



# Virtuozzo Hybrid Server 7

Installation on ASRock Rack

May 18, 2022

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2022 Virtuozzo International GmbH. All rights reserved.

This product is protected by United States and international copyright laws. The product's underlying technology, patents, and trademarks are listed at <https://www.virtuozzo.com/legal.html>.

Microsoft, Windows, Windows Server, Windows NT, Windows Vista, and MS-DOS are registered trademarks of Microsoft Corporation.

Apple, Mac, the Mac logo, Mac OS, iPad, iPhone, iPod touch, FaceTime HD camera and iSight are trademarks of Apple Inc., registered in the US and other countries.

Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective owners.

# Contents

<b>1. About This Guide</b>	<b>1</b>
<b>2. Installing Virtuozzo Hybrid Server</b>	<b>2</b>
2.1 Preparing Installation Media	2
2.2 Booting from USB	4
2.3 Performing Installation	5
2.3.1 CLI Installation	6
2.3.2 GUI Installation	14
<b>3. Enabling SR-IOV on ASRock Rack</b>	<b>25</b>
3.1 About SR-IOV	25
3.2 Enabling SR-IOV and IOMMU Support	27
3.3 Configuring SR-IOV VFs Network Adapters	28
3.4 Assigning SR-IOV VF Network Adapter to Linux KVM VM	30

## CHAPTER 1

# About This Guide

This guide explains how to install Virtuozzo Hybrid Server 7 on the ASRock Rack X570D4U-2L2T server powered by AMD Ryzen 9 5950X 16-Core @ 3.40GHz with CLI management or GUI (Virtuozzo Automator) support.

### **About the Hardware**

The ASRock Rack X570D4U-2L2T features:

- CPU AMD Ryzen 9 5950X 16-Core @3.40GHz.
- Four 288-pin DDR4 DIMM Slots. Support for up to 128GB DDR4 ECC/UDIMM (conditionally supports ECC error reporting function).
- Four hot swap 3.5-inch HDD/SSD bays.
- Three fixed 2.5-inch HDD/SSD bays.
- Two M.2 Slots on the Motherboard, our test server came with 2 NVMe Drives on 1 x each slot.
- Dual Intel Gigabit Ethernet ports (Intel i210 controller).
- IPMI connectivity (ASpeed AST2500) BMC.
- Two 10 Gigabit Ethernet ports.

## CHAPTER 2

# Installing Virtuoizzo Hybrid Server

This chapter explains how to install Virtuoizzo Hybrid Server 7 on the AMD Ryzen 9 5950X 16-Core @ 3.40GHz.

---

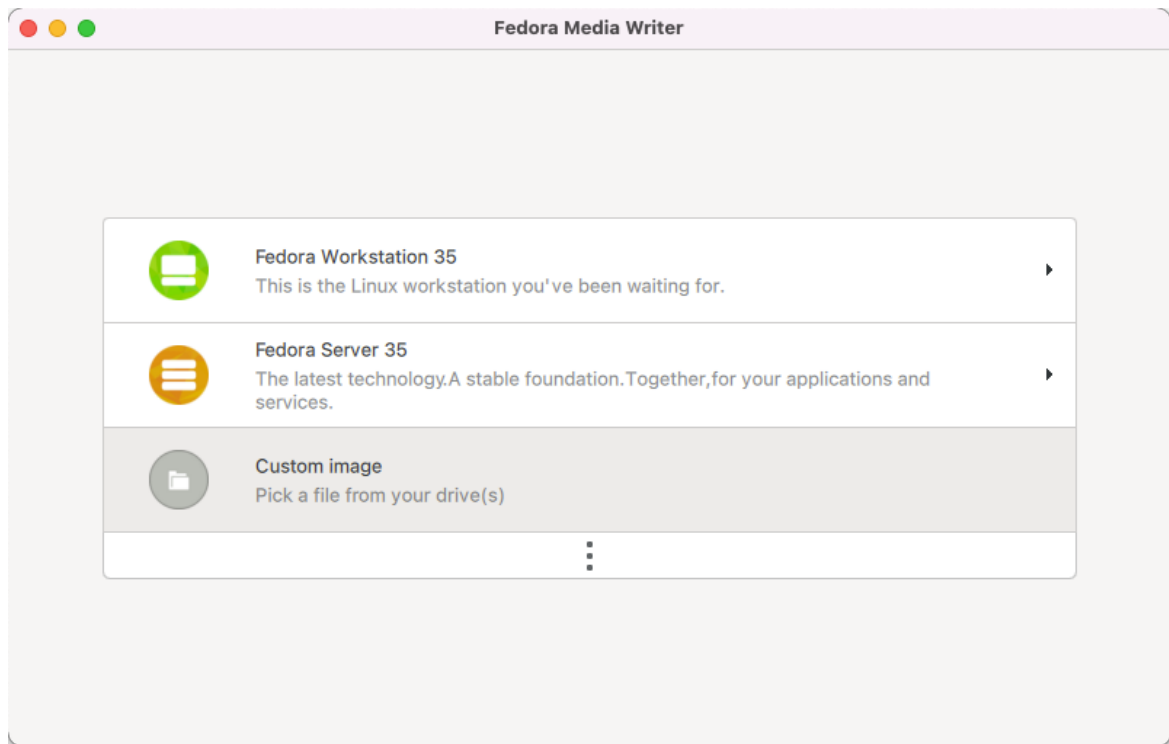
**Note:** You can check the general [Virtuoizzo Hybrid Server 7 Installation Guide](#).

---

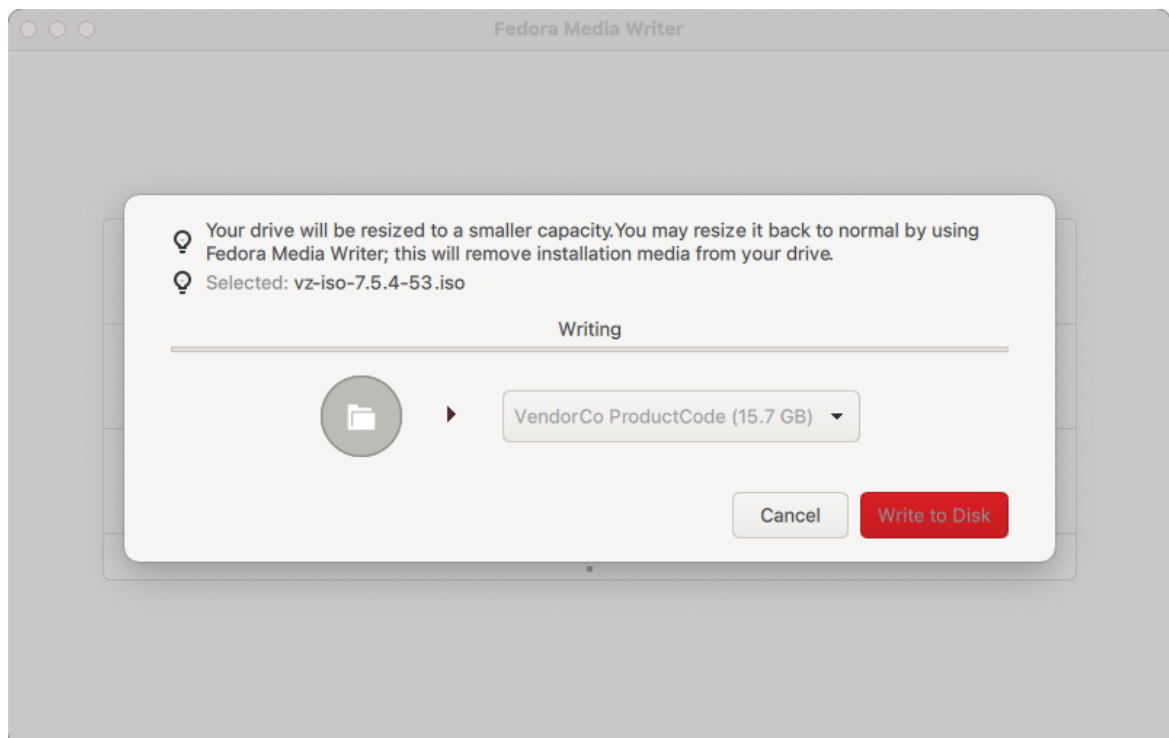
## 2.1 Preparing Installation Media

In this guide, we will perform a USB-based installation. The **Fedora media writer** software will be used to create the installation USB.

1. Download the Fedora media writer for your OS from here:
  - MacOS: <https://getfedora.org/fmw/FedoraMediaWriter-osx-latest.dmg>
  - Windows: <https://getfedora.org/fmw/FedoraMediaWriter-win32-latest.exe>
2. Download the Virtuoizzo Hybrid Server 7 ISO by following the steps in the linked [Obtaining the Distribution Image](#) guide. Also, in the received email, you will find a license key number to activate a Virtuoizzo Hybrid Server.
3. Open the Fedora media writer and click the *Custom Image* option.



4. Select the VHS ISO and click the **Write to Disk** button.



5. You can close once the media is ready.

## 2.2 Booting from USB

Once the installation USB is ready, the next step will be starting the ASRock Rack X570D4U-2L2T server and going to the boot menu.

