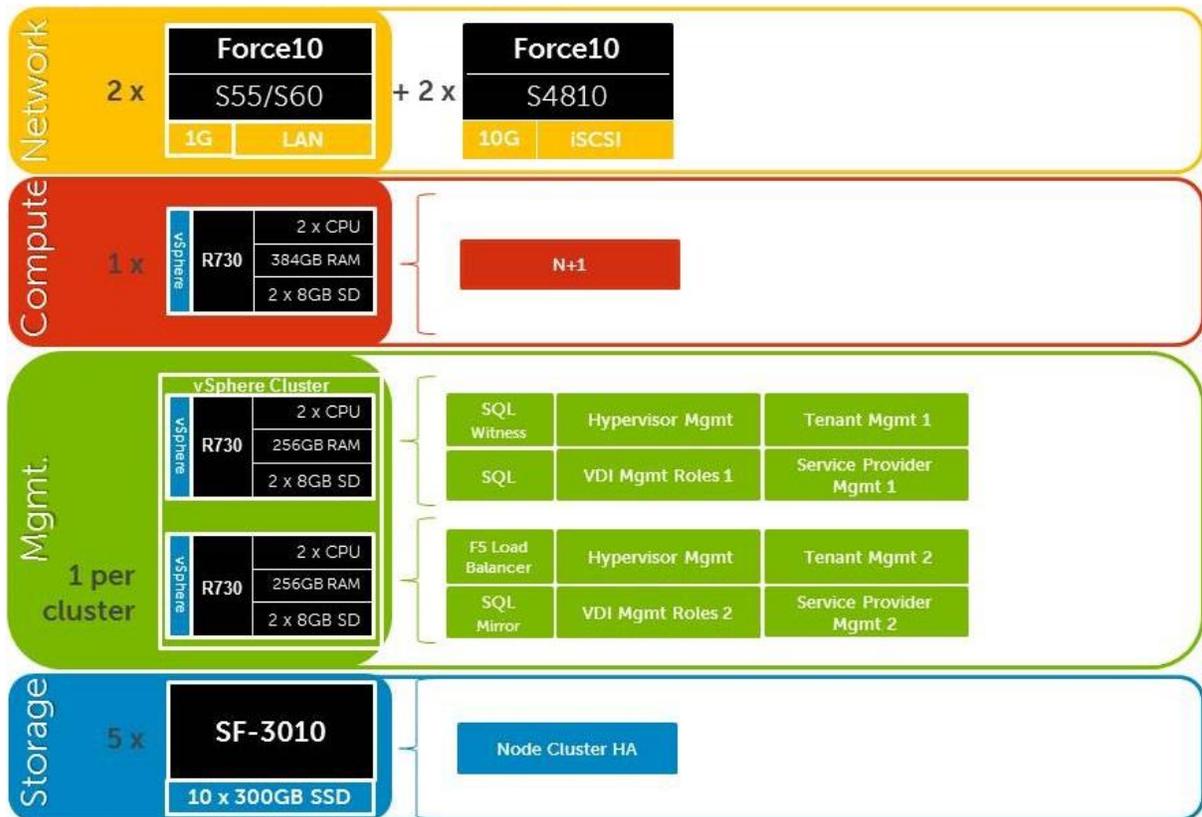
The following is a sample depiction of an overall datacenter wiring diagram.



6.6 Solution high availability

High availability (HA) is offered to protect each layers of the solution architecture, individually if desired. Following the N+1 model, additional ToR switches for LAN, iSCSI, or FC are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere clustering is introduced in the management layer, SQL is mirrored or clustered, an F5 device can be leveraged for load balancing, all management appliances in the Horizon DaaS architecture are only deployed in HA pairs, full node failure HA storage clusters and lastly backup strategies can be implemented to protect user data.





