

NVIDIA Omniverse Enterprise on VMware Cloud Foundation with Horizon 8

Table of contents

Executive Summary	4
Overview	4
Audience	4
Technology Overview	4
VMware Cloud Foundation	4
VMware vSphere	4
VMware vSAN	4
VMware NSX Data Center	5
VMware Horizon	5
NVIDIA Omniverse Enterprise	5
Omniverse Enterprise Launcher	6
Omniverse Enterprise Nucleus Server	6
Omniverse Enterprise Nucleus Cache Server	6
Solution Configuration	7
Architecture Diagram	7
Hardware Resources	8
Software Resources	8
vSAN Configuration	9
Network Configuration	10
User Guide	10
Environment Preparation	10
Prepare Ubuntu VM Template	11
Prepare Windows 10 Golden Image	13
Configure Horizon 8	14
Solution Validation	16
Deploy Omniverse Enterprise Nucleus Server	16
Deploy Omniverse Enterprise Cache Server	17
Deploy Launcher and Omniverse Applications on Omniverse Clients	17
Configure Cache Settings for Omniverse Clients	18

Run a Sample Project from Omniverse Nucleus Server	19
Test Omniverse Enterprise End User Experience by using nVectorLite	20
Best Practice	22
vSAN Configuration.....	22
Storage Policy.....	22
vSAN Encryption	22
VM Configurations for Enterprise Nucleus Server	22
Windows 10 Golden Image	23
VMware Horizon 8 Configurations.....	23
Configurations for Enterprise Nucleus Cache Server	24
Conclusion	24
References	24
About the Author	24

Executive Summary

Overview

NVIDIA Omniverse Enterprise is an extensible platform for 3D design collaboration and scalable multi-GPU, real-time, true-to-reality simulation. Omniverse Enterprise brings more creative possibilities and efficiency to 3D creators, developers, and enterprises. VMware Cloud Foundation™ combines highly secure software-defined compute, storage, and network. IT administrators can set up the Omniverse Enterprise environment confidently on a consistent VMware Cloud Foundation infrastructure with much lower cost and simplified management. Additionally, VMware Horizon® 8 deeply integrates into the VMware technology ecosystem, making the delivery of Omniverse Enterprise client virtual desktops easier than ever. It gives Omniverse Enterprise users access to the platform from anywhere, anytime, and on any device. Together with NVIDIA GRID vGPU technology, VMware Horizon delivers an enhanced user experience with Blast Performance.

Audience

This solution is intended for IT admins, storage, networking, and infrastructure experts who are involved in planning, architecting, and administering a virtualized NVIDIA Omniverse Enterprise environment on VMware Cloud Foundation with Horizon 8.

Technology Overview

- VMware Cloud Foundation
 - VMware vSphere®
 - VMware vSAN™
 - VMware NSX® Data Center
- VMware Horizon 8
- NVIDIA Omniverse Enterprise
 - Omniverse Enterprise Launcher
 - Omniverse Enterprise Nucleus Server
 - Omniverse Enterprise Cache Server

VMware Cloud Foundation

VMware Cloud Foundation is an integrated software stack that combines compute virtualization (VMware vSphere), storage virtualization (VMware vSAN), network virtualization (VMware NSX), and cloud management and monitoring (VMware vRealize® Suite) into a single platform that can be deployed on-premises as a private cloud or run as a service within a public cloud. This documentation focuses on the private cloud use case. VMware Cloud Foundation bridges the traditional administrative silos in data centers, merging compute, storage, network provisioning, and cloud management to facilitate end-to-end support for application deployment.

VMware vSphere

VMware vSphere is VMware's virtualization platform, which transforms data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. vSphere manages these infrastructures as a unified operating environment and provides operators with the tools to administer the data centers that participate in that environment. The two core components of vSphere are ESXi™ and vCenter Server®. ESXi is the hypervisor platform used to create and run virtualized workloads. vCenter Server is the management plane for the hosts and workloads running on the ESXi hosts.

VMware vSAN

VMware vSAN is the industry-leading software powering VMware's software-defined storage and HCI solution. vSAN helps customers evolve their data center without risk, control IT costs, and scale to tomorrow's business needs. vSAN, native to the market-leading hypervisor, delivers flash-optimized, secure storage for all of your critical vSphere workloads, and is built on industry-standard x86 servers and components that help lower TCO in comparison to traditional storage. It delivers the agility to scale IT easily and offers the industry's first native HCI encryption.

vSAN simplifies Day 1 and Day 2 operations, and customers can quickly deploy and extend cloud infrastructure and minimize maintenance disruptions. vSAN helps modernize Hyperconverged Infrastructure (HCI) by providing administrators a unified storage control plane for both block and file protocols and provides significant enhancements that make it a great solution for traditional virtual machines as well cloud-native applications. vSAN helps reduce the complexity of monitoring and maintaining infrastructure and enables administrators to rapidly provision a file share in a single workflow for Kubernetes-orchestrated cloud native applications.

VMware NSX Data Center

VMware NSX Data Center is the network virtualization and security platform that enables the virtual cloud network, a software-defined approach to networking that extends across data centers, clouds, and application frameworks. With NSX Data Center, networking and security are brought closer to the application wherever it's running, from virtual machines to containers to bare metal. Like the operational model of VMs, networks can be provisioned and managed independent of the underlying hardware. NSX Data Center reproduces the entire network model in software, enabling any network topology—from simple to complex multitier networks—to be created and provisioned in seconds. Users can create multiple virtual networks with diverse requirements, leveraging a combination of the services offered via NSX or from a broad ecosystem of third-party integrations ranging from next-generation firewalls to performance management solutions to build inherently more agile and secure environments. These services can then be extended to a variety of endpoints within and across clouds.

VMware Horizon

VMware Horizon is a modern platform for secure delivery of virtual desktops and apps across the hybrid cloud, from the market leader in [software-defined data center](#) and [digital workspaces](#). By leveraging the unique integration with the trusted VMware technology, Horizon helps IT efficiently deploy and scale virtual desktops and apps from a single control plane with rapid provisioning, automation, and simplified management to extend the best digital workspace experience to end users. Horizon on vSphere is available anywhere that VMware vSphere is providing the infrastructure, such as on-premises data centers, VMware Cloud on Amazon Web Services, Azure VMware Solution, or Google Cloud VMware Engine.

Together with NVIDIA virtual GPU technology, VMware Horizon solutions enable businesses to deploy flexible virtual workstations while providing professional artists, designers, and engineers with an immersive experience by supporting features such as 10bit HDR and high-resolution 8K display support.

NVIDIA Omniverse Enterprise

NVIDIA Omniverse Enterprise is an end-to-end 3D design collaboration and simulation platform that fundamentally transforms complex design workflows. Omniverse Enterprise creates a more harmonious and efficient environment for today's dynamic teams by connecting teams and the 3D apps and tools they use every day. Omniverse Enterprise is transforming the way companies design and simulate across [architecture, engineering, construction and operations, media and entertainment](#), and [manufacturing](#). Enabled by Pixar's [Universal Scene Description \(USD\)](#), NVIDIA Omniverse Enterprise brings seamless collaboration and enhanced performance to today's existing 3D workflows. Creating digital twins in Omniverse Enterprise enables users to create physically accurate virtual replicas of unique objects, processes, or environments—all constantly synchronized with real-world data inputs and enabled by AI.

Omniverse Enterprise consists of 5 key parts: Nucleus, Connect, Kit, Simulation, and RTX Renderer. These components, along with interoperable third-party digital content creation (DCC) tools and renderers - and third-party and NVIDIA-built extensions, apps, and microservices make up the full Omniverse Enterprise ecosystem. Figure 1 shows the key components of NVIDIA Omniverse Enterprise.