1. Boot the server.
2. Press *F11* to access the boot menu.



3. Choose the installer USB to boot.

```

Please select boot device:

UEFI: Built-in EFI Shell
NVMe B1D0F0: M.2 (P80) 3TE6
NVMe B23D0F0: M.2 (P80) 3TE6
UEFI: PXE IP4 P0 Intel(R) Ethernet Controller X550
UEFI: PXE IP6 P0 Intel(R) Ethernet Controller X550
UEFI: PXE IP4 P1 Intel(R) Ethernet Controller X550
UEFI: PXE IP6 P1 Intel(R) Ethernet Controller X550
UEFI: PXE IP4 P1 Intel(R) I210 Gigabit Network Connection
UEFI: PXE IP6 P1 Intel(R) I210 Gigabit Network Connection
UEFI: PXE IP4 P0 Intel(R) I210 Gigabit Network Connection
UEFI: PXE IP6 P0 Intel(R) I210 Gigabit Network Connection
Virtuozzo Hybrid Infrastructure (M.2 (P80) 3TE6)
UEFI OS (M.2 (P80) 3TE6)
USB: VendorCoProductCode
UEFI: VendorCoProductCode, Partition 2
Enter Setup

↑ and ↓ to move selection
ENTER to select boot device
ESC to boot using defaults

```

## 2.3 Performing Installation

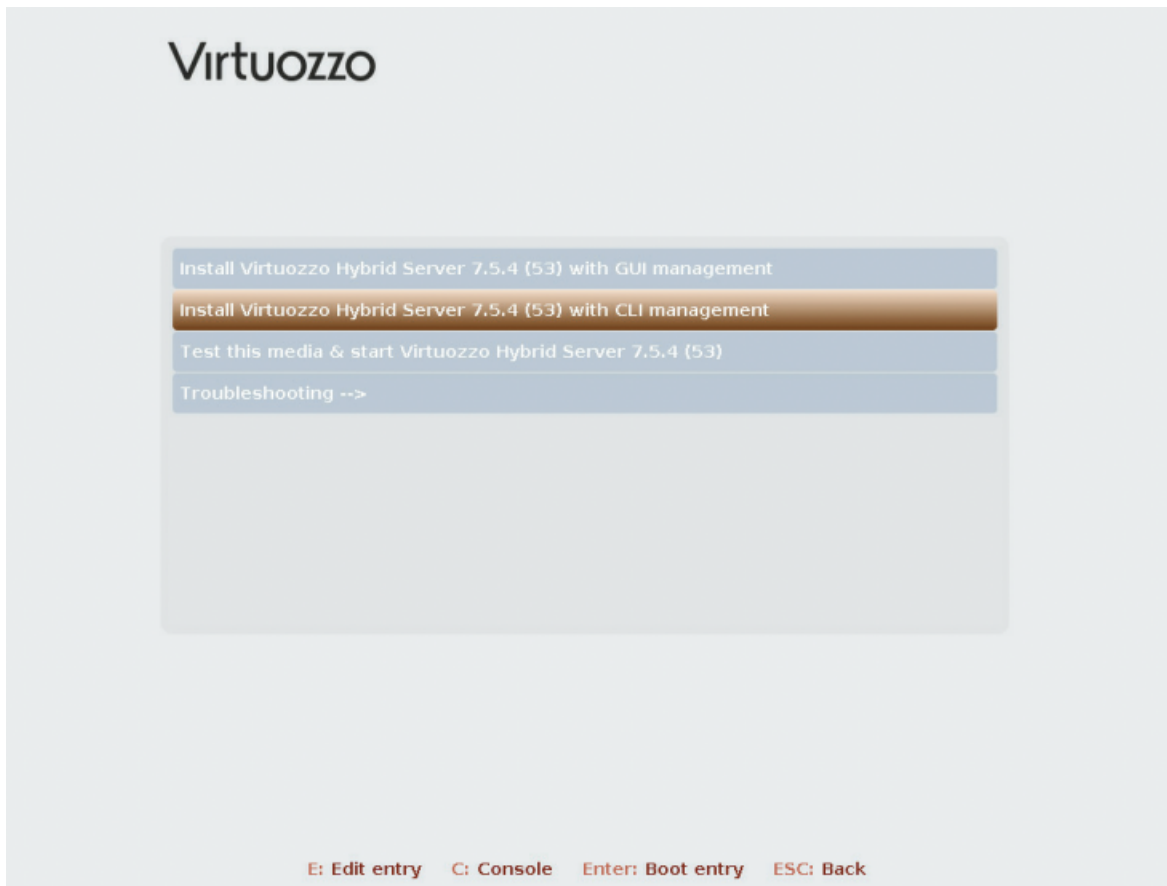
Based on your preferences, you can install Virtuozzo Hybrid Server 7 with:

- **CLI support** - no GUI/Automator will be installed; VHS administration will be performed from the command line
- **GUI support** - Automator is installed; VHS administration will be done from the Control Plane (Web UI)



## 2.3.1 CLI Installation

1. Click the *Install Virtuozzo Hybrid Server 7 with CLI management* option.



2. Within the installation window, you need to configure the following points:
  - accept the EULA
  - select storage type
  - choose the installation destination
  - configure the network interface



3. Click **EULA** and read the agreement. Accept it and click **Done** to proceed.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

EULA

us Done

Virtuozzo End-User License Agreement

This End-User License Agreement (this "Agreement") is a legal contract between you, as either an individual or an Entity (as defined below), and Virtuozzo International GmbH ("Virtuozzo"). READ THE TERMS AND CONDITIONS OF THIS AGREEMENT CAREFULLY BEFORE DOWNLOADING, INSTALLING, OBTAINING A LICENSE KEY, OR OTHERWISE ACCESSING OR USING VIRTUOZZO'S PROPRIETARY SOFTWARE ACCOMPANIED BY THIS AGREEMENT (the "Software"). THE SOFTWARE IS COPYRIGHTED AND IT IS LICENSED TO YOU UNDER THIS AGREEMENT, NOT SOLD TO YOU. BY DOWNLOADING, INSTALLING, OBTAINING A LICENSE KEY, OR OTHERWISE ACCESSING OR USING THE SOFTWARE, YOU ACKNOWLEDGE THAT YOU HAVE READ THIS AGREEMENT, THAT YOU UNDERSTAND IT, AND THAT YOU ACCEPT AND AGREE TO BE BOUND BY ITS TERMS. IF YOU ARE ACCEPTING THIS AGREEMENT ON BEHALF OF A COMPANY, ORGANIZATION, EDUCATIONAL INSTITUTION, OR AGENCY, INSTRUMENTALITY OR DEPARTMENT OF A GOVERNMENT (AN "ENTITY") AS ITS AUTHORIZED LEGAL REPRESENTATIVE, THEN YOU REPRESENT AND WARRANT THAT YOU HAVE THE POWER AND AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS, AND REFERENCES TO "YOU" HEREIN REFER TO BOTH YOU, THE INDIVIDUAL END USER, AND THE ENTITY ON WHOSE BEHALF YOU ARE ACCEPTING THIS AGREEMENT. IF AT ANY TIME YOU ARE NOT WILLING TO BE BOUND BY THE TERMS OF THIS AGREEMENT, YOU SHOULD CLICK THE "I DO NOT ACCEPT" OR SIMILAR BUTTON, TERMINATE THE DOWNLOAD AND/OR INSTALLATION PROCESS, IMMEDIATELY CEASE AND REFRAIN FROM ACCESSING OR USING THE SOFTWARE AND DELETE ANY COPIES YOU MAY HAVE. THIS AGREEMENT, ALONG WITH ANY ADDITIONAL TERMS OR POLICIES INCORPORATED HEREIN BY REFERENCE, REPRESENTS THE ENTIRE AGREEMENT BETWEEN YOU AND VIRTUOZZO CONCERNING THE SOFTWARE, AND THIS AGREEMENT SUPERSEDES AND REPLACES ANY PRIOR PROPOSAL.

☒ Accept ☐ Decline

- Return to the main menu and go to **Select Storage Type**. Ensure that "*Basic Storage*" is selected and click the **Done** button.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

SELECT STORAGE TYPE

us Done

Select type of a Storage:

☒ **Basic Storage**  
Install or upgrade to basic storage.

☐ **Virtuozzo Storage**  
Install or upgrade to Virtuozzo Storage.

Create or join Virtuozzo Storage:

☒ **Join**  
You can join an existing Virtuozzo Storage in your network, if you have already set up one.

☐ **Create**  
You can create a new Virtuozzo Storage cluster. Choose this option if you do not have Virtuozzo Storage in your network or want to set up a separate Virtuozzo Storage cluster for this installation.

Specify the Virtuozzo Storage cluster to join:

Name:

Enter the name of the Virtuozzo Storage cluster to join this server to.

Select functional roles for this server:

☐ **Metadata Server Role**  
A Metadata Server stores metadata about Chunk Servers. The specified IP address will be used to access the Metadata Server and cannot be changed later.

☐ **Chunk Server Role**  
A Chunk Server stores the contents of virtual machines and Containers on its local drives.