The HA options provides redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

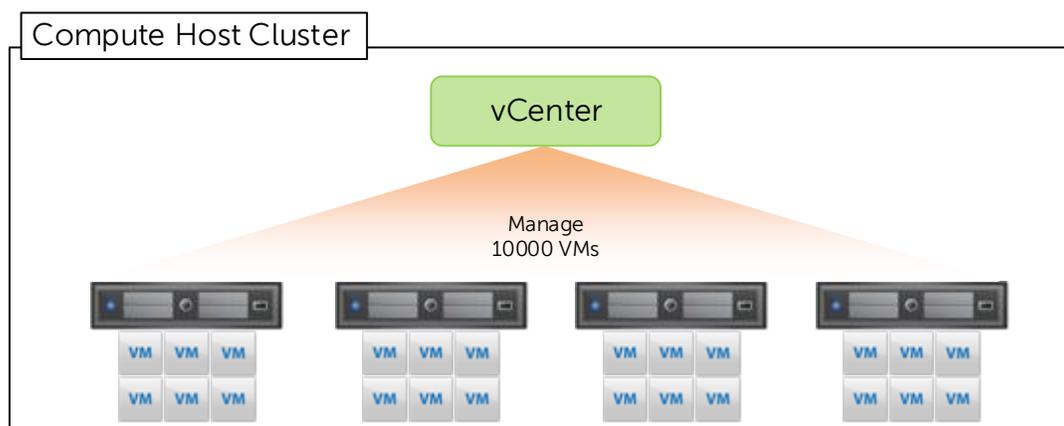
- An additional switch is added at the network tier which will be configured with the original as a stack and equally spreading each host's network connections across both.
- At the tenant compute tier an additional ESXi host is added to provide N+1 protection provided by vSphere for server tier protection. vSphere will provide HA plus load balancing (DRS) in this N+1 clustered scenario.
- A number of enhancements occur at the Management tier, the first of which is the addition of another host. The Management hosts will then be configured in an HA cluster. All applicable Horizon DaaS management roles are built in HA pairs and each VM appliance should be set to run on a separate host within the collective management cluster. SQL will also receive greater protection through the addition and configuration of a SQL mirror with a witness.
- Management appliance tiers should have their respective databases backed with proper backup strategies. Since the appliances can be recreated much quicker than a full restore the database is the better backup strategy.

6.6.1 vSphere HA

Both Compute and Management hosts are identically configured, within their respective tiers and leverage shared storage so can make full use of vSphere HA. The Compute hosts can be configured in an HA



cluster following the boundaries of vCenter with respect to limits imposed by VMware. This will result in multiple HA clusters managed by multiple vCenter servers.



A single HA cluster will be sufficient to support the Management layer up to 10K users. An additional host can be used as a hot standby or to thin the load across all hosts in the cluster.

6.6.2 Horizon DaaS infrastructure protection

VMware Horizon DaaS infrastructure data protection with Dell Data Protection. As mentioned before, Horizon DaaS management appliances are deployed in HA pairs so are inherently redundant. The appliances can be recreated if needed much easier than a full backup restore. The databases that support these appliances should be backed via recommended Horizon DaaS best practices. VMware provides backup and restore scripts for the embedded appliance databases with the Horizon DaaS platform. The backup scripts should be executed at least daily and onsite/offsite backup strategies should be implemented to ensure an expedited restore process. More information to support vCenter, SQL Server (non-embedded Horizon DaaS databases) and other similar auxiliary infrastructure can be protected via Dell Data Protection technology. Information is at: <http://dell.to/1ed2dQf>

6.6.3 Management server high availability

The applicable core Horizon DaaS generic roles can be load balanced via DNS by default. In environments requiring HA, F5 can be introduced to manage load-balancing efforts. vCenter configurations (optionally vCenter Update Manager) are stored in SQL Server which will be protected via the SQL mirror.

If the customer desires, some Role VMs can be optionally protected further via the form of a cold stand-by VM residing on an opposing management host. A vSphere scheduled task can be used, for example, to clone the VM to keep the stand-by VM current.

The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a vSphere cluster.
- SQL Server mirroring is configured with a witness to further protect SQL.

