

Virtualization & Servers

VMware vSphere and ESXi, Veeam Backup, and notes for HPE and Dell hardware.

- [VMware vSphere & ESXi](#)
 - [VMWare - Delete inaccessible datastore](#)
 - [VMWare - Install Windows 11 on vCenter 8 without TMP 2.0](#)
 - [VMWare - Copy and Paste from Console to VM](#)
 - [VMWare - Fix VMware View Horizon the redo log of .vmdk is corrupt](#)
 - [ESXi - UPS Nut Installation](#)
- [Veeam Backup](#)
 - [Veeam Upgrade](#)
 - [Veeam Duplicate UUID](#)
- [Server Hardware \(HPE & Dell\)](#)
 - [HPE - iLO4 Reflash after Not Responsive 380p Gen8](#)
 - [HPE - ProLiant DL380p Gen8 GPU Install](#)
 - [Dell - Hard Drive Capacity Difference Between Drives](#)
 - [DELL - XPS Killer Wifi](#)

VMware vSphere & ESXi

Installing, upgrading and managing VMware vSphere and ESXi hosts.

VMWare - Delete inaccessible datastore

Rename the store

Use the web client to rename each store and a DEL in front of the name

Access the shell of the appliance server and then the database

VMware vCenter Server 8.0.2.00300

Type: vCenter Server with an embedded Platform Services Controller

Connected to service

* List APIs: "help api list"

* List Plugins: "help pi list"

* Launch BASH: "shell"

Command> shell

Shell access is granted to root

root@vcenter [~]#

```
shell:/opt/vmware/vpostgres/current/bin/psql -d VCDB -U postgres
```

View the stores

This will give you the ID that is needed to clean up

```
SELECT * FROM vpx_entity;
```

Delete the store

This is the delete for cleaning up the inaccessible datastore

```
DELETE FROM vpx_ds_assignment WHERE ds_id=1015;  
DELETE FROM vpx_datastore WHERE id=1015;  
DELETE FROM vpx_vm_ds_space WHERE ds_id=1015;  
DELETE from vpx_entity where id=1015;
```

Change the ID and then run the following

```
DELETE FROM vpx_ds_assignment WHERE ds_id=5015;DELETE FROM vpx_datastore WHERE id=5015;DELETE  
FROM vpx_vm_ds_space WHERE ds_id=5015;DELETE from vpx_entity where id=5015;
```


Restart the services

```
\q;service-control --stop --all && service-control --start --all
```

VMWare - Install Windows 11 on vCenter 8 without TMP 2.0

Create a VMWare to host Windows 11 without having a virtual TMP or device installed.

Once you are at this step you have to hit "Shift-F10" to open up the command prompt




Once the command prompt window is opened type "regedit"


```
C:\> Administrator: X:\windows\system32\cmd.exe
Microsoft Windows [Version 10.0.22000.318]
(c) Microsoft Corporation. All rights reserved.

X:\Sources>regedit
```


Once opened then navigate to "HKEY_LOCAL_MACHINE, SYSTEM, Setup" and then right-click and add a "New, Key"




Add the key called "LabConfig" then right-click on the "LabConfig" and add a new "DWORD Value"



Add the new value called "ByPassTPMCheck" with a hexadecimal of "1"



Close every window then click on "Install Now" but do not reboot or the new added key will be lost.




VMWare - Copy and Paste from Console to VM

On any Powered Off VM, click the "**Advanced Parameters**" tab.

1. Click filter icon of Attribute column. Search for 3 items' names, check if they already exist. If yes, edit their corresponding values. If not, add them from the textboxes above.

Image not found or type unknown




Name	Value
isolation.tools.copy.disable	FALSE

isolation.tools.paste.disable	FALSE
isolation.tools.setGUIOptions.enable	TRUE

2. Click "**OK**".
3. Power on the VM
4. Then use Copy/Paste directly on Windows/Linux/any other platform.

VMWare - Fix VMware View Horizon the redo log of .vmdk is corrupt



To disable the VAAI on the host follow the following steps:

1. Open Vsphere Client and connect to Vcenter Server.
2. Navigate to the ESXi Host and then Click on the Configuration Tab.
3. Click on Advanced Settings.
4. Click on DataMover.
5. Change the value of DataMover.HardwareAcceleratedMove setting from 1 to 0.
6. Change the value of DataMover.HardwareAcceleratedInit setting from 1 to 0.
7. Click on VMFS3.
8. Change the value of VMFS3.HardwareAcceleratedLocking setting from 1 to 0.
9. Click Ok.
10. Repeat the steps 1-9 on all ESXi Hosts in the cluster.