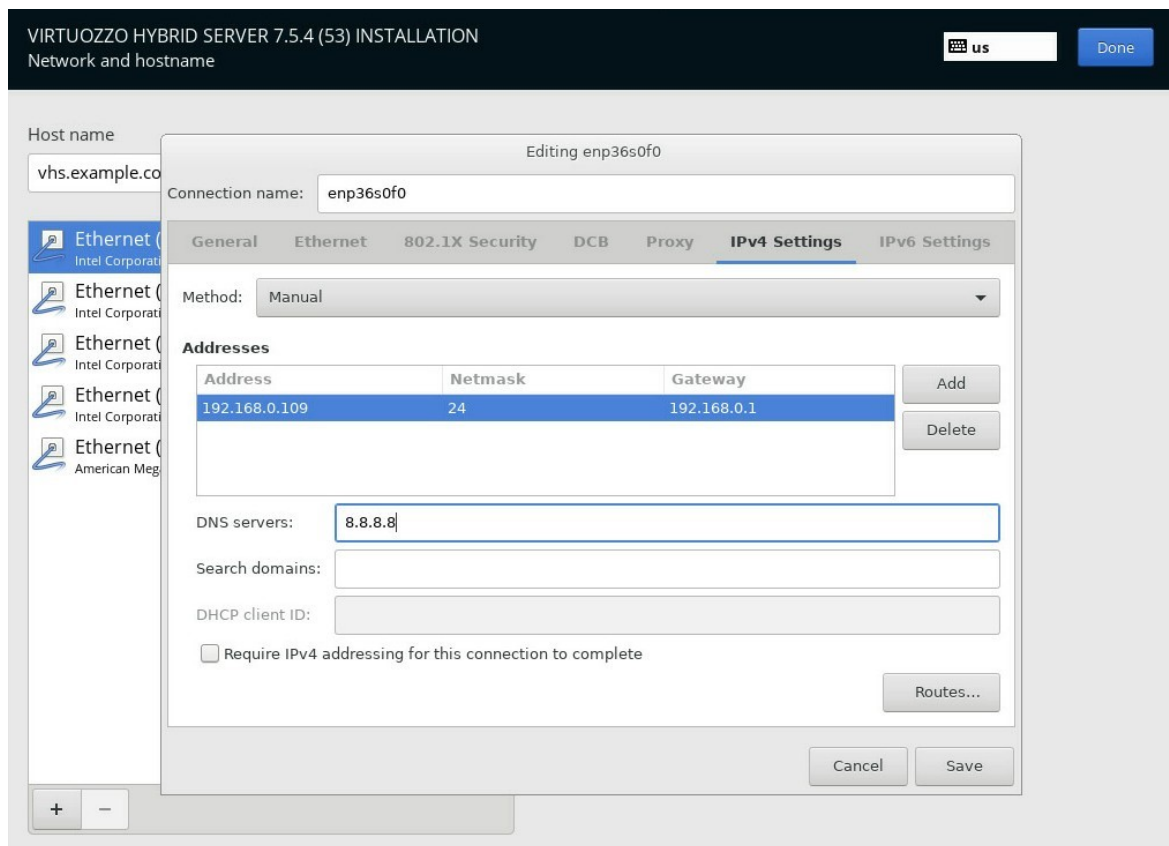
☐ **Client Server Role**  
A Client allows you to access Virtuozzo Storage from this server and run virtual machines and containers hosted on Virtuozzo Storage.

Internal IP address:

All ports on this interface will be open to let storage components reach other cluster nodes.

5. Next, configure the network interface by clicking **Network & Host Name**. We use the *vhs.example.com* hostname as an example. Choose the required interface (*enp36s0f0* in our case). Navigate to *Configure > IPv4 Settings > Method > Manual*, click **Add** and enter the following data:

- LAN IP address
- Netmask
- Gateway
- DNS servers



Click **Save** and **Done** when ready.

6. **Installation Destination** is the last option you need to configure. We have two disks, select *System* one (to install the OS) and tick both for *Datastore*. At the bottom of the window, select the *"I will configure partitioning"* radio button and click **Done**.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

Installation destination

us

Done

Select system disk(s)

Choose where to install the system. No data on disks will be touched until you click 'Begin installation'.

Disk	Type	Size	System	Datastore	Purpose
nvme0n1 / M.2 (P80) 3TE6	HDD	894.25 GiB	<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	Used by operating system
nvme1n1 / M.2 (P80) 3TE6	HDD	894.25 GiB	<input type="radio"/>	<input checked="" type="checkbox"/>	Available for storage

2 disks selected; 1788.51 GiB capacity; 2669 KiB free

Disks left unselected here will be cleaned.

☐ Automatically configure partitioning. ☒ I will configure partitioning.

Refresh

7. By default, a small partition is created for the OS (64GB), and the remaining space is used as a *Datastore*. Click the **Done** button when ready.

---

**Note:** Everything on those drives will be erased.

---

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION  
Manual partitioning

us Done

▼ New Virtuozzo Hybrid Server 7.5.4 (53) Installation

DATA

Partition	Size
/vz	1660.5 GiB
vhs_vhs-vz	

SYSTEM

Partition	Size
/boot	1024 MiB
nvme0n1p2	
BIOS Boot	1024 KiB
nvme0n1p1	
/	64 GiB
vhs_vhs-root	
swap	63 GiB
nvme0n1p3	

+ - ↺

AVAILABLE SPACE: 2669 KiB TOTAL SPACE: 1788.51 GiB

2 storage devices selected

vhs\_vhs-vz

Mount Point: /vz Device(s): M.2 (P80) 3TE6 (nvme0n1) and 1 other

Desired Capacity: 1660.5 GiB

Modify...

Device Type: LVM ☐ Encrypt

File System: ext4 tuned for VZ ☒ Reformat

Volume Group: vhs\_vhs (0 B free) Modify...


Label: Name: vz

Update Settings

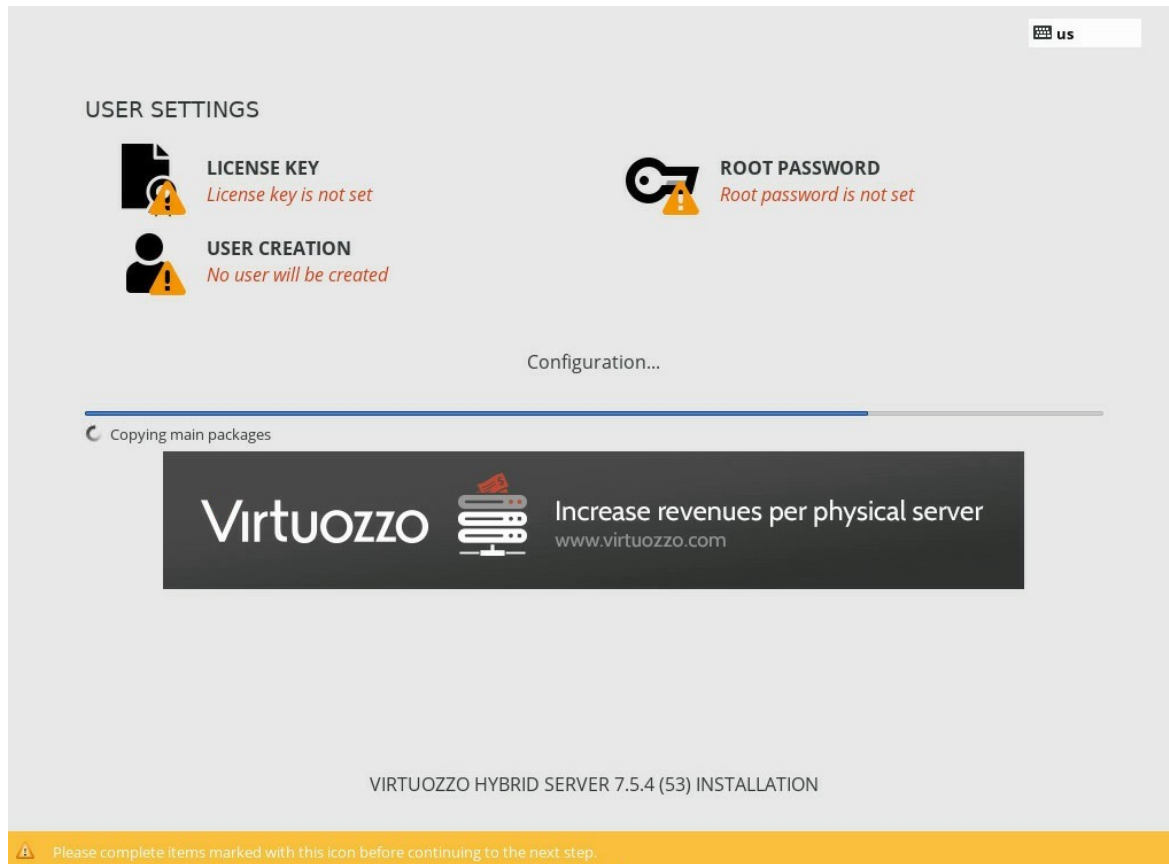
Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.

Reset All

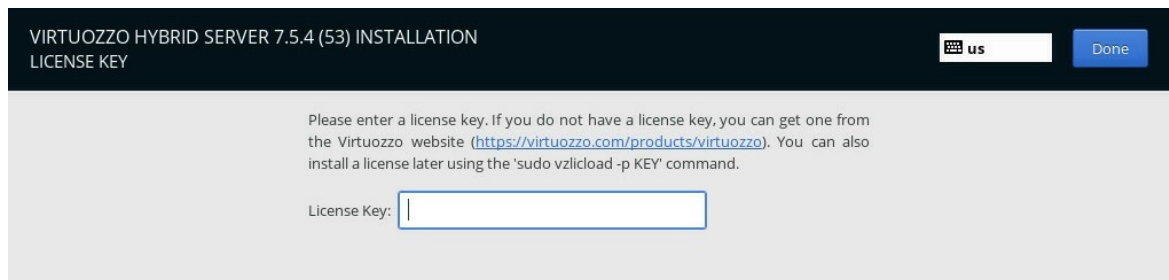
8. Now, you can click the **Begin Installation** button.



9. While the installation is in progress, you can provide the license key, set the root password, and create a user.



10. In the **License Key** section, enter the license key you've got via email alongside the VHS image and click **Done**.



11. Go to the **Root Password** section and create a root account password. Click **Done** to proceed.



VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

Root password

Specify password

The root account is used for administering the system.

Create password

Empty

Confirm password

The password must be at least 8 characters long, with at least one capital letter and one digit. The password can contain letters (a-z), numbers (0-9), dashes (-), underscores (\_), apostrophes ('), and periods (.).

12. The **Create User** section helps create a new user. Click **Done** when ready.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

Create user

Full name

User name

**Tip:** Keep your user name shorter than 32 characters and do not use spaces.

☒ Make this user administrator

☒ Require a password to use this account


Password


Weak


Confirm password

13. Once the installation is done and all the configurations are provided, click **Finish configuration** and **Reboot** the server.