6.6.4 Windows File Services high availability

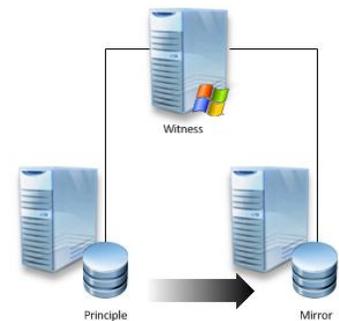
High availability for file services will be provided by the Dell FS7600, FS8600 or PowerVault NX3300 clustered NAS devices. To ensure proper redundancy, distribute the NAS cabling between ToR switches.

Unlike the FS8600, the FS7600 and NX3300 do not support for 802.1q (VLAN tagging) so configure the connecting switch ports with native VLANs, both iSCSI and LAN/ VDI traffic ports. Best practice dictates that all ports be connected on both controller nodes. The back-end ports are used for iSCSI traffic to the storage array as well as internal NAS functionality (cache mirroring and cluster heart beat). Front-end ports can be configured using Adaptive Load Balancing or a LAG (LACP).

The Dell Wyse Solutions Engineering recommendation is to configure the original file server VM to use RDMs to access the storage LUNs, therefore migration to the NAS will be simplified by changing the presentation of these LUNs from the file server VM to the NAS.

6.6.5 SQL Server high availability

HA for SQL will be provided via a 3-server synchronous mirror configuration that includes a witness (High safety with automatic failover). This configuration will protect all critical data stored within the database from physical server as well as virtual server problems. DNS will be used to control access to the active SQL server. Place the principal VM that will host the primary copy of the data on the first Management host. Place the mirror and witness VMs on the second or later Management hosts. Mirror all critical databases to provide HA protection.



The following article details the step-by-step mirror configuration:

<http://www.sqlserver-training.com/how-to-setup-mirroring-in-sql-server-screen-shots>

Additional resources can be found in TechNet:

<http://technet.microsoft.com/en-us/library/ms189047.aspx>

<http://technet.microsoft.com/en-us/library/ms188712.aspx>

7 Solution performance and testing

7.1 Load generation and monitoring

7.1.1 VMware View Planner

View Planner, currently at version 3, operates in two modes; Benchmark mode and Flexible mode. Benchmark mode is used in a locked down or fixed mode and is designed to run a workload that cannot be changed for the purpose of determining a standardized result on a given set of hardware. Flexible mode is more of a traditional version designed to allow customization of workloads to try and emulate target scenarios and produce results designed to determine the scale and density of a given infrastructure. Similar to Login VSI, “helper” or “launcher” systems aid in the testing process with View Planner 3.0. This tool is required for testing Dell Wyse Datacenter for VMware Horizon View Solutions in VMware’s Rapid Desktop Program.

7.1.2 Login VSI – Login Consultants

Login VSI is the de-facto industry standard tool for testing VDI environments and server-based computing / terminal services environments. It installs a standard collection of desktop application software (e.g. Microsoft Office, Adobe Acrobat Reader etc.) on each VDI desktop; it then uses launcher systems to connect a specified number of users to available desktops within the environment. Once the user is connected the workload is started via a logon script which starts the test script once the user environment is configured by the login script. Each launcher system can launch connections to a number of ‘target’ machines (i.e. VDI desktops), with the launchers being managed by a centralized management console, which is used to configure and manage the Login VSI environment. Important to note is that there are some performance changes between VSI 3.7 and 4.0. Namely, desktop read/write IO has decreased a bit while CPU utilization has increased in 4.0.



7.1.3 Liquidware Labs Stratusphere UX

Stratusphere UX was used during each test run to gather data relating to User Experience and desktop performance. Data was gathered at the Host and Virtual Machine layers and reported back to a central server (Stratusphere Hub). The hub was then used to create a series of “Comma Separated Values” (.csv) reports which have then been used to generate graphs and summary tables of key information. In addition the Stratusphere Hub generates a magic quadrant style scatter plot showing the Machine and IO experience of the sessions. The Stratusphere hub was deployed onto the core network therefore its

monitoring did not impact the servers being tested. This core network represents an existing customer environment and also includes the following services:

- Active Directory
- DNS
- DHCP
- Anti-Virus

Stratusphere UX calculates the User Experience by monitoring key metrics within the Virtual Desktop environment, the metrics and their thresholds are shown in the following screen shot:

Machine Experience Indicators						
	Weight (%)	Good	Fair	Poor		
Login Delay : Time it takes to login (sec.) ?	20	0 <= 15	<= 60	<= unbounded		
Application Load Time : Avg. startup time for applications (sec.) ?	20	0 <= 10	<= 30	<= unbounded		
CPU Queue Length : Length of CPU queue at inspection time ?	20	0 <= 3	<= 6	<= unbounded		
Page Faults : Number of page faults during inspection interval ?	20	0 <= 2,000	<= 10,000	<= unbounded		
Non-Responding Applications : Number of unresponsive applications at inspection time ?	20	0 <= 2	<= 3	<= unbounded		

I/O Experience Indicators						
	Weight (%)	Good	Fair	Poor		
Disk Load : Avg. disk IO per second ?	25	0 <= 25	<= 75	<= unbounded		
Disk Queue Length : Avg. length of disk queue(s) ?	25	0 <= 1	<= 3	<= unbounded		
Network Latency : Avg. network roundtrip time (ms) ?	25	0 <= 150	<= 300	<= unbounded		
Failed Connections : Number of outgoing connection attempts that failed ?	25	0 <= 5	<= 15	<= unbounded		

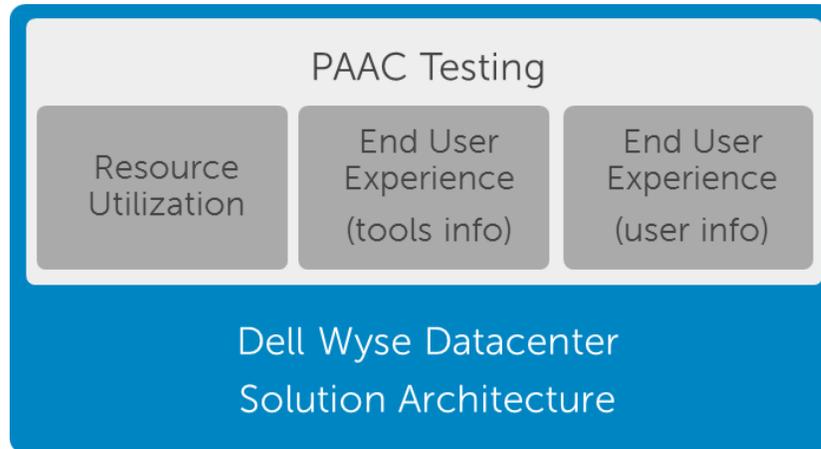
7.1.4 VMware vCenter

VMware vCenter has been used for VMware vSphere-based solutions to gather key data (CPU, memory and network usage) from each of the desktop hosts during each test run. This data was exported to .csv files for each host and then consolidated to show data from all hosts. While the report does not include specific performance metrics for the Management host servers, these servers were monitored during testing and were seen to perform above the expected acceptable level.



7.2 Performance analysis methodology

In order to ensure the optimal combination of end user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell Wyse Datacenter solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing. This methodology is based on the three factors below. Login VSI is currently the load-testing tool used during PAAC of Dell Wyse Datacenter solutions; Login VSI is the de-facto industry standard for VDI and server-based computing (SBC) environments and is discussed in more detail below.



7.2.1 Resource utilization

Poor end user experience is one of the main risk factors when implementing desktop virtualization but the root cause for poor end user experience is resource contention – hardware resources at some point in the solution have been exhausted causing poor performance. In order to ensure that this has not happened (and that it is not close to happening), PAAC on Dell Wyse Datacenter solutions monitors the relevant resource utilization parameters and applies relatively conservative thresholds as shown in the table below. As discussed above, these thresholds are carefully selected to deliver an optimal combination of good end user experience and cost-per-user, while also providing burst capacity for seasonal / intermittent spikes in usage. These thresholds are used to decide the number of virtual desktops (density) that can be hosted by a specific hardware environment (i.e. combination of server, storage and networking) that forms the basis for this Dell Wyse Datacenter for VMware Horizon View Reference Architecture.