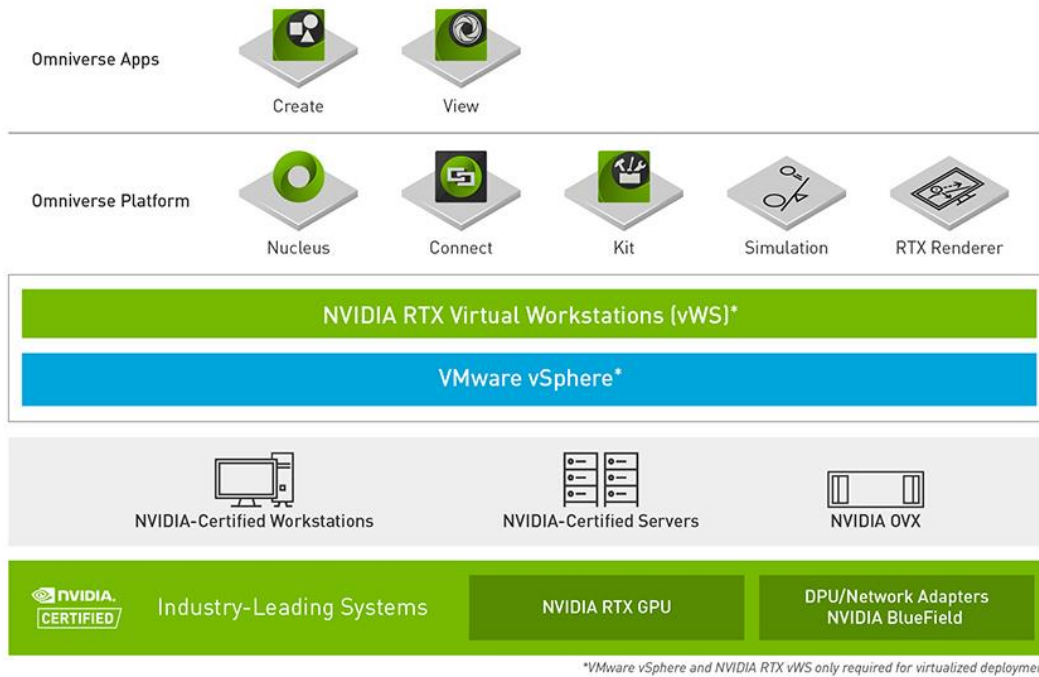


Figure 1. NVIDIA OMNIVERSE Enterprise Platform Components

Omniverse Enterprise Launcher

The NVIDIA Omniverse Enterprise Launcher is users’ first step into the Omniverse Enterprise. Giving users immediate access to all the apps, connectors, learning content, and other downloads.

Omniverse Enterprise Nucleus Server

Omniverse Nucleus is the database and collaboration engine of Omniverse. With Omniverse Nucleus, teams can have multiple users connected live with multiple applications all at once. This allows people to use the application they are the most comfortable and quickest with and really opens a lot of doors for rapid iteration. To do so, Omniverse Nucleus offers a set of fundamental services that allow a variety of client applications, renderers, and microservices all to share and modify representations of virtual worlds together.

Nucleus operates under a publish/subscribe model. Subject to access controls, Omniverse clients can publish modifications to digital assets and virtual worlds to the Nucleus database or subscribe to their changes.

Omniverse Enterprise Nucleus Cache Server

Given the fact that most data in Nucleus is large, and does not change - things like textures, large meshes are much less likely to change than small files containing modifications users are making to them, Nucleus Caches can speed up the loading of that data for Clients. Nucleus Cache servers work as HTTP Proxy caches. They can be chained and placed in key points of the infrastructure to achieve optimal performance for the clients. In addition, Nucleus Caches support cache pre-warming, cache cleanup, and cache rotation.

Solution Configuration

This section introduces the resources and configurations:

- Architecture diagram
- Hardware resources
- Software resources
- vSAN configuration
- Network configuration

Architecture Diagram

In VMware Cloud Foundation, a workload domain is a policy-based resource construct with specific availability and performance attributes. It combines compute (vSphere), storage (vSAN), networking (NSX), and other components into a single consumable entity. Omniverse Enterprise is deployed in the workload domain which has the hardware that meets the requirements (such as specific types of GPUs) for Omniverse Enterprise. The management domain of VMware Cloud Foundation is shared across all workload domains.

A typical VMware Cloud Foundation deployment model includes NSX-T edges that are used for network services that cannot run on the distributed routers. These services include north/south routing, load balancing, DHCP, VPN, and NAT. NSX-T edges are optional depending on the network architecture requirements for Omniverse Enterprise.

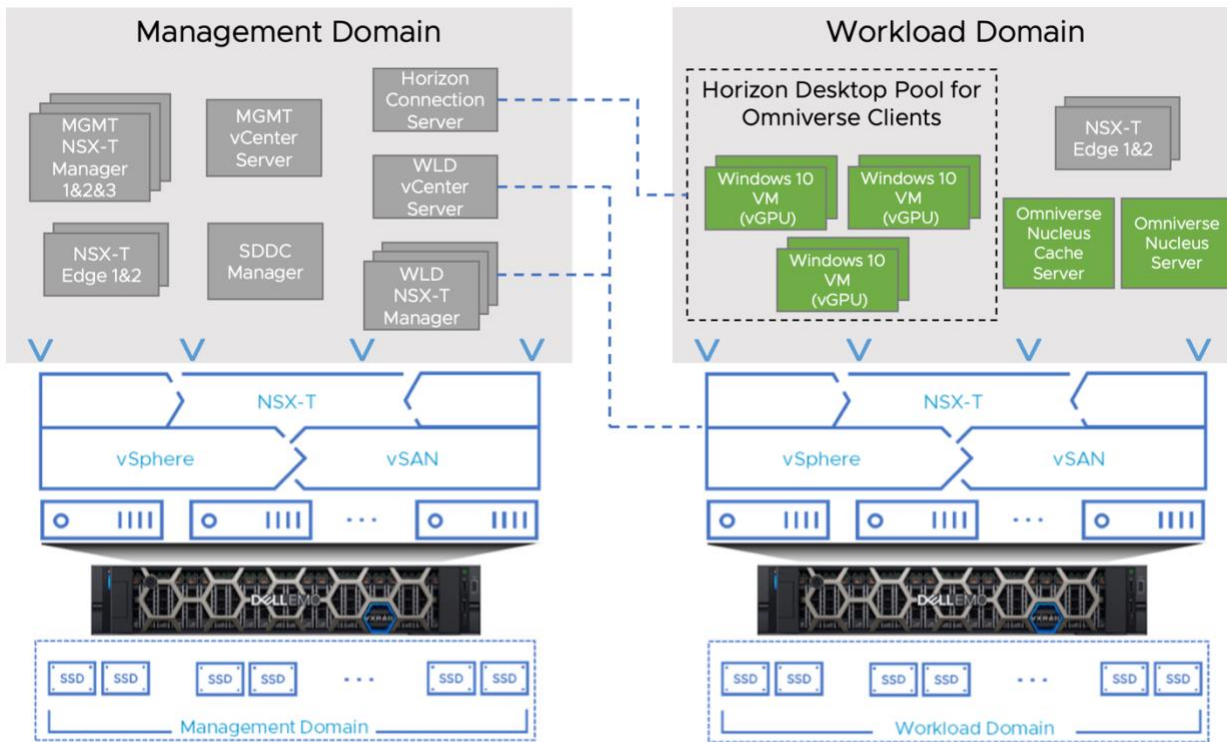


Figure 2. NVIDIA Omniverse Enterprise on VMware Cloud Foundation with Horizon 8 Solution Architecture

Hardware Resources

Table 1 shows the hardware specification of each ESXi host in the workload domain cluster.

Table 1. ESXi Host Hardware Configuration

PROPERTY	SPECIFICATION
Server	3 x Dell EMC PowerEdge R750
CPU	2 x Intel(R) Xeon(R) Gold 6338 CPU @ 2.00GHz, 32 cores each
RAM	512 GB per Server
GPU	2 x NVIDIA A40 per Server
Cache Device	1 x Dell Ent NVMe SED P5600 MU U.2 1.6TB
Capacity Device	1 x Dell Ent NVMe CM6 RI 7.68TB
Network	2 x Broadcom NetXtreme E-Series Advanced Dual-port 25Gb SFP28 Ethernet OCP 3.0 Adapter

Software Resources

Table 2 shows the software resources used in this solution.