USER SETTINGS

 **LICENSE KEY**  
License key is not set

 **ROOT PASSWORD**  
Root password is set

 **USER CREATION**  
Administrator tester will be created

Configuration...

Complete!

Virtuozzo Hybrid Server is now successfully installed and ready for you to use!  
Go ahead and reboot to start using it!



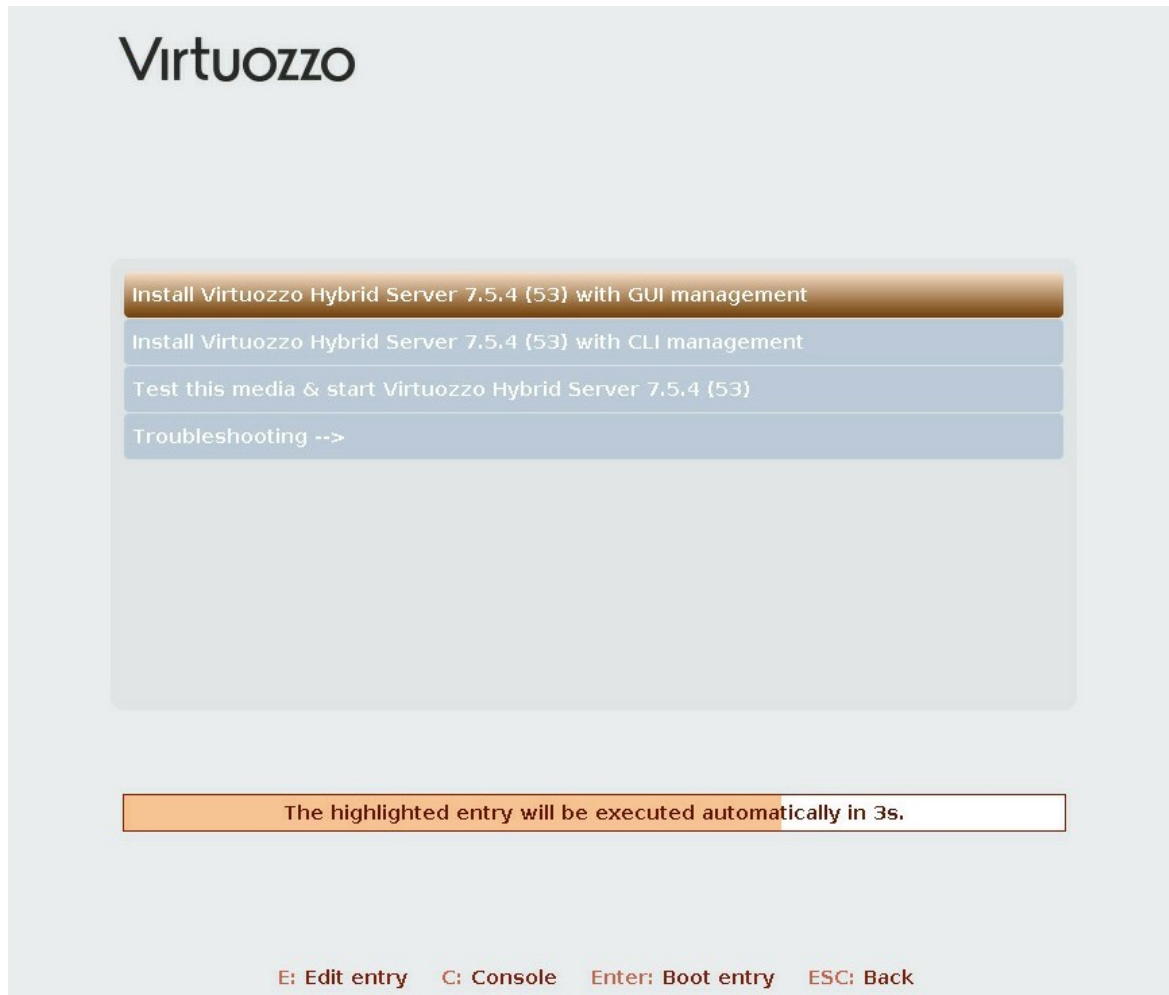
14. Verify installation by accessing the server over SSH via the LAN IP assigned during the installation.

```
jbustos@gibson ~ % ssh root@192.168.0.109
The authenticity of host '192.168.0.109 (192.168.0.109)' can't be established.
ED25519 key fingerprint is SHA256:o2zzZHAo8uJwb3NIL0yBm4YX3mp+eb4ddzMEeTmgyXE.
This key is not known by any other names
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '192.168.0.109' (ED25519) to the list of known hosts.
root@192.168.0.109's password:
MOTD generated at: 03:09:06
Uptime:          1 min
OS:              Virtuozzo release 7.5.4 (53)
IP:              192.168.0.109 10.37.130.2 fdb2:2c26:f4e4::1
Hostname:        vhs.example.com
Kernel:          3.10.0-1160.53.1.vz7.185.3 GNU/Linux
System Load:     0.11
/vz Usage:        0% of 1.6T
Swap Usage:       0%
RAM Free:         98% of 125.7GB

[root@vhs ~]#
```