This should resolve the error once and for all. More info about VAAI can be found [here](#).

Source link [here](#)

ESXi - UPS Nut Installation

Vmware ESXi NUT Client Installation and Configuration

Requirements

- You must have a NUT server connected to a UPS and configured in your environment.
- SSH must be enabled on your ESXi installation
- Community acceptance level must be enabled on your ESXi installation in order to install the client

Download

Download the latest **NutClient-ESXi-x.x.x-x.x.x-offline_bundle.zip** package from:

<https://github.com/rgc2000/NutClient-ESXi/releases>.

Install

- Secure copy the **NutClient-ESXi-x.x.x-x.x.x-offline_bundle.zip** to your ESXi server's /tmp directory by either using WinSCP/pscp in Windows or scp in Linux.
- Set ESXi Community Acceptance level:

```
esxcli software acceptance set --level=CommunitySupported
```

If you get the following error:

```
[AcceptanceConfigError]  
Secure Boot enabled: Cannot change acceptance level to community.  
Please refer to the log file for more details.
```

You must disable Secure Boot in your computer BIOS and re-try the installation again.

- Install bundle:

```
esxcli software vib install -d /tmp/NutClient-ESXi-x.x.x-x.x.x-offline_bundle.zip
```

- If installation was successful you should see the following output:

```
Installation Result
Message: Operation finished successfully.
Reboot Required: false
VIBs Installed: Margar_bootbank_upsmon_x.x.x-x.x.x
VIBs Removed:
VIBs Skipped:
```

- You can now delete the files in the /tmp directory and disable the SSH service if desired.

ESXi Configuration

- In the ESXi Web client, navigate to **Host —> Manage —> System —> Advanced Settings**. In the **Search** box enter **UserVars.Nut (Figure 1)**.

Figure 1

image-1647104283497.png

- Configure the following variables:

NutUpsName: Name of the UPS on the NUT server (in the form of inverter_name@server_name or server_ip). Several inverters can be entered separated by a space. There will be no system shutdown until the last UPS still standing has given the shutdown command.

NutUser: Name of the NUT server login account

NutPassword: NUT Server Connection Account Password

NutFinalDelay: Seconds to wait after receiving the low battery event to shut down the system

NutSendMail: Set to 1 for the NUT client to send an e-mail to each important event of the UPS

NutMailTo: E-mail address to send UPS events to

- In the ESXi Web client, navigate to **Host —> Manage —> Services —> NutClient —> Actions —> Policy —> Start and Stop with Host (Figure 2)**.

Figure 2

image-1647104294686.png

- In the ESXi Web client, navigate to **Host —> Manage —> Services —> NutClient —> Actions —> Start (Figure 3).**

Figure 3

[image-1647104306878.png](#)

Tips

- Use the ESXi host configuration tab in the vSphere Client to decide how to start and stop (or suspend) virtual machines. This order will be respected by the UPS shutdown procedure.
- The clean shutdown of the OS in the virtual machines is only possible if the vmware tools are installed.
- To uninstall the NUT client, use the following command:

```
esxcli software vib remove -n upsmon
```

- To update the NUT client, use the following command:

```
esxcli software vib update -d /tmp/NutClient-ESXi-x.x.x-x.x.x-offline_bundle.zip
```

- To estimate the time needed for the server to shut down on UPS alert, type the command below on the host ESXi (by ssh or on the console). The shutdown procedure is immediately started:

```
/opt/nut/sbin/upsmon -c fsd
```

- If the NUT Client is configured correctly, the ESXi /var/log/syslog.log should have a message similar to below where ups@UPSHOST is the ups name and the UPS host you setup earlier :

```
2019-09-22T13:28:07Z upsmon[2111424]: Communications with UPS ups@UPSHOST established
```

Veeam Backup

Veeam Backup & Replication setup and fixes.

Veeam Upgrade

https://www.reddit.com/r/Veeam/comments/1b84qg4/migrate_11_to_12_on_new_hardware/

- Disable all jobs, perform necessary backups and snapshots of existing v11 server
- Perform upgrade to v12
- Post upgrade, take configuration backup
- On new v12 server, perform configuration restore
- Enable jobs
- Decom old VBR server/DB.

Or you can install the same v11 version on the new server, perform the config backup and restore first, then do an in place upgrade of the new server to v12.