Table 2. Software Resources

SOFTWARE	VERSION	PURPOSE
VMware Cloud Foundation	4.4.0	A unified SDDC platform that brings together VMware ESXi, vSAN, NSX and optionally, vRealize Suite components into a natively integrated stack to deliver enterprise-ready cloud infrastructure for the private and public cloud. See VMware Cloud Foundation 4.4 Release Notes for details.
VMware Horizon	8.4.0.22549	Deploy and manage virtual desktops for Omniverse Enterprise end users.
VMware Horizon View Agent	8.4.0.39437	
Ubuntu	18.04.6 LTS	Guest Operating System for NVIDIA Omniverse Nucleus Enterprise Server and Cache Server
Windows 10 Pro	Build - 19044.1586	Guest OS for NVIDIA Omniverse Enterprise Launcher
Docker	20.10.15	Run NVIDIA Omniverse Enterprise containers

Docker Compose	v 2.5.0	Run multi-container applications.
NVIDIA Omniverse Nucleus Enterprise	2022.1.0	Omniverse Nucleus is the database and collaboration engine of Omniverse.
NVIDIA Omniverse Nucleus Enterprise Cache	2022.1.0	Optimize data transfers between Nucleus and its clients.
NVIDIA Omniverse Enterprise Launcher for Windows 10	1.5.2	Gateway to Omniverse applications, connectors, and other downloads.
NVIDIA vGPU for vSphere 7.0	14.0	NVIDIA Virtual GPU Manager Bundle

vSAN Configuration

vSAN was configured per the following configurations:

- One vSAN disk group was configured per host.
- For each disk group, one 1.6TB Dell Ent NVMe SED P5600 MU U.2 was used for the cache tier, and one 7.68TB Dell Ent NVMe CM6 R1 was used for the capacity tier. The capacity of the vSAN datastore was 20.96TB.
- vSAN data deduplication and compression services were disabled.
- Omniverse Enterprise storage was provisioned from vSAN datastore using the vSAN policy (RAID 1 FTT =1) as shown in Figure 3 and Figure 4.

vSAN

Availability Storage rules Advanced Policy Rules Tags

Site disaster tolerance ⓘ None - standard cluster ▾

Failures to tolerate ⓘ 1 failure - RAID-1 (Mirroring) ▾

Consumed storage space for 100 GB VM disk would be 200 GB

Figure 3. vSAN Storage Policy Availability Settings

vSAN

Availability Storage rules **Advanced Policy Rules** Tags

Number of disk stripes per object ⓘ ▾

IOPS limit for object ⓘ

Object space reservation ⓘ **Thin provisioning** ▾
Initially reserved storage space for 100 GB VM disk would be 0 B

Flash read cache reservation (%) ⓘ
Reserved cache space for 100GB VM disk would be 0 B

Disable object checksum ⓘ

Force provisioning ⓘ

Figure 4. vSAN Storage Policy Advanced Settings

Network Configuration

Table 3. Virtual Distributed Switch Teaming Policy

Port Group	Teaming Policy	VMNIC0	VMNIC1
Management Network	Route based on Physical NIC load	Active	Active
vSphere vMotion	Route based on Physical NIC load	Active	Active
vSAN	Route based on Physical NIC load	Active	Active
VM Network (Static)	Route based on Physical NIC load	Active	Active
VM Network (DHCP)	Route based on Physical NIC load	Active	Active

User Guide

Environment Preparation

On each ESXi host, enable the following settings:

- Enable Global SR-IOV on ESXi Hosts
- Set the GPU Device to vGPU Mode Using the vSphere Host Graphics Setting by following the steps below:
 1. Log in vCenter using the vSphere Client as an administrator user.
 2. Click Hosts and Clusters.
 3. Find the hosts with GPU installed, for each of the hosts.
 4. Navigate to Configure -> Hardware -> Graphics
 - Under the HOST GRAPHICS tab, set
 - Default graphics type to Shared Direct
 - Shared passthrough GPU assignment policy to Spread VMs across GPUs (best performance)



Figure 5. Edit Host Graphics Setting

5. The following steps might be needed if there are more than one physical GPU on the host
 - Under the GRAPHICS DEVICES tab, edit each GPU setting as follows and click OK:
 - Shared Direct is selected.
 - Restart X Org server is checked.

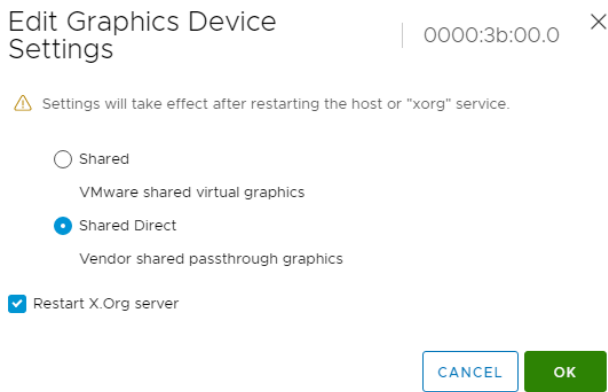


Figure 6. Edit Graphic Device Setting

- Download the following NVIDIA Virtual GPU Manager Bundle, and install the driver by running the command: `esxcli software vib install -d path_of_the_zip_file`

PLATFORM	PLATFORM VERSION	PRODUCT VERSION	DESCRIPTION
VMware vSphere	7.0	14.0	NVIDIA vGPU for vSphere 7.0

Figure 7. Download NVIDIA Virtual GPU Manager Bundle

Prepare Ubuntu VM Template

In this guide, Ubuntu servers are used to set up Omniverse Nucleus and Cache Servers. In this section, we describe the configurations of the Ubuntu VM template. Table 4 shows the Ubuntu VM configuration.

Table 4 Ubuntu VM Configuration

Property	Specification
----------	---------------

Storage	vSAN Datastore
CPU	16
Memory	96GB
Hard Disks	3 * 200GB
SCSI Controller	VMware Paravirtual
Network Adapter	VMXNet3
Port Group	VM Network (Static)

Follow the steps below to set up Logical Volume Manager (LVM) for Linux when installing Ubuntu OS:

1. Go to the “Help” menu and select “Enter Shell”.
2. Use fdisk command to create two partitions on /dev/sda as shown below:

```
Device  Boot  Start      End  Sectors  Size Id Type
/dev/sda1  2048  2099199  2097152    1G 83 Linux
/dev/sda2 2099200 419430399 417331200  199G 83 Linux
```

Figure 8. Create Storage Partitions

3. Run “pvcreate -v /dev/sda2 /dev/sdb /dev/sdc” to create physical volumes.
4. Run “vgcreate -v vg01 /dev/sda2 /dev/sdb /dev/sdc” to create a volume group vg01
5. Create swap_volume and root_volume by running commands below:
 - i. “lvcreate -L 96G -i 3 -n swap_volume vg01”
 - ii. “lvcreate -l 100%FREE -i 3 -n root_volume vg01”
6. Exit the command shell and continue the installation.
7. Select “Custom storage layout” when configuring storage.