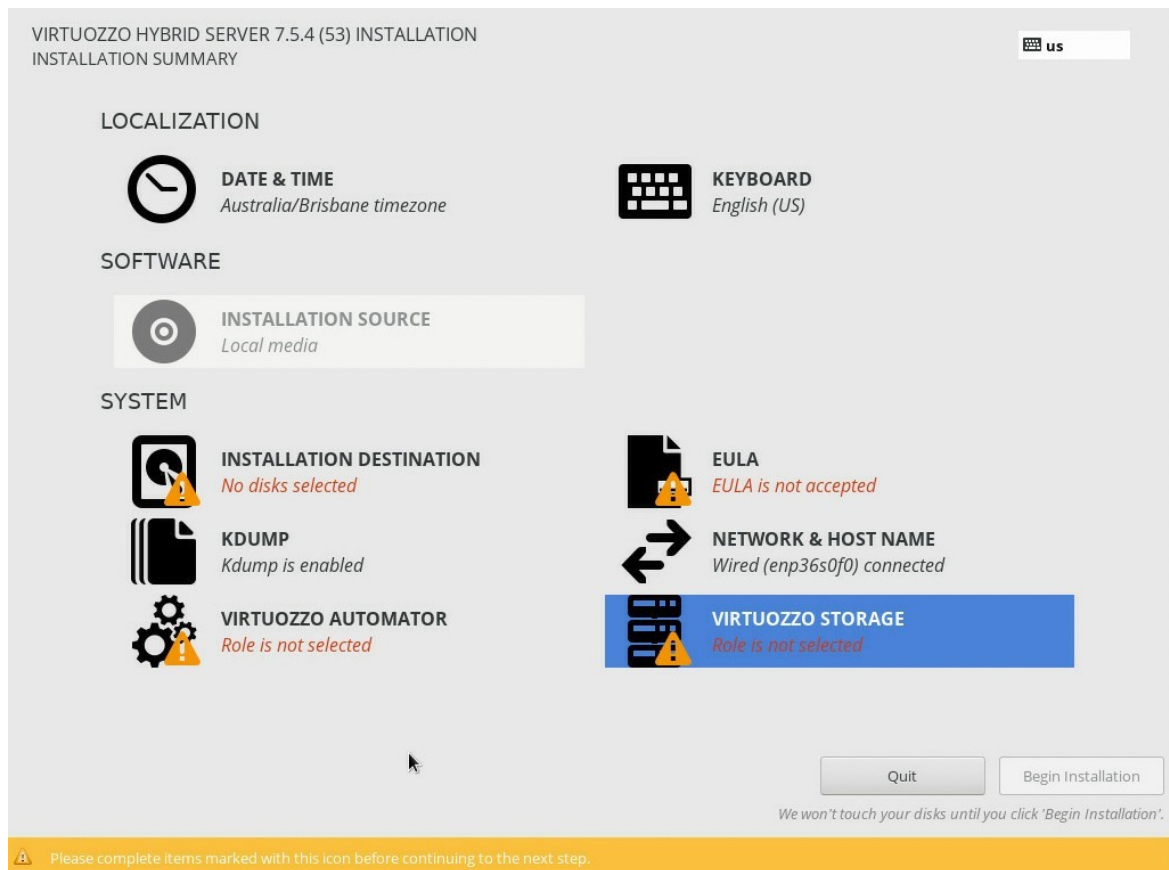
## 2.3.2 GUI Installation

1. Click the *Install Virtuozzo Hybrid Server 7 with GUI management* option.

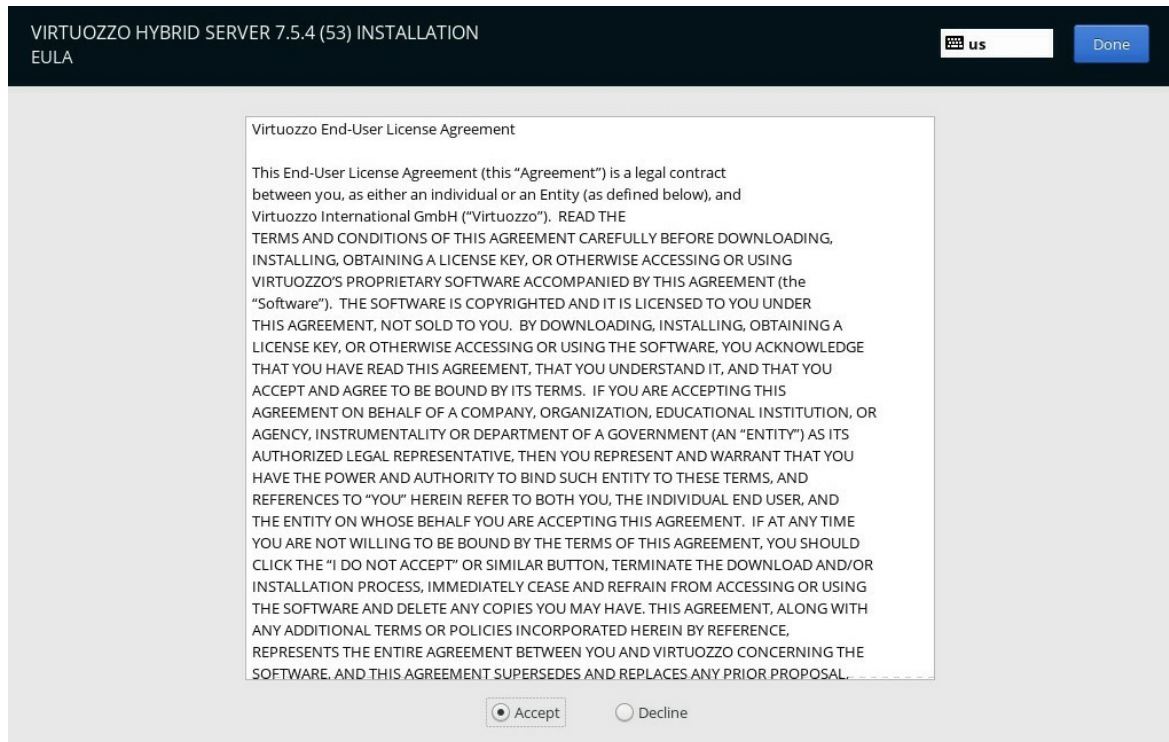


2. Within the installation window, you need to configure the following points:

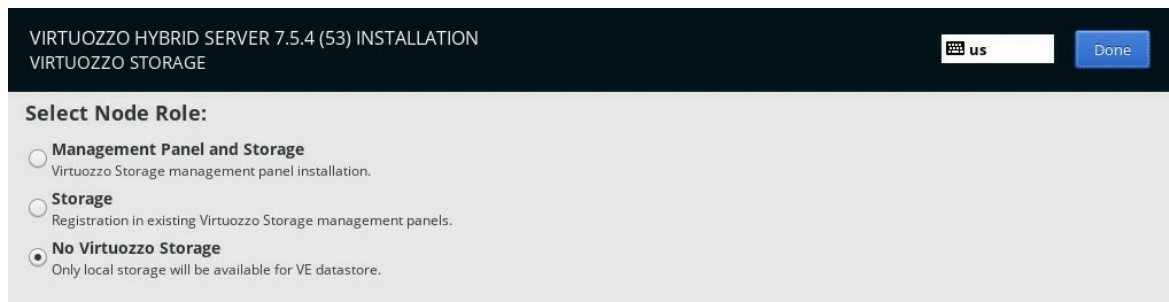
- accept the EULA
- disable Virtuozzo Storage (use local storage)
- configure the network interface
- define the IP of the Virtuozzo Automator (GUI)
- choose the installation destination



3. Click **EULA** and read the agreement. *Accept* it and click **Done** to proceed.

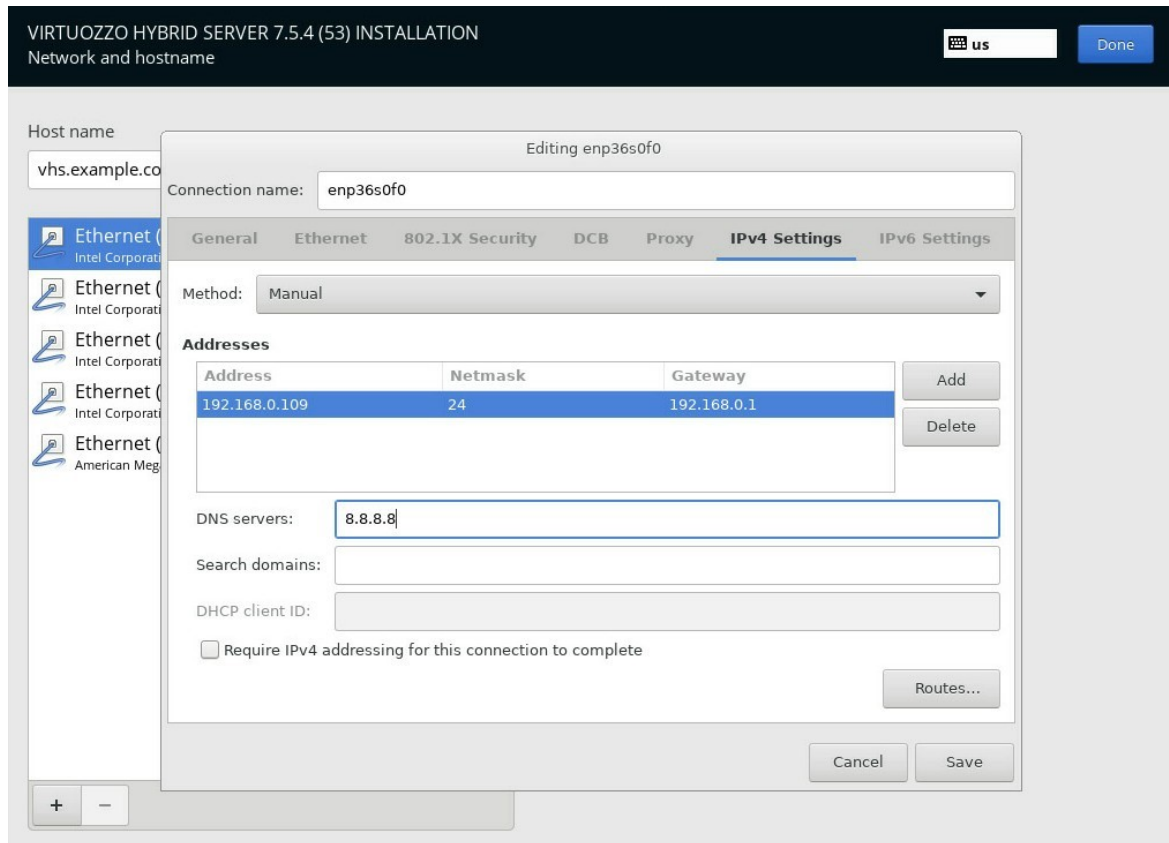


4. Return to the main menu and go to **Virtuozzo Storage**. Select the "*No Virtuozzo Storage*" option and click the **Done** button.



5. Next, configure the network interface by clicking **Network & Host Name**. We use the *vhs.example.com* hostname as an example. Choose the required interface (*enp36s0f0* in our case). Navigate to *Configure > IPv4 Settings > Method > Manual*, click **Add** and enter the following data:

- *IP address*
- *Netmask*
- *Gateway*
- *DNS servers*



Click **Save** and **Done** when ready.

6. In the **Virtuozzo Automator** section, choose the "Management Panel and Compute" option, enter the IP and hostname for the Automator GUI and click the **Done** button.

This action will deploy a Virtuozzo system container with the Automator GUI. The network interface specified in the previous step will be used to create a Linux bridge so that Automator GUI IP will be available through the bridge on the same LAN.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION  
VIRTUOZZO AUTOMATOR

Select Node Role:

☒ **Management Panel and Compute**  
Virtuozzo Automator installation.

☐ **Compute**  
Registration in existing Virtuozzo Automator.

**Important: Only one Virtuozzo Automator management panel is required, so choose this role for the first node only.**

VA Management Node IP address:\*

VA Management Node hostname:

A new container with Virtuozzo Automator will be created with the specified IP address and hostname.

To access Virtuozzo Automator management panel, visit `http://<IP_or_hostname>` in a web browser and log in as root.  
You will set a root password later during installation.

Done

7. **Installation Destination** is the last option you need to configure. We have two disks, select *System* one (to install the OS) and tick both for *Datastore*. At the bottom of the window, select the *"I will configure partitioning"* radio button and click **Done**.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION  
Installation destination

Select system disk(s)

Choose where to install the system. No data on disks will be touched until you click 'Begin installation'.

Disk	Type	Size	System	Datastore	Purpose
nvme0n1 / M.2 (P80) 3TE6	HDD	894.25 GiB	<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	Used by operating system
nvme1n1 / M.2 (P80) 3TE6	HDD	894.25 GiB	<input type="radio"/>	<input checked="" type="checkbox"/>	Available for storage

2 disks selected; 1788.51 GiB capacity; 2669 KiB free

☐ Automatically configure partitioning. ☒ I will configure partitioning.

Disks left unselected here will be cleaned.

Refresh

8. By default, a small partition is created for the OS (64GB), and the remaining space is used as a *Datastore*. Click the **Done** button when ready.

---

**Note:** Everything on those drives will be erased.

---

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION  
Manual partitioning

us Done

▼ New Virtuozzo Hybrid Server 7.5.4 (53) Installation

DATA

Partition	Size
/vz	1660.5 GiB
vhs_vhs-vz	

SYSTEM

Partition	Size
/boot	1024 MiB
nvme0n1p2	
BIOS Boot	1024 KiB
nvme0n1p1	
/	64 GiB
vhs_vhs-root	
swap	63 GiB
nvme0n1p3	

+ - ↺

AVAILABLE SPACE: 2669 KiB  
TOTAL SPACE: 1788.51 GiB

2 storage devices selected

vhs\_vhs-vz

Mount Point: /vz

Desired Capacity: 1660.5 GiB

Device(s): M.2 (P80) 3TE6 (nvme0n1) and 1 other

Modify...

Device Type: LVM ☐ Encrypt

File System: ext4 tuned for VZ ☒ Reformat

Volume Group: vhs\_vhs (0 B free) Modify...


Label: Name: vz

Update Settings


Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.