Make sure you upgrade any VeeamONE and Enterprise Manager first.

Licensing, you can use the same license on both servers as long as its part of the migration process...only running the same license on two different VBR server in production at the same time is an issue.

Veeam Duplicate UUID

This is to bypass the UUID check if you have restored an old back to new hardware for example a laptop or physical PC.

REGEDIT

To 'EnableSystemUuidFailover' key should be added on each Veeam Agent host with non-unique ID.

This is the key enabling Veeam failover logic.

Name: EnableSystemUuidFailover
Path: HKLM\SOFTWARE\Veeam\
Type: DWORD
Value: 1

Then a unique ID should be assigned manually on each Agent station.
Another Veeam registry should be added on each host for it:

Name: SystemUuid
Path: HKLM\SOFTWARE\Veeam\
Type: REG_SZ
Value: {a custom UUID}

Custom UUID should correspond with the standard UUID format rules, brackets are required.

You can use some online generator for it (<https://www.uuidgenerator.net/>) or put some your own strings there.

Rescan protection group and you are deployed now

Disable "Windows recovery image file not found" warning

You can add AgentsDisableRECollectionWarning key here
HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication to disable this warning.
Name: AgentsDisableRECollectionWarning
Path: HKLM\SOFTWARE\Veeam\
Type: DWORD
Value: 1

Server Hardware (HPE & Dell)

Firmware, management and driver notes for HPE and Dell hardware.

HPE - iLO4 Reflash after Not Responsive 380p Gen8

iLO 4 firmware in Direct Mode

Download Online ROM Flash Component for Linux - HPE Integrated Lights-Out 4
(CP#####.scexe)

https://support.hpe.com/connect/s/product?language=en_US&ismnp=0&l5oid=5219994&cep=on&kmpmoid=5219994&tab=driversAndSoftware&driversAndSoftwareFilter=8000029

Attached is 2.82 version of iLO4 [CP053894.scexe](#)

Download Service Pack for ProLiant 2016+ (search on torrents) version 8.1 or use this ISO

Using HP USB Key Utility for Windows, we create a bootable flash drive from the SPP image.

For Windows use RUFUS <https://rufus.ie/en/> to create a usb boot from the SPP .iso file

For Linux installation use <https://unetbootin.github.io/> to create a usb boot

Copy [CP053894.scexe](#) to the finished flash drive in /hp/swpackages

Turn off the power, remove the case. On the system board, find the main dip switch (System maintenance switch) and turn the first (S1) switch to the ON position

Load the server from the SPP flash drive

Once on the image screen you need to select Interactive mode and when the main menu loads very slowly

press ctrl + alt + d + b + x to launch the console

Run:


You need to enter a simple magic command:

```
rmmod hpilo
```

then Run:

```
sh /mnt/bootdevice/hp/swpackages/CP027578.scexe --direct
```

After successful iLO flashing, turn off the power and return the first (S1) switch back (OFF).



HPE - ProLiant DL380p Gen8

GPU Install

Supported GPUs for the HP ProLiant DL380p Gen8 are limited by the server's specific PCIe risers, 225W power constraints, and passive cooling requirements. Official and community-tested options include the **NVIDIA GRID K1, K2, Tesla P100, Quadro K6000**, or consumer cards like the **NVIDIA Tesla M4/M60**

Because the DL380 G8 requires specific hardware components for proper installation, you will need to pay attention to the following details:

- **Hardware & Riser Kits:** You will need the HP Double Wide Riser Cage Kit (662885-B21) and the Secondary GPU Riser (728543-B21) to accommodate double-wide or full-height GPUs.
- **Power Cables:** Standard server power cables will not fit most NVIDIA GPUs. You will need a custom 10-pin to 8-pin GPU power cable to draw power from the motherboard.
- **Power Supply:** Operating dual GPUs requires dual 1200W power supply units (PSUs) and two physical processors installed
- **BIOS Configuration:** For advanced GPUs (like the Tesla P100), you must enter the BIOS, press **Ctrl + A** at the main screen to reveal the "Service Options," and enable **PCI Express 64-bit BAR Support**.
- **iLO & POST Issues:** Many modern GPUs (like RTX series) can trigger POST code 5808, halting the boot process. Ensure your server has the latest BIOS update and iLO firmware to minimize compatibility issues.
- **Card Compatibility:** The server officially supports specific enterprise cards like the NVIDIA Quadro K4000/K5000/P2000 and Grid K1/K2. Consumer cards (like the GTX 1050, GTX 1650, or RTX series) often require BIOS modifications, external power adapters, and custom chassis modification

Dell - Hard Drive Capacity Difference Between Drives

The capacity difference between the two drives is caused by **differing logical block/sector formatting** or **OEM vendor provisioning (HPA / Max LBA limits)**.