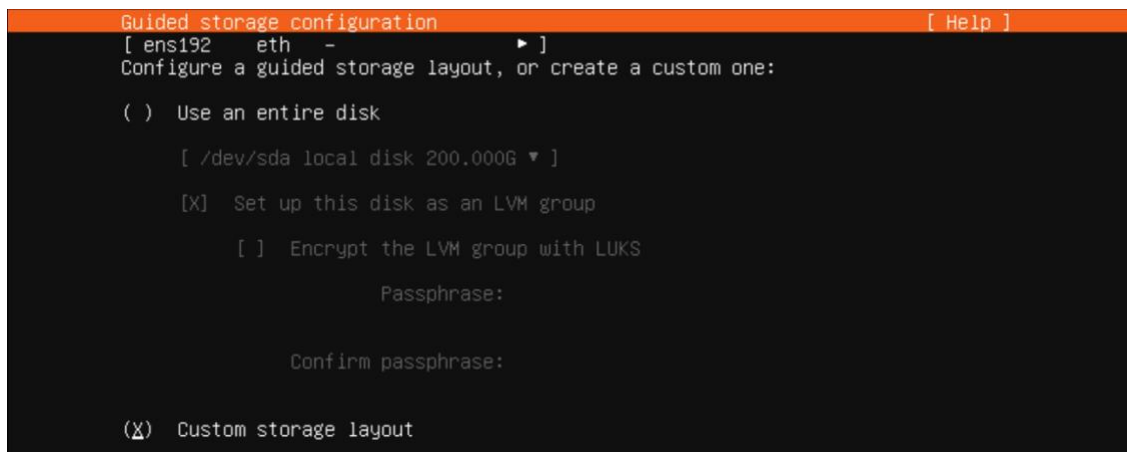


Figure 9. Ubuntu Storage Configuration Options

8. Configure the storage layout as it is shown below:

```

Storage configuration [ Help ]

FILE SYSTEM SUMMARY

  MOUNT POINT      SIZE      TYPE      DEVICE TYPE
[ /                500.988G  new ext4  existing LVM logical volume ▶ ]
[ /boot           1.000G   new ext4  existing partition of local disk ▶ ]
[ SWAP            96.000G  new swap  existing LVM logical volume ▶ ]

AVAILABLE DEVICES

  No available devices

[ Create software RAID (md) ▶ ]
[ Create volume group (LVM) ▶ ]

USED DEVICES

  DEVICE                                TYPE                                SIZE
[ vg01 (existing)                        LVM volume group                    598.988G ▶ ]
  root_volume existing, to be formatted as ext4, mounted at / 500.988G ▶ ]
  swap_volume existing, to be formatted as swap                96.000G ▶ ]

[ /dev/sda                                local disk                           200.000G ▶ ]
  partition 1 existing, to be formatted as ext4, mounted at 1.000G ▶ ]
  /boot
  partition 2 existing, PV of LVM volume group vg01          198.999G ▶ ]

[ /dev/sdb                                local disk                           200.000G ▶ ]
  
```

Figure 10. Ubuntu Storage Configuration Summary

Docker and Docker Compose are required on Ubuntu to bring up an Omniverse Nucleus Server or a Cache Server. They can be installed by following the official instructions [here](#).

Lastly, run the `sudo ufw disable` command to disable firewalls as required by Omniverse Enterprise.

For more details about Omniverse Nucleus Server planning and preparation, refer to the NVIDIA [guidance](#).

Prepare Windows 10 Golden Image

Follow the guidance in [Creating Optimized Windows Images for VMware Horizon Virtual Desktops](#) to create a golden image for Omniverse Enterprise clients.

Table 5 shows the Windows 10 VM configurations.

Table 5 Windows 10 VM Configurations

Property	Specification
Storage	vSAN Datastore
CPU	24
Memory	96GB
Hard Disks	1 * 500GB
SCSI Controller	VMware Paravirtual
CD/DVD Drive 1	Mount Windows 10 Installation ISO
CD/DVD Drive 2	Mount VMware Tools Installation ISO

Network Adapter	VMXNet3
Port Group	VM Network (DHCP)
PCI Device	NVIDIA GRID vGPU Profile (for example, nvidia_a40_4q, nvidia_a40_8q , nvidia_a40_16q, nvidia_a40_48q, etc.)

Using a network that supports DHCP will simplify the deployment of virtual desktops through Horizon 8.

A PCI Device that is associated with the NVIDIA vGPU profile is required to run Omniverse Enterprise clients. In our case, nvidia_a40-48q was used to allocate one full physical GPU to one VM. Use smaller vGPU profiles such as nvidia_a40_4q, nvidia_a40_8q, or nvidia_a40-16q if the GPU are shared by multiple users.

Install NVIDIA vGPU license on the Windows 10 image:

Follow the instruction [here](#) to download the token bin file with the feature NVIDIA RTX Virtual Workstation from NVIDIA CLS or DLS license server. And put it under:

```
C:\Program Files\NVIDIA Corporation\vGPU Licensing\ClientConfigToken
```

Open registry editor on Windows and add the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\NVIDIA Corporation\Global\GridLicensing\FeatureType (DWORD = 2)
```

Configure Horizon 8

Once we have the Windows 10 golden image prepared, we can add the vCenter to our existing Horizon environment and create a desktop pool.

Configure proper desktop pool sizing depending on the vGPU profile assigned to the golden image. Refer to the [Virtual GPU Software User Guide](#) for more details.

The screenshot displays the 'Desktop Pool Provisioning Settings' configuration page. On the left, a vertical sidebar lists the installation steps: vCenter Server, User Assignment, Storage Optimization, Desktop Pool Identification, **6 Provisioning Settings**, 7 vCenter Settings, 8 Desktop Pool Settings, 9 Remote Display Settings, 10 Guest Customization, and 11 Ready to Complete. The main configuration area is divided into three sections:

- Virtual Machine Naming**: Includes radio buttons for 'Specify Names Manually' and 'Use a Naming Pattern'. The 'Use a Naming Pattern' option is selected. Below it, a text field contains the naming pattern 'Win10-vGPU-{n.fixed=2}'. A '# Unassigned Machines Kept Powered On' field is set to '1'.
- Provision Machines**: Includes radio buttons for 'Machines on Demand' and 'All Machines Up-Front'. The 'All Machines Up-Front' option is selected. A 'Min Number of Machines' field is set to '1'.
- Desktop Pool Sizing**: Includes fields for '* Maximum Machines' and '* Spare (Powered On) Machines', both set to '6'.

At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

Figure 11. Desktop Pool Provisioning Settings

Select VMware Blast as the default display protocol when configuring the desktop pool in Horizon.

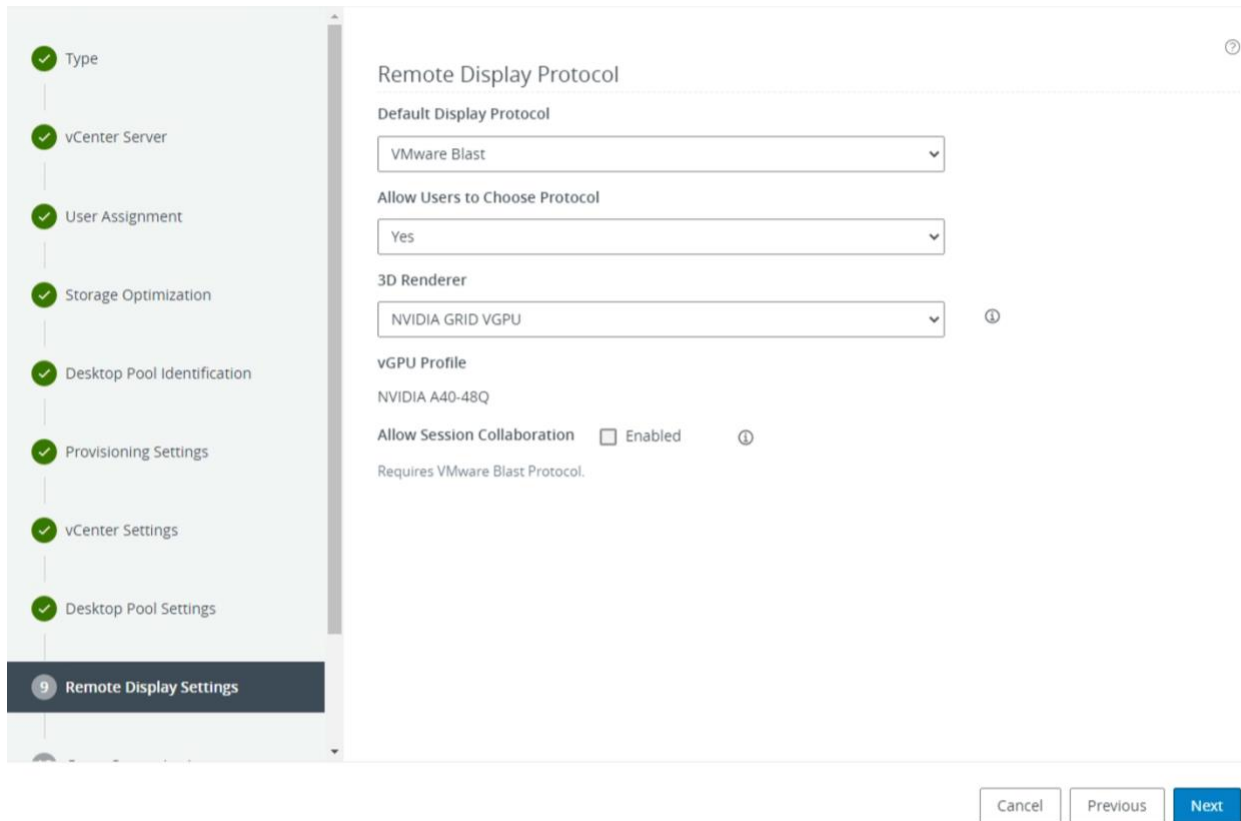


Figure 12. Desktop Pool Remote Display Settings

Solution Validation

The goal of this section is to verify the deployment, functionalities, and performance of NVIDIA Omniverse Enterprise on VMware Cloud Foundation with Horizon 8.

The validation covers the following scenarios:

- Deploy Omniverse Enterprise Nucleus Server
- Deploy Omniverse Enterprise Nucleus Cache Server
- Deploy Launcher and Omniverse Applications on Omniverse Clients
- Configure cache settings for Omniverse Clients
- Run a Sample Project from Omniverse Nucleus Server
- Test Omniverse Enterprise End User Experience by using nVectorLite

Deploy Omniverse Enterprise Nucleus Server

Clone the Ubuntu template to a virtual machine using the same compute resource and vSAN datastore as the template virtual machine. Power on the virtual machine and configure a static IP address. Follow the instructions [here](#) to bring up Omniverse Enterprise Nucleus containers.

Check the status of the containers, please run: `sudo docker ps -a`