Reset All


9. Now, you can click the **Begin Installation** button.



10. While the installation is in progress, you can provide the license key, set the root password, and create a user.



11. In the **License Key** section, enter the license key you've got via email alongside the VHS image and click **Done**.



12. Go to the **Root Password** section and create a root account password. Click **Done** to proceed.



VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

Root password

Specify password

The root account is used for administering the system.

Create password

Empty

Confirm password

The password must be at least 8 characters long, with at least one capital letter and one digit. The password can contain letters (a-z), numbers (0-9), dashes (-), underscores (\_), apostrophes ('), and periods (.).

13. The **Create User** section helps create a new user. Click **Done** when ready.

VIRTUOZZO HYBRID SERVER 7.5.4 (53) INSTALLATION

Create user

Full name

User name

**Tip:** Keep your user name shorter than 32 characters and do not use spaces.

☒ Make this user administrator

☒ Require a password to use this account


Password


Weak


Confirm password

14. Once the installation is done and all the configurations are provided, click **Finish configuration** and **Reboot** the server.

USER SETTINGS

 **LICENSE KEY**  
License key is not set

 **ROOT PASSWORD**  
Root password is set

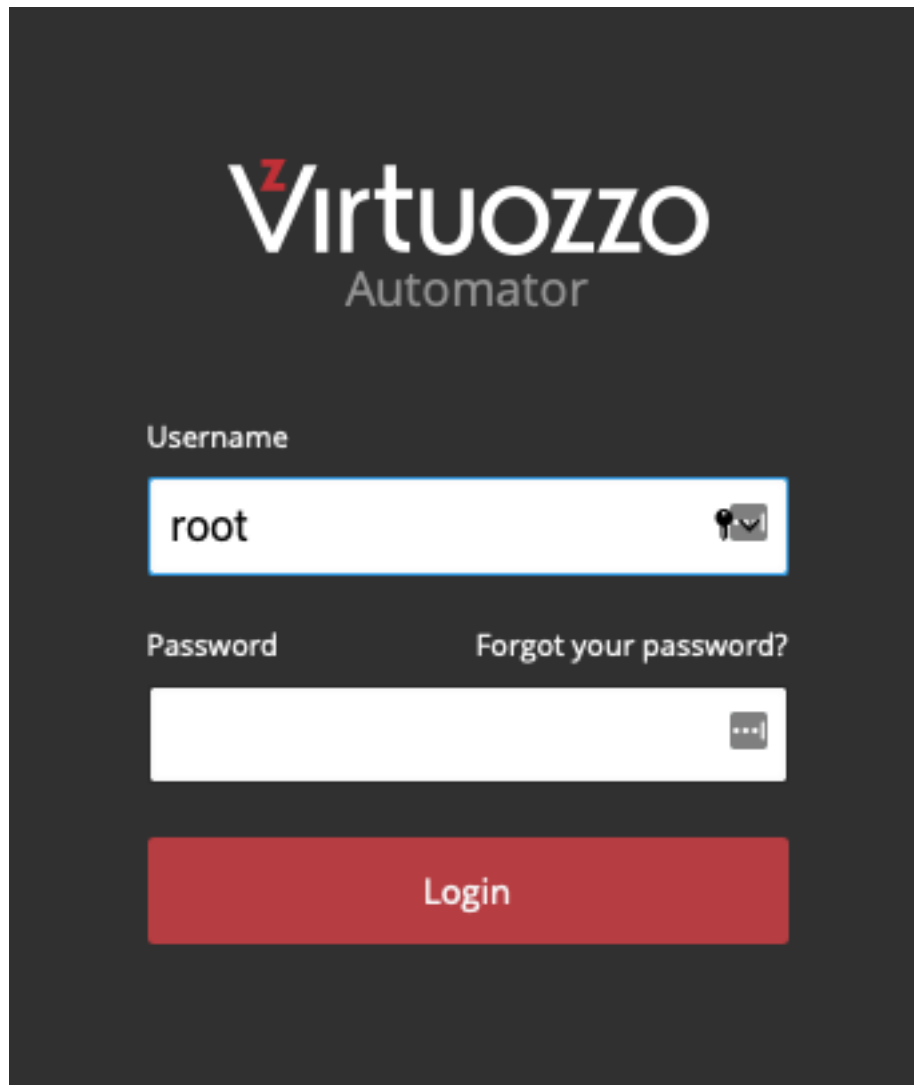
 **USER CREATION**  
Administrator tester will be created

Configuration...

Complete!

Virtuozzo Hybrid Server is now successfully installed and ready for you to use!  
Go ahead and reboot to start using it!

15. Once the system is booted, go to the IP specified during the Virtuozzo Automator configuration (<https://ip-of-the-automator-gui>). Use the root user with the password you've specified in the 12th step to log in.

The image shows the login interface of the Virtuozzo Automator. It has a dark gray background. At the top, the 'Virtuozzo Automator' logo is displayed in white, with a red 'z' in the 'Virtuozzo' part. Below the logo, there are two input fields. The first is labeled 'Username' and contains the text 'root'. The second is labeled 'Password' and is empty. To the right of the password field is a link that says 'Forgot your password?'. Below the password field is a large red button with the word 'Login' in white text. There are also small icons for password visibility (an eye) and a key icon next to the username field.

Or you can verify the installation by accessing the server over SSH via the LAN IP assigned during the installation.

```
jbustos@gibson ~ % ssh root@192.168.0.109
The authenticity of host '192.168.0.109 (192.168.0.109)' can't be established.
ED25519 key fingerprint is SHA256:o2zzZHAo8uJwb3NIL0yBm4YX3mp+eb4ddzMEeTmgyXE.
This key is not known by any other names
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '192.168.0.109' (ED25519) to the list of known hosts.
root@192.168.0.109's password:
MOTD generated at: 03:09:06
Uptime:                1 min
OS:                    Virtuozzo release 7.5.4 (53)
IP:                    192.168.0.109 10.37.130.2 fdb2:2c26:f4e4::1
Hostname:              vhs.example.com
Kernel:                3.10.0-1160.53.1.vz7.185.3 GNU/Linux
System Load:          0.11
/vz Usage:             0% of 1.6T
Swap Usage:            0%
RAM Free:              98% of 125.7GB

[root@vhs ~]#
```

## CHAPTER 3

# Enabling SR-IOV on ASRock Rack

This chapter explains how to enable and configure SR-IOV on Virtuozzo Hybrid Server 7 and assign the PCI SR-IOV devices to virtual machines.

## 3.1 About SR-IOV

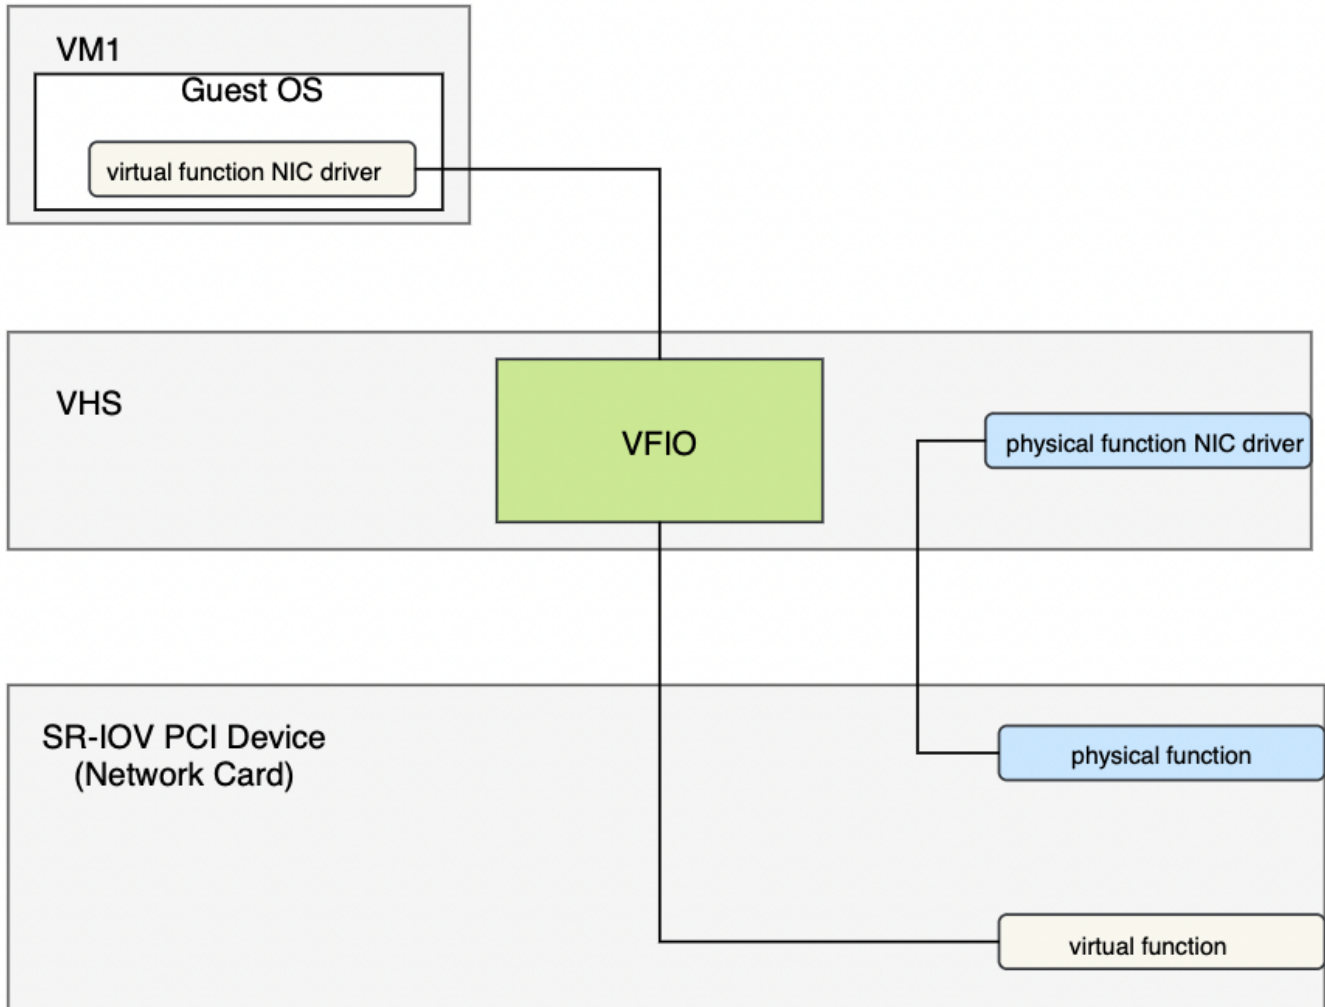
**Single Root I/O Virtualization (SR-IOV)** is a hardware specification that allows a single PCI Express (PCIe) endpoint to function as numerous devices (e.g., a single Ethernet port will appear as multiple, separate, physical devices). *Physical Functions (PFs)* and *Virtual Functions (VFs)* are two PCIe functions that enable this functionality. *PCI passthrough* is another functionality provided by the IOMMU driver, which allows assigning PCIe devices directly to Virtual Hosts. It allows virtual machines to benefit from direct PCI device assignment while only requiring one slot on the host physical computer.