Resource Utilization Thresholds	
Parameter	Pass / Fail Threshold
Physical host CPU utilization	85%
Physical host memory utilization	85%
Network throughput	85%
Storage IO latency	20ms

7.2.2 EUE tools information

Good EUE is one of the primary factors in determining the success of a VDI implementation. As a result, a number of vendors have developed toolsets that monitor the environmental parameters that are relevant to EUE. PAAC on Dell Wyse Datacenter solutions uses the Liquidware Labs Stratusphere UX tool to ensure that good EUE is delivered for the density numbers defined in our RAs. More specifically, our PAAC analysis uses a scatter plot provided by Stratusphere UX which presents end-user experience for all load-testing users. Stratusphere UX does this by algorithmically combining relevant parameters in relation to virtual machine experience (e.g. login duration) and virtual desktop IO experience (e.g. disk queue length) to provide a plot that shows end user experience as good, fair or poor using a golden-quadrant type approach.

7.2.3 EUE real user information

To complement the tools-based end user experience information gathered using Stratusphere UX (as described above) and to provide further certainty around the performance of Dell Wyse Datacenter solutions, PAAC on our solutions also involves a user logging into one of the solutions while it is fully populated (based on the density specified in the relevant Reference Architecture) and executing user activities that are representative of the user type being tested (e.g. task, knowledge or power user). An example would be a knowledge worker executing various activities in Excel. The purpose of this activity is to verify that the end-user experience is as good as the user would expect on a physical laptop or desktop.

7.2.4 Workload and profile definitions

It is important to understand user workloads and profiles when designing a desktop virtualization solution in order to understand the density numbers that the solution can support. For our testing, we use three workload / profile levels, each of which is bound by specific metrics and capabilities. In addition, we use workloads and profiles that are targeted at graphics-intensive use cases. We have presented detailed information for these workloads and profiles below, however it is useful to define the terms “workload” and “profile” as they are used in this document.

- **Profile** – The configuration of the virtual desktop – the number of vCPUs and amount of RAM configured on the desktop (i.e. visible to the user).
- **Workload** – The set of applications used for PAAC of Dell Wyse Datacenter solutions (e.g. Microsoft Office applications, PDF reader, Internet Explorer etc.).

7.2.5 Profile specifications

The table below presents the profiles used during PAAC of the Dell Wyse Datacenter solutions. These profiles have been carefully selected to provide the optimal level of resources for the most common use cases.

Profile Name	vCPUs per VM	Memory per VM	Use Case
Standard	1	2 GB	Task worker
Enhanced	2	3 GB	Knowledge worker
Professional	2	4 GB	Power user



Shared Graphics	2 + shared GPU	3 GB	Knowledge worker with high graphics requirements
Dedicated Graphics	4 + dedicated GPU	32 GB	Workstation-type user producing complex 3D models

7.2.6 Workload specifications

Load testing on each of the profiles described in the above table is carried out using an appropriate workload that is representative of the relevant use case. In the case of the non-graphics use cases, the workloads are Login VSI workloads. In the case of graphics use cases, the workloads are specially designed workloads that stress the VDI environment to a level that is appropriate for the relevant use case. This information is summarized in the table below.

Profile Name	Workload	OS Image
Standard	Login VSI Light	Shared
Enhanced	Login VSI Medium	Shared
Professional	Login VSI Heavy	Shared + profile virtualization
Shared Graphics	Fishbowl / eDrawings workload	Shared + profile virtualization
Dedicated Graphics	eDrawings / AutoCAD – SPEC Viewperf workload	Persistent

With respect to the table above, additional information for each of the workloads is given below. It should be noted that for Login VSI testing, the following login and boot paradigm was used:

- For single-server / single-host testing (typically carried out to determine the virtual desktop capacity of a specific physical server), users were logged in every 30 seconds.
- For multi-host / full solution testing, users were logged in over a period of 1-hour, to replicate the normal login storm in an enterprise environment.
- All desktops were fully booted prior to each login attempt.

For all testing, virtual desktops ran an industry-standard anti-virus solution (McAfee VirusScan Enterprise) in order to replicate a typical customer environment.