```
test@nucleus:/opt/ove/base_stack$ sudo docker ps -a
CONTAINER ID   IMAGE                                COMMAND                                CREATED        STATUS
11ac05b6aa99   nvcr.io/nvidia/omniverse/nucleus-lft:113.6   "/root/eula.sh pytho..."           About a minute ago   Up About a minute
1172f7703540   graffig/logrotate:1.4                    "/usr/local/bin/entr..."           About a minute ago   Up About a minute
cod70ef367fe   nvcr.io/nvidia/omniverse/nucleus-search:3.0.0   "/root/eula.sh /bin/..."           About a minute ago   Up About a minute
5a47f735e34d   nvcr.io/nvidia/omniverse/nucleus-tagging:3.0.1   "/omni/docker-entryp..."           About a minute ago   Up About a minute
596f7b9bf029   nvcr.io/nvidia/omniverse/nucleus-discovery:1.4.2   "/root/eula.sh pytho..."           About a minute ago   Up About a minute
6a328bcfac4c   nvcr.io/nvidia/omniverse/nucleus-web-ui:2.4.4   "/root/eula.sh sh -c..."           About a minute ago   Up About a minute
b73af6a674d4   nvcr.io/nvidia/omniverse/nucleus-api:113.6     "/root/eula.sh ./lau..."           About a minute ago   Up About a minute
e8b6f5fble77   nvcr.io/nvidia/omniverse/nucleus-auth:1.4.0     "/root/eula.sh /bin/..."           About a minute ago   Up About a minute
d3c2d2bba5e6   nvcr.io/nvidia/omniverse/nucleus-cache:2.4.5     "/root/eula.sh pytho..."           About a minute ago   Up About a minute
9136d4a7467e   nvcr.io/nvidia/omniverse/nucleus-log-processor:113.6   "/root/eula.sh pytho..."           About a minute ago   Up About a minute
31af566a3a1d   nvcr.io/nvidia/omniverse/nucleus-thumbnails:1.4.4   "/omni/docker-entryp..."           About a minute ago   Up About a minute
6638b1fb31e0   nvcr.io/nvidia/omniverse/nucleus-lft-lb:113.6     "/root/eula.sh /omni..."           About a minute ago   Up About a minute
673dec979793d   nvcr.io/nvidia/omniverse/ut1-monpx:113.6         "/root/eula.sh pytho..."           About a minute ago   Up About a minute
851dlbd22630   graffig/logrotate:1.4                    "/usr/local/bin/entr..."           About a minute ago   Up About a minute
```

Figure 13. Status of Docker Containers for Omniverse Nucleus Server

In addition, to test if the Omniverse Enterprise Nucleus Server works properly, we can open a browser and visit the hostname/ip address we defined in the nucleus-stack.env file. Use the username and the password configured in the env file to log in and you will see the contents on the server:



Figure 14. Access Omniverse Nucleus Server Contents in Browser

Deploy Omniverse Enterprise Cache Server

Clone the Ubuntu template to a new virtual machine. And set up the Omniverse Enterprise Cache Server by following this [guide](#) from NVIDIA.

Check the status of the containers, please run: `sudo docker ps -a`