When SR-IOV is used in conjunction with the PCI passthrough functionality in its most basic form (**SR-IOV VF PCI**), we can inject an SR-IOV network VF into a Linux KVM VM.

**Physical Functions** are Full PCIe devices that include SR-IOV capabilities (PCIe Devices).

**Virtual Functions** are PCIe functions that only compute I/O. These are the resulting VFs from the PCIe PFs device (Virtual NICs).

The direct assignment (**PCI passthrough**) approach also bypasses the Linux virtual bridges that are usually in place for virtual machine networking.



### Why use SR-IOV on VHS?

Virtual Functions offer **better performance** than paravirtual drivers or emulated devices. The performance is nearly **close to native**. An overlooked feature when using VFs is **data protection**; all the data is managed and controlled by the hardware, which improves data protection between virtual machines on the same host. These features allow the host to have an increased virtual machine density. It is important to mention that Virtual Machines with assigned PCI devices cannot be migrated to other hosts, which in our case works well as we will be using a dedicated host with local storage, and machines will always be running on this host.

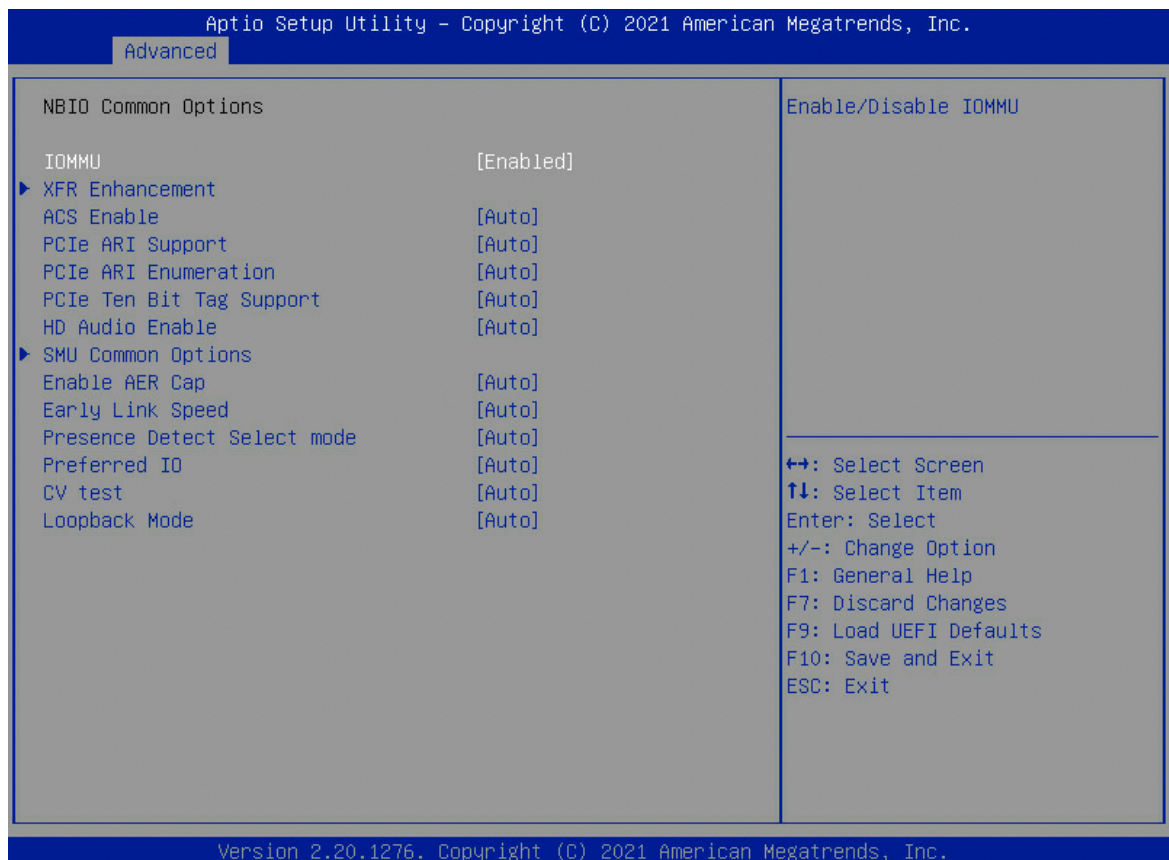
## 3.2 Enabling SR-IOV and IOMMU Support

**SR-IOV VF PCI** relies heavily on hardware support. Here is a list of hardware requirements for SR-IOV:

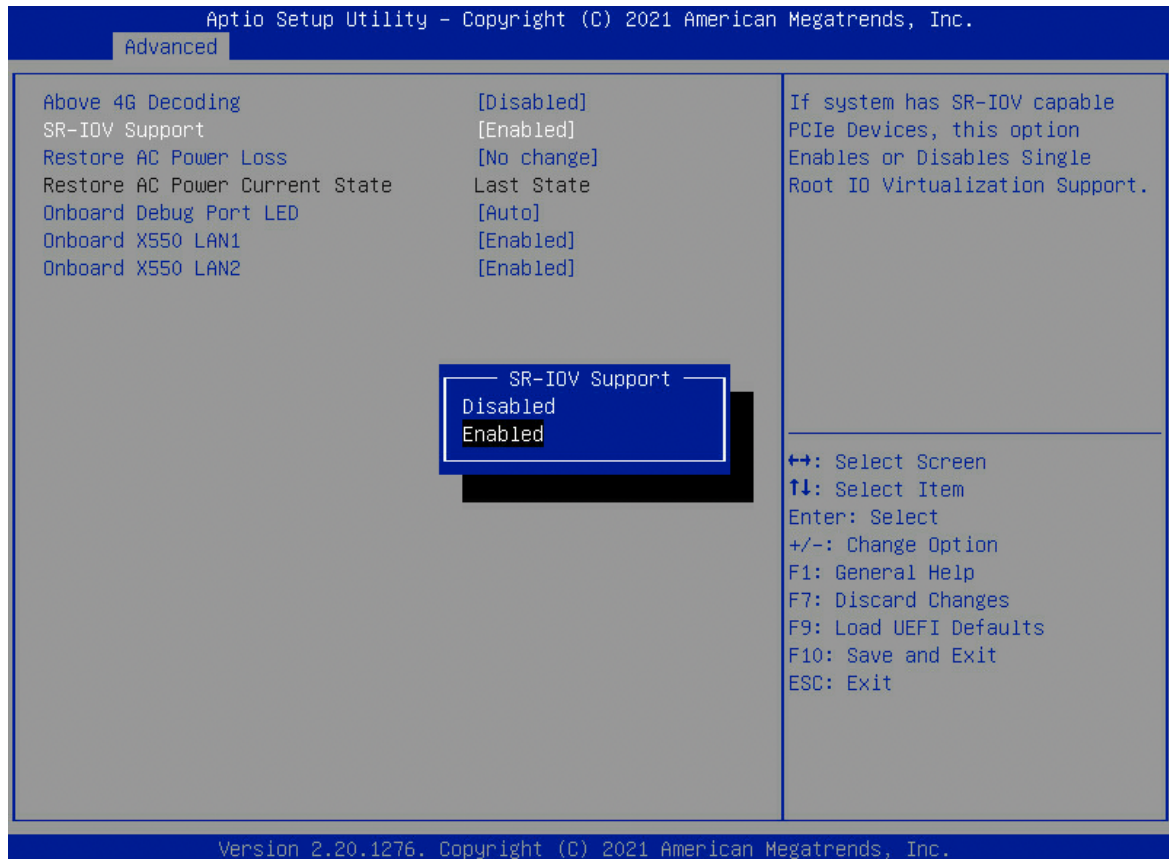
- Your firmware (BIOS or UEFI) must support SR-IOV.
- Your PCIe devices (e.g., Ethernet Port) must support SR-IOV.
- Root Ports or the PCIe switch must support ARI (alternative routing ID interpretation).
- Device assignment requires IOMMU (I/O Memory Management Unit) support in the CPU and firmware.

The **ASRock Rack X570D4U-2L2T server powered by AMD Ryzen 9 5950X 16-Core @ 3.40GHz** was certified for compatibility with Virtuozzo Hybrid Server version 7.5.4 by the Virtuozzo Technology Alliances team. All validation tests available for the standalone server configuration were completed with a 100% pass rate. A wide range of installation scenarios and functional tests were covered, including SR-IOV and PCIe passthrough support.