7.2.6.1 Login VSI light workload

Compared to the Login VSI medium workload described below, the light workload runs fewer applications (mainly Excel and Internet Explorer with some minimal Word activity) and starts/stops the applications less frequently. This results in lower CPU, memory and disk IO usage.

7.2.6.2 Login VSI medium workload

The Login VSI medium workload is designed to run on 2 vCPUs per desktop VM. This workload emulates a medium knowledge worker using Office, IE, PDF and Java/FreeMind. The Login VSI medium workload has the following characteristics



- Once a session has been started the workload will repeat (loop) every 48 minutes.
- The loop is divided in 4 segments, each consecutive Login VSI user logon will start a different segments. This ensures that all elements in the workload are equally used throughout the test.
- The medium workload opens up to 5 applications simultaneously.
- The keyboard type rate is 160 ms for each character.
- Approximately 2 minutes of idle time is included to simulate real world users.

Each loop will open and use:

- Outlook, browse messages.
- Internet Explorer, browsing different webpages and a YouTube style video (480p movie trailer) is opened three times in every loop.
- Word, one instance to measure response time, one instance to review and edit a document.
- Doro PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel, a very large randomized sheet is opened.
- PowerPoint, a presentation is reviewed and edited.
- FreeMind, a Java based Mind Mapping application.

7.2.6.3 Login VSI heavy workload

The heavy workload is based on the medium workload except that the heavy workload:

- Begins by opening 4 instances of Internet Explorer. These instances stay open throughout the workload loop.
- Begins by opening 2 instances of Adobe Reader. These instances stay open throughout the workload loop.
- There are more PDF printer actions in the workload.
- Instead of 480p videos a 720p and a 1080p video are watched.
- Increased the time the workload plays a flash game.
- The idle time is reduced to 2 minutes.

7.3 Testing and validation

7.3.1 Testing process

The purpose of the single server testing is to validate the architectural assumptions made about the server stack. Each user load is tested against 4 runs. A pilot run to validate that the infrastructure is functioning and valid data can be captured and 3 subsequent runs allowing correlation of data. Summary of the test results are listed below in tabular format.

At different stages during testing, the team completed manual "User Experience" testing while the environment was under load. This involved a team member logging into a session during the testing run and completing tasks similar to the User Workload description. While this experience was subjective, it helped to provide a better understanding of the end user experience of the desktop sessions, particularly under high load to ensure that the data gathered was reliable.

Login VSI has two modes for launching user's sessions:

Parallel – Sessions are launched from multiple launcher hosts in a round robin fashion; this mode is recommended by Login Consultants when running tests against multiple host servers. In parallel mode the VSI console is configured to launch a number of sessions over a specified time period (specified in seconds)

Sequential - Sessions are launched from each launcher host in sequence; sessions are only started from a second host once all sessions have been launched on the first host and is repeated for each launcher host. Sequential launching is recommended by Login Consultants when testing a single desktop host server. The VSI console is configured to launch a specific number of sessions at a specified interval specified in seconds

All test runs were conducted using the Login VSI "Parallel Launch" mode and all sessions were launched over an hour in an attempt to simulate a logon storm. Once the last user session has connected, the sessions are left to run for 15 minutes prior to the sessions being instructed to logout at the end of the current task sequence, this allows every user to complete a minimum of two task sequences within the run before logging out. The single server test runs were configured to launch user sessions every 60 seconds, as with the full bundle test runs sessions were left to run for 15 minutes after the last user connected prior to the sessions being instructed to log out.

7.4 VMware Horizon DaaS test results

This validation was designed to evaluate the capabilities of the Ivy Bridge processors when used in a VMware Horizon 6 DaaS environment with Windows 8.1. Intel E5-2697 v3 processors were used in rack servers and the Horizon DaaS solutions were deployed on the ESXi 5.1 hypervisor.

This validation was performed on single server compute host solutions, running on a PowerEdge R730 host. The compute hosts had 384GB of RAM and dual Intel E5-2697 v3 2.6GHz 14 core processors.

All the results presented below were performed on SolidFire All-Flash storage to provide inline deduplication, QoS storage performance for multi-tenancy, and high performance storage.