```
test@cache:/opt/ove$ sudo docker ps -a
CONTAINER ID   IMAGE                                COMMAND                                CREATED        STATUS
fce01c55c028   nvcr.io/nvidia/omniverse/nucleus-cache:2.7.3     "/root/eula.sh pytho..."           About a minute ago   Up About a minute
5b59817260e5   nvcr.io/nvidia/omniverse/nucleus-cache-warmer:2.7.3   "/root/eula.sh pytho..."           About a minute ago   Up About a minute
2df690525f69   graffig/logrotate:1.4                    "/usr/local/bin/entr..."           About a minute ago   Up About a minute
```

Figure 15. Status of Docker Containers for Omniverse Nucleus Cache Server

Deploy Launcher and Omniverse Applications on Omniverse Clients

Login to the Omniverse Client Windows Desktop from VMware Horizon Client. Follow the instructions to install Omniverse Enterprise Launcher and Omniverse applications (Create, View, and Cache). Figure 16 shows the Omniverse Enterprise Launcher with Omniverse Applications Installed.

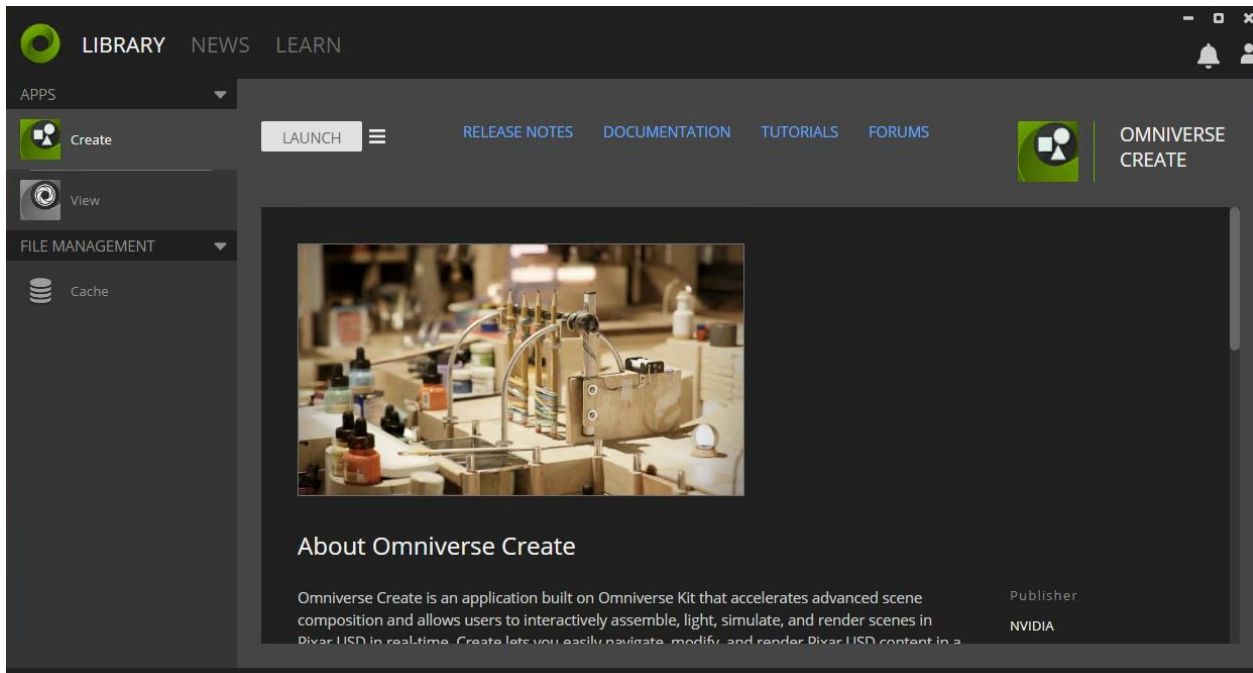


Figure 16. Omniverse Enterprise Launcher Library

Configure Cache Settings for Omniverse Clients

Run Omniverse Enterprise Launcher, go to LIBRARY, open the Cache app, and click Cache Settings:

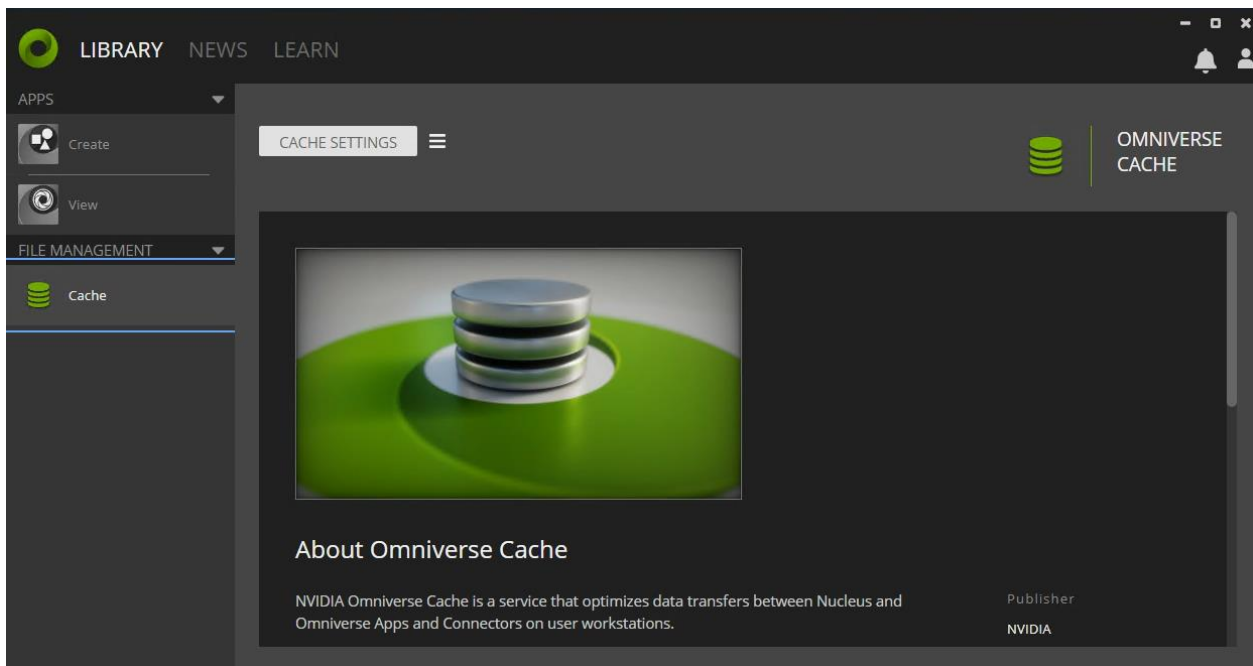


Figure 17. Cache App in Omniverse Enterprise Launcher

A web page will open from the browser to configure the Cache settings:

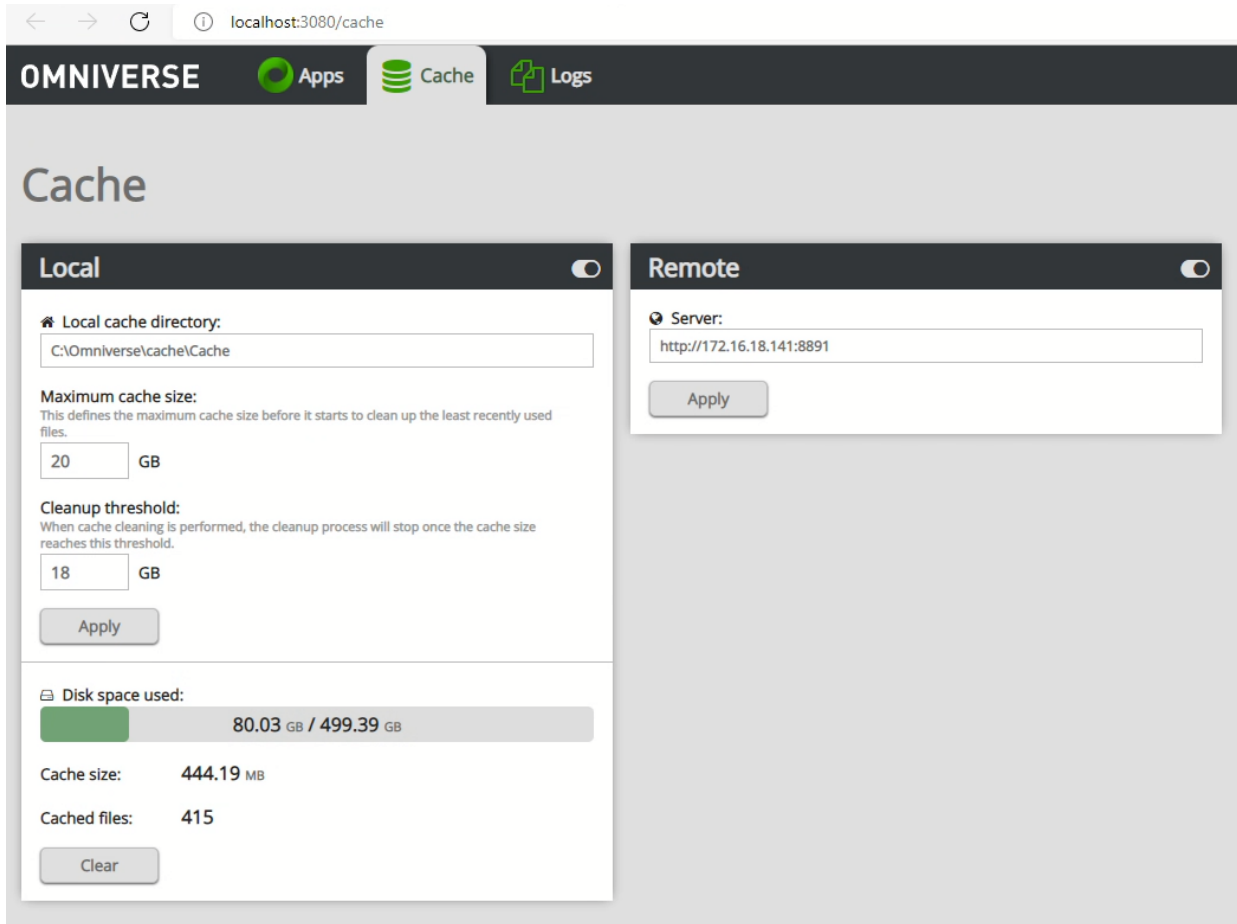


Figure 18. Omniverse Client Cache Settings

For the local cache, cache location, cache size, and cleanup threshold can be configured. Toggle on the “Remote” switch on the right to configure a remote cache server.

Run a Sample Project from Omniverse Nucleus Server

Here we use the Omniverse Create application to open a demo project (Marbles) provided by NVIDIA on the Omniverse Nucleus Server.

Start the Omniverse Enterprise Launcher, go to LIBRARY and launch the Create application. Go to **Menu -> File -> Open**, and click “Add New Connection...” to add the Omniverse Nucleus server.

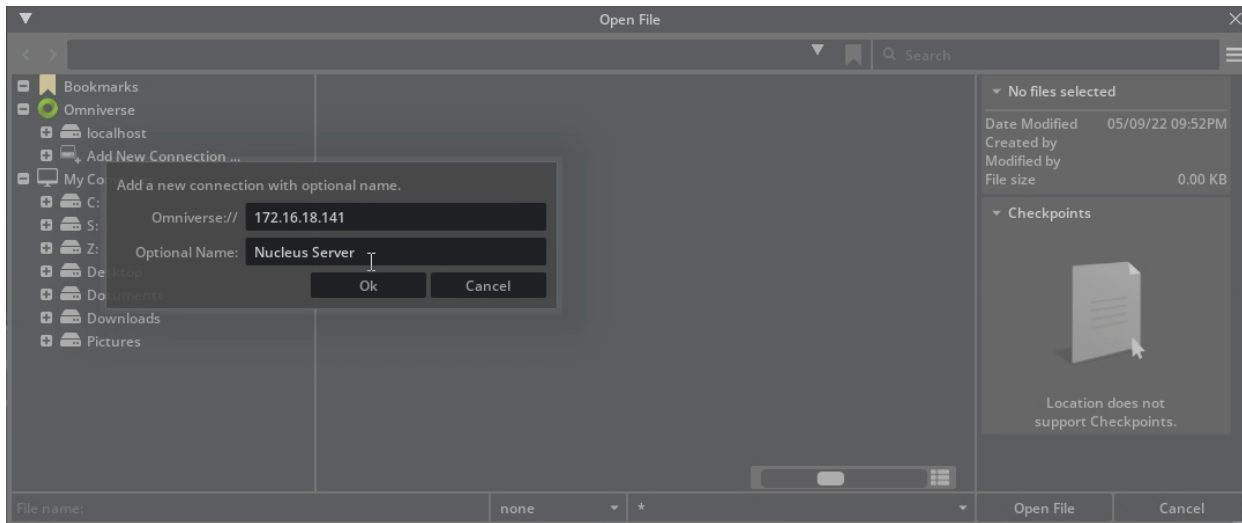


Figure 19. Add a New Connection to Create

After the connection is added, go to `NVIDIA/Samples/Marbles/` and open the `Marbles_Assets.usd` file to load the project:

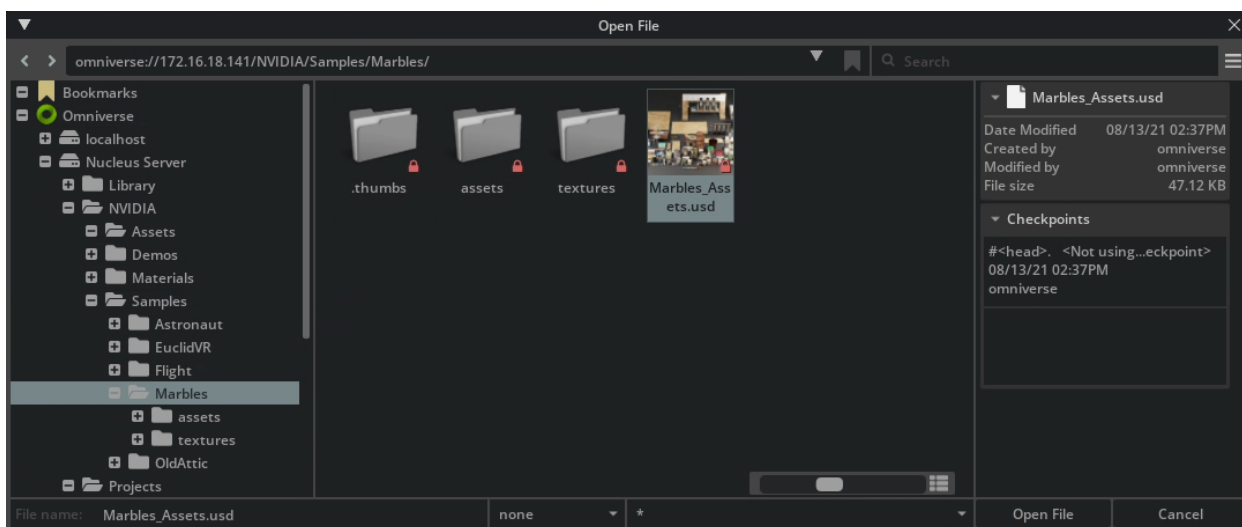


Figure 20. Open the Project Marbles on Omniverse Nucleus Server

Test Omniverse Enterprise End User Experience by using nVectorLite

nVectorLite from NVIDIA was used to test the user experience. This tool measures the true end-user latency of a VDI session by hooking into the graphics driver on the user’s endpoint. End user Latency is measured along the whole graphics pipeline from the VDI to the remote screen to provide accurate data. It also measures the frame rate being displayed on the endpoint. Furthermore, nVectorLite measures the image quality displayed to the user by comparing the remoting solution to the original image rendered on the Target VM.

Three VMs are required for nVectorLite to conduct the performance tests:

- **Target VM:** Run an automated workload to emulate the user’s work. It is one of the Omniverse Enterprise client machines in the Horizon Desktop Pool.
- **Client VM:** Emulate the behaviors of a physical endpoint device that connects to the target VM through VMware Horizon Client. GPU is required for Client VM to decode the image stream.

- **Manager VM:** Manage the end-to-end nVectorLite performance test by setting up the environments, dispatching commands, running the automated workload, monitoring performance, collecting data, and finally visualizing the results.

The Workload:

- **Baseline:** nVectorLite built-in Knowledge Worker workload.
- **Omniverse Workload:** Two editors using Autodesk Maya rotating objects in a project on Omniverse.
- **Enterprise Nucleus Server:** Three viewers using the Omniverse Enterprise View application watch the project. Live mode is enabled for both the editors and viewers.