1. Check if **IOMMU** is enabled on the BIOS or UEFI. Go to *Advanced > AMD CBS > NBIO Common Options > IOMMU* and enable the option, if needed.



2. Check that **SR-IOV** support is enabled on the BIOS or UEFI. Go to *Advanced > Chipset Configuration > SR-IOV Support* and enable the option, if needed.



### 3.3 Configuring SR-IOV VFs Network Adapters

There are a few configuration steps required to create SR-IOV VFs network adapters:

- Identify network cards on your system and verify SR-IOV support.
- Verify that IOMMU is supported and loaded by your kernel when booting.
- Define how many SR-IOV VFs network adapters are created.
- Enable persistent device creation after reboot.

1. Identify the available network cards.

```
# lshw -c network -businfo
```

```
[root@vhs ~]# lshw -c network -businfo
Bus info          Device          Class          Description
=====
pci@0000:24:00.0  enp36s0f0       network        Ethernet Controller 10G X550T
pci@0000:24:00.1  enp36s0f1       network        Ethernet Controller 10G X550T
pci@0000:26:00.0  enp38s0         network        I210 Gigabit Network Connection
pci@0000:27:00.0  enp39s0         network        I210 Gigabit Network Connection
```

2. Use the **Bus info** (without the *pci@* part) from the previous step output to check for SR-IOV support. For example, on the *enp360f0* device:

```
# lspci -vs 0000:24:00.0
```

```
[root@vhs ~]# lspci -vs 0000:24:00.0
24:00.0 Ethernet controller: Intel Corporation Ethernet Controller 10G X550T (rev 01)
Subsystem: ASRock Incorporation Device 1563
Flags: bus master, fast devsel, latency 0, IRQ 38
Memory at fc800000 (64-bit, prefetchable) [size=4M]
Memory at fcc04000 (64-bit, prefetchable) [size=16K]
Expansion ROM at fb280000 [disabled] [size=512K]
Capabilities: [40] Power Management version 3
Capabilities: [50] MSI: Enable- Count=1/1 Maskable+ 64bit+
Capabilities: [70] MSI-X: Enable+ Count=64 Masked-
Capabilities: [a0] Express Endpoint, MSI 00
Capabilities: [100] Advanced Error Reporting
Capabilities: [140] Device Serial Number 00-00-00-ff-ff-00-00-00
Capabilities: [150] Alternative Routing-ID Interpretation (ARI)
Capabilities: [160] Single Root I/O Virtualization (SR-IOV)
Capabilities: [1a0] Transaction Processing Hints
Capabilities: [1b0] Access Control Services
Capabilities: [1c0] Latency Tolerance Reporting
Capabilities: [1d0] #19
Kernel driver in use: ixgbe
Kernel modules: ixgbe
```

3. Create **SR-IOV VFs network adapter**. You need to set the required adapters' number (*one*, in our example):

```
# echo 1 > /sys/class/net/enp36s0f0/device/sriov_numvfs
```

4. Verify that adapters were correctly created.

```
# lshw -c network -businfo
```



```
[root@vhs ~]# lshw -c network -businfo
```

Bus info	Device	Class	Description
pci@0000:24:00.0	enp36s0f0	network	Ethernet Controller 10G X550T
pci@0000:24:00.1	enp36s0f1	network	Ethernet Controller 10G X550T
pci@0000:26:00.0	enp38s0	network	I210 Gigabit Network Connection
pci@0000:27:00.0	enp39s0	network	I210 Gigabit Network Connection
pci@0000:25:10.0	enp37s16	network	X550 Virtual Function
	br0	network	Ethernet interface
	virbr0-nic	network	Ethernet interface
	host-routed	network	Ethernet interface
usb@3:5.3	enp42s0f3u5u3c2	network	Ethernet interface
	virbr0	network	Ethernet interface

```
[root@vhs ~]#
```

5. To make these changes permanent, you can create the following *udev rules* file (use the name of your interface as the file name).

```
# vim /etc/udev/rules.d/enp36s0f0.rules
ACTION=="add", SUBSYSTEM=="net", ENV{ID_NET_DRIVER}=="ixgbe",ATTR{device/sriov_numvfs}="1"
```

## 3.4 Assigning SR-IOV VF Network Adapter to Linux KVM VM

We have a virtual machine running CentOS 7. This virtual machine already has a primary network interface with IP address *192.168.0.223*. Our goal is to assign a function network interface created previously to this running CentOS 7 virtual machine.

We have identified the new virtual function network adapter as *enp37s16* and can see that it appears as an available network interface on our hypervisor.

```
[root@vhs ~]# ip link show enp37s16
8: enp37s16: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 12:42:f7:4b:2e:9e brd ff:ff:ff:ff:ff:ff
```

Now, let's see how to pass through this network virtual function to a virtual machine in a persistent way.

1. We will be using the **Bus info** in the next steps. However, you need to replace ":" and "@" with "\_" when querying the device. For example, *pci@0000:25:10:0* » *pci\_0000\_25\_10\_0*.

```
[root@vhs ~]# lshw -c network -businfo
```

Bus info	Device	Class	Description
pci@0000:24:00.0	enp36s0f0	network	Ethernet Controller 10G X550T
pci@0000:24:00.1	enp36s0f1	network	Ethernet Controller 10G X550T
pci@0000:26:00.0	enp38s0	network	I210 Gigabit Network Connection
pci@0000:27:00.0	enp39s0	network	I210 Gigabit Network Connection
pci@0000:25:10.0	enp37s16	network	X550 Virtual Function
	br0	network	Ethernet interface
	virbr0-nic	network	Ethernet interface
	host-routed	network	Ethernet interface
usb@3:5.3	enp42s0f3u5u3c2	network	Ethernet interface
	virbr0	network	Ethernet interface

```
[root@vhs ~]#
```

2. Gather information about the **VF network adapter**. Use *virsh* and the *pci* address in the format from the previous step.

```
# virsh nodedev_dumpxml pci_0000_25_10_0
```

```
[root@vhs ~]# virsh nodedev-dumpxml pci_0000_25_10_0
<device>
  <name>pci_0000_25_10_0</name>
  <path>/sys/devices/pci0000:00/0000:00:01.2/0000:20:00.0/0000:21:02.0/0000:25:10.0</path>
  <parent>pci_0000_21_02_0</parent>
  <driver>
    <name>ixgbevf</name>
  </driver>
  <capability type='pci'>
    <class>0x020000</class>
    <domain>0</domain>
    <bus>37</bus>
    <slot>16</slot>
    <function>0</function>
    <product id='0x1565'>X550 Virtual Function</product>
    <vendor id='0x8086'>Intel Corporation</vendor>
    <capability type='phys_function'>
      <address domain='0x0000' bus='0x24' slot='0x00' function='0x0' />
    </capability>
    <iommuGroup number='35'>
      <address domain='0x0000' bus='0x25' slot='0x10' function='0x0' />
    </iommuGroup>
    <pci-express>
      <link validity='cap' port='0' speed='8' width='4' />
      <link validity='sta' width='0' />
    </pci-express>
  </capability>
</device>

[root@vhs ~]#
```

3. Using the information highlighted in the image from the previous step, create an *xml* file to define the

network interface.

```
[root@vhs ~]# vim /tmp/vf.xml

<interface type='hostdev' managed='yes'>
  <mac address='52:54:00:6d:90:02' />
  <driver name='vfio' />
  <source>
    <address type='pci' domain='0' bus='37' slot='16' function='0' />
  </source>
  <model type='virtio' />
  <alias name='hostdev0' />
</interface>
```

Attach the device using *virsh*. You can use the following flags:

- **-live** - attach the interface to the running virtual machine
- **-persistent** - use this network interface definition for future *guest* reboots (it doesn't persist on *host* reboots)
- **-config** - apply the changes after VM reboot

```
[root@vhs ~]# virsh attach-device MyVM vf.xml --persistent
Device attached successfully

[root@vhs ~]# █
```

Verify that the device has been attached to the virtual machine correctly.

```
[root@vhs ~]# virsh dumpxml MyVM | grep -w hostdev -A9
  <interface type='hostdev' managed='yes'>
    <mac address='52:54:00:6d:90:02' />
    <driver name='vfio' />
    <source>
      <address type='pci' domain='0x0000' bus='0x25' slot='0x10' function='0x0' />
    </source>
    <model type='virtio' />
    <alias name='hostdev0' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x07' function='0x0' />
  </interface>
[root@vhs ~]# █
```

4. We have attached a virtual function network card to the virtual machine as a secondary network NIC. Let's see the status of the network interfaces on the virtual machine **before** assigning the device.

```
[root@localhost ~]# lshw -c network -businfo
Bus info          Device          Class          Description
=====
pci@0000:00:03.0          network        Virtio network device
virtio@0            eth0           network        Ethernet interface
[root@localhost ~]#
```

5. **After** assigning the device, we can see the new virtual function network adapter inside the virtual machine at the guest layer.

```
[root@localhost ~]# lshw -c network -businfo
Bus info          Device          Class          Description
=====
pci@0000:00:03.0          network        Virtio network device
virtio@0            eth0           network        Ethernet interface
pci@0000:00:07.0    ens7           network        X550 Virtual Function
[root@localhost ~]# lspci -vmmks 00:07.0
Slot:    00:07.0
Class:   Ethernet controller
Vendor:  Intel Corporation
Device:  X550 Virtual Function
SVendor: ASRock Incorporation
SDevice: Device 1563
PhySlot: 7
Driver:  ixgbevf
Module:  ixgbevf
[root@localhost ~]#
```

That's all! Now, you can proceed with network interface configuration as in any regular virtual machine.