Validation was performed using Dell Wyse Datacenter standard testing methodology using LoginVSI load generation tool for VDI benchmarking that simulates production user workloads. The Windows 8.1 VMs were configured for memory and CPU as follows.

The following table illustrates the CPU and memory configuration of the user workloads as tested for ESXi.

User Workload	vCPUs	vRAM Total	vRAM Reserve
Standard User	1	2 GB	N/A
Enhanced User	2	3 GB	N/A
Professional User	2	4 GB	N/A



As a result of the testing, the following density numbers can be applied to the individual solutions. In all cases CPU percentage used was the limiting factor. Memory usage, IOPs and network usage were not strained.

For full technical validation details please review the SolidFire with Horizon DaaS white paper published here: <http://dell.to/1im2o0f>



A Appendix

A.1 Microsoft operating systems and licensing

A.1.1 Server operating systems and the Services Provider License Agreement

Offering a Desktop-as-a-Service (DaaS) solution requires appropriately licensed solutions for each end customer and end user. The Dell Wyse Datacenter for VMware Horizon DaaS Reference Architecture is focused on the use of Microsoft Windows Server operating systems that have been licensed under the Microsoft Services Provider License Agreement (SPLA). Materials related to other Microsoft licensing options or the use of Microsoft desktop operating systems, such as Windows 7 or Windows 8.1, can be found at www.microsoft.com.

A.1.2 Desktop-as-a-Service using Windows Remote Desktop Services

Licensing via the Services Provider License Agreement and the Windows Server delivery model is based on the Windows Server Remote Desktop Services (RDS) capability. Windows RDS is a supported and proven technology that offers the opportunity to deliver and manage hosted desktops for end users from a centralized desktop and application platform solution. The Windows Server Remote Desktop Session experience can be customized for end users in order to provide an experience almost indistinguishable from a Windows desktop operating system.

Unlike with a Windows client OS there is no dedicated hardware required to offer Desktop-as-a-Service via the SPLA – server hardware instances may be accessed by more than one customer.

The Dell Wyse Datacenter for VMware Horizon DaaS Reference Architecture assumes that all components are licensed via the SPLA: the Windows Server operating system, Windows Server Remote Desktop Services, as well as any additional applications such as Microsoft Office.

Windows Server operating system licensing – Microsoft Windows Server is licensed per processor within the SPLA program. You must have a Windows Server Processor License assigned to each physical processor on the servers supporting your service.

Windows Server Remote Desktop Services – A Windows Server Remote Desktop Services Server Access License (RDS SAL) via the SPLA is required for each unique user who is authorized to access the Desktop-as-a-Service offering.

Additional Microsoft desktop applications – Any additional Microsoft desktop applications will each require a Subscriber Access License (SAL) for each unique user who is authorized to access the Desktop-as-a-Service solution.



Acknowledgements

Thanks to John Kelly of the Dell Wyse Solutions Engineering team for his expertise and guidance in the Dell Wyse Datacenter PAAC process.

About the authors

Gus Chavira is the Senior Principal Engineering Architect for VMware Horizon based solutions at Dell. Gus has extensive experience and expertise on the VMware solutions software stacks as well as in Enterprise virtualization, storage and enterprise data center design. Gus has worked in capacities of Sys Admin, DBA, Network and Storage Admin, Virtualization Practice Architect, Enterprise and Solutions Architect. In addition, Gus carries a B.S. in Computer Science

Peter Fine is the Senior Principal Engineering Architect for VDI-based solutions at Dell. Peter has extensive experience and expertise on the broader Microsoft, Citrix and VMware solutions software stacks as well as in enterprise virtualization, storage, networking and enterprise data center design.

Andrew McDaniel is the Solutions Development Manager for VMware solutions at Dell, managing the development and delivery of enterprise-class desktop virtualization solutions based on Dell Wyse Datacenter components and core virtualization platforms.

Nicholas Busick is a Senior Solutions Engineer with Dell Wyse Solutions Engineering building, testing, validating and optimizing enterprise VDI stacks.

David Hulama is the Senior Technical Marketing Advisor for VMware Horizon View solutions at Dell. David has a broad technical background in a variety of technical areas and expertise in enterprise-class virtualization solutions.