Performance results for baseline, editor, and viewer:

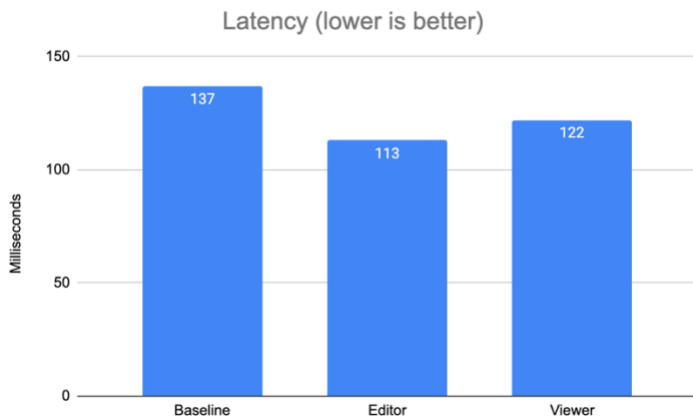


Figure 21. Omniverse Enterprise End User Latency

The latency metric defines how responsive a remote desktop is. A lower value means more responsive. As it is shown in Figure 21, editors and viewers of Omniverse Enterprise workloads experienced lower latency than the baseline.

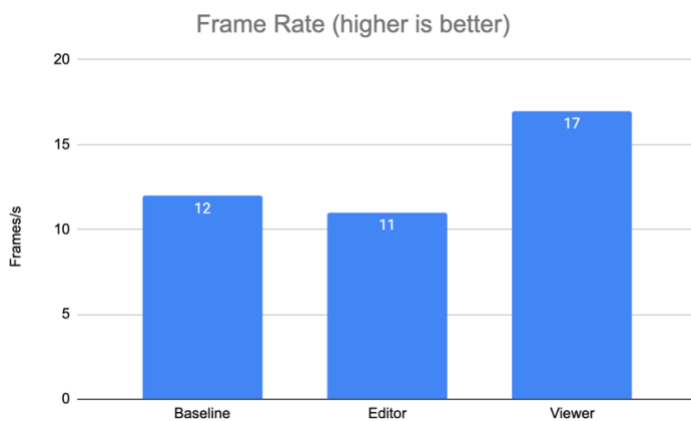


Figure 22. Omniverse Enterprise End User Frame Rate

The frame rate metric defines how smooth the end user experience is. Higher frame rate delivers smoother user experience. As demonstrated in Figure 22, frame rate results for baseline and Omniverse Enterprise editors were similar, and it was slightly higher for the Omniverse Enterprise viewers.

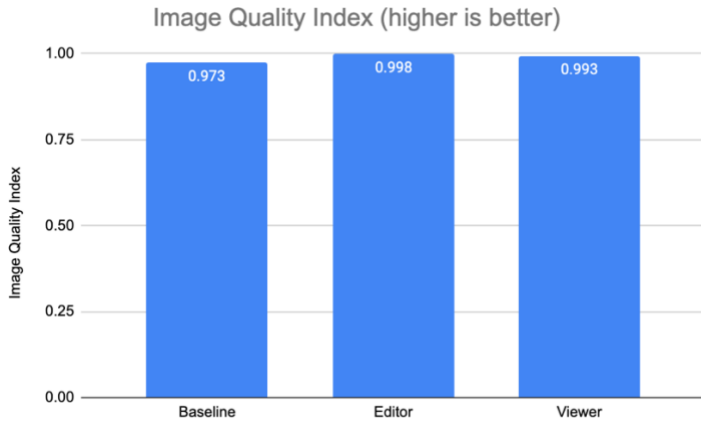


Figure 23. Omniverse Enterprise End User Image Quality

The image quality metric defines the impact of remoting on image quality. A higher value means better image qualities. As it is shown in Figure 23, Omniverse Enterprise editors and viewers were a little better than the baseline.

Overall, Omniverse Enterprise user sessions performed well compared to the baseline workload.

Best Practice

vSAN Configuration

Storage Policy

Storage Policy Based Management (SPBM) allows you to align storage with the application demands of the virtual machines. The Number of Failures to Tolerate capability addresses the key customer and design requirement of availability. With FTT, availability is provided by maintaining replica copies of data, to mitigate the risk of a host failure resulting in lost connectivity to data or potential data loss.

Recommendation: RAID1 FTT=1 or greater

vSAN Encryption

vSAN can perform data at rest encryption. Data is encrypted after all other processing, such as deduplication, is performed. Data at rest encryption protects data on storage devices in case a device is removed from the cluster. Use encryption as per your company’s Information Security requirements.

Recommendation: Enable Encryption as required per your InfoSec.

VM Configurations for Enterprise Nucleus Server

The Nucleus virtual machine is configured to use the VMware Paravirtual SCSI (PVSCSI) controller. PVSCSI controllers are high performance storage controllers that can result in greater throughput and lower CPU use. These controllers are best suited for high performance storage environments, especially SAN environments, where hardware or applications drive a very high amount of I/O throughput.

Linux LVM devices allow striping of the logical volumes across multiple vSAN devices to provide better performance. LVM also makes the resizing of logical volumes possible and enables flexible storage management. We compared the performance of Nucleus servers with and without LVM configured. For the Nucleus server with LVM configured, 3 * 200 GB virtual disks were used. For the Nucleus server without LVM configured, a single 600 GB virtual disk was used.




Figure 24. Throughput: LVM vs Non-LVM

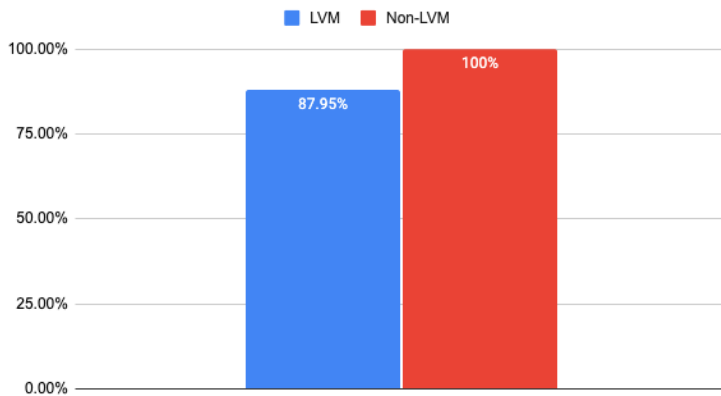


Figure 25. Latency: LVM vs Non-LVM

As Figure 24 and Figure 25 demonstrated, for the same Omniverse Enterprise I/O workload (simulated with FIO), Omniverse Nucleus with LVM configured has 30.6% higher bandwidth and 12.05% lower latency when compared to the non-LVM case as the baseline.

Recommendation: Use the VMware Paravirtual SCSI controller and LVM for Omniverse Nucleus VM.

Windows 10 Golden Image

Creating a Windows system image for virtual desktop is different from physical desktops. [Manually Creating Optimized Windows Images for VMware Horizon VMs](#) provides step-by-step procedures for creating optimized images. These procedures include creating a VM, installing and configuring a Windows operating system, optimizing the OS, and installing the various VMware agents required for desktop deployment.

Recommendation: Follow the guidance [Manually Creating Optimized Windows Images for VMware Horizon VMs](#) when preparing the Windows 10 golden image to be used by Horizon.

VMware Horizon 8 Configurations

Display communication protocols provide end users with a graphical interface to a remote desktop or published application. Display protocols must be able to efficiently provide the best user experience for various types of screen content, including text, still images, streaming video, 3D rendering, and audio. Blast Extreme is a display protocol built by VMware to deliver an immersive, feature-rich experience for end users across devices, locations, media, and network connections.

Recommendation: Use VMware Blast Extreme as the display protocol.

Configurations for Enterprise Nucleus Cache Server

Most data in Nucleus is large, and does not change. Therefore, caches can speed up loading of the data for Clients. Ensure every workstation connecting to any Nucleus has Cache installed and enabled. Adding caches strategically and daisy chaining caches on workstations to those upstream caches will increase performance. Refer to [Designing your Caching Infrastructure](#) for more details.

Recommendation: Enable Cache on workstations connecting to the Nucleus server. Carefully design Caching infrastructure as needed.

Conclusion

VMware Cloud Foundation combines highly secure software-defined compute, storage, and network. With VMware Cloud Foundation, we can easily manage the lifecycle of the hybrid cloud environment. IT administrators can set up the Omniverse Enterprise environment confidently on a consistent infrastructure with much lower cost and simplified management. VMware Horizon 8 makes the delivery of Omniverse Enterprise client virtual desktops easier than ever. With NVIDIA GRID vGPU technology, VMware Horizon can deliver an enhanced user experience with Blast Performance.

In this solution paper, we demonstrated the architecture of running Omniverse Enterprise with VMware Cloud Foundation and VMware Horizon 8. We showed the configuration details, the hardware resources, and the software resources used in the solution validation. We showed the various configuration options in addition to the best practices. All the above lead to a consolidated solution of running Omniverse Enterprise with VMware Cloud Foundation and VMware Horizon 8.

References

- [Omniverse Deployment](#)
- [VMware Cloud Foundation](#)
- [VMware vSphere](#)
- [VMware vSAN](#)
- [VMware NSX Data Center](#)
- [VMware Cloud Foundation Admin Guide](#)
- [VMware Cloud Foundation Architecture and Deployment Guide](#)

About the Author

Aaron Chu, Sr. Solution Engineer in the Workload Application Solutions team of the Cloud Infrastructure Business Group, wrote the original version of this paper.

The following reviewers also contributed to the paper contents:

- Charles Lee, Sr. Technical Marketing Architect in the Workload Application Solutions, VMware
- Ting Yin, Sr. Technical Marketing Architect in the Workload Application Solutions team, VMware
- Chen Wei, Director in the Workload Application Solutions team, VMware
- Jim Senicka, Director in the VMware Cloud Foundation Technical Marketing in VMware
- Catherine Xu, Group Manager in the Workload Application Solutions team, VMware
- Andy Ju, Manager in the Technical Product Marketing, NVIDIA Corporation
- Alice Hwang, Sr. Technical Marketing Program Manager, Omniverse, NVIDIA
- David Sun, Technical Marketing Engineer, Omniverse, NVIDIA
- Sunny Sehmi, Technical Marketing Engineer, Omniverse, NVIDIA
- Madison Huang, Product Marketing Manager, Omniverse, NVIDIA



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.
Copyright © 2022 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at vmware.com/go/patents.
VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-wp-tech-temp-word-102-proof 5/19