Even within the same Western Digital / Ultrastar DC HC520 12TB family, SAS drives are deployed with different low-level format structures:

1. Different Sector Formats (512e vs. 4Kn vs. Enterprise 520/528 Bytes)

- **Standard 12TB:** Typically reports around **11,175.5 GB** (\$11.17\text{ TiB}\$) when formatted with standard 512-byte emulation (**512e**) or 4096-byte native (**4Kn**) sectors.
- **Enterprise SAN/EMC Formats:** Drives pulled from proprietary SAN enclosures (EMC, NetApp, Huawei) often use **520-byte or 528-byte sectors** (or Type 1/2 Data Integrity Field protection). When translated or truncated by a Dell PERC/HBA, the usable LBA count drops, showing up around **10,948.5 GB**.

2. OEM Firmware & Provisioning (HPA / Max LBA)

If one drive came from an OEM vendor (e.g., NetApp, EMC, Cisco, HPE) and the other is a generic Dell/retail drive, the OEM firmware may restrict the drive's total accessible Logical Block Addressing (LBA) count via Host Protected Area (HPA) or factory provisioning.

How to Fix / Match Them

To get disk **0:1:3** to match **0:1:4** at full capacity:

1. Check the Sector Size:

Boot into a Linux live environment (or pass the drive through via HBA/IT-mode) and run:

Bash

■ ■

```
smartctl -i /dev/sdX

# or

sg_scan -i
```

Check the `Logical block size` and `Formatted with type 1/2/3 protection` values.

2. Low-Level Format to Standard 512e / 4Kn:

Using the `sg3_utils` package in Linux, reformat the drive to standard 512-byte or 4096-byte sectors:

Bash

■ ■

```
# Reformat to standard 512-byte sectors:

sg_format --format --size=512 /dev/sdX
```

(Note: This is a low-level format that will erase all data on that specific disk and takes several hours for a 12TB drive.)

3. Reset Max LBA (if sector size is already 512):

ATA/SATA drives only. `hdparm -N` sends ATA commands, so on a SAS drive (like the HC520 SAS) it fails with 'SG_IO: bad/missing sense data, sb[]: 72 05 20 ...' (Illegal Request, invalid command) and prints bogus output like 'max sectors = 0/513, HPA is enabled'. Ignore that; SAS drives have no HPA. Use step 4 instead.

If the drive was already 512 bytes, reset any HPA/LBA clamp:

Bash

■ ■

```
hdparm -N /dev/sdX

# To restore full native capacity:

hdparm -N p[max_lba_value] /dev/sdX
```

4. SAS drives: verify and fix capacity with sg3_utils

Find the right device first (sg numbers can change after a reboot):

```
lsscsi -g

# or

sg_scan -i
```

Check capacity and protection:

```
sg_readcap --long /dev/sgX
```

A full 12TB HC520 at 512-byte sectors shows: Number of logical blocks=23437770752, Logical block length=512 bytes, prot_en=0, Device size 12.00 TB (about 11175.5 GB in the PERC).

If the block count is short (sg_format keeps the existing block count when the sector size doesn't change), resize to the maximum. This takes seconds, not hours:

```
sg_format --resize --count=-1 /dev/sgX
```

If prot_en=1 (Type 1/2 protection still on, which PERC controllers often reject), do a full format that clears protection and sets max size:

```
sg_format --format --size=512 --fmtinfo=0 --count=-1 /dev/sgX
```

5. Get the PERC to pick up the new size

Power-cycle the server (or rescan), then in the PERC config utility or iDRAC storage page clear any foreign config on the drive and confirm it shows 11175.5 GB before adding it to a virtual disk.

6. Optional: format to 4096-byte native sectors (4Kn)

Before you start, check that everything using the drive supports 4Kn. Don't mix 4Kn and 512e drives in the same PERC virtual disk; to boot from a 4Kn disk the server must use UEFI, not BIOS mode; and some older OS and hypervisor versions don't support 4Kn disks. SAS drives (sg3_utils). This erases the whole drive and takes several hours on a 12TB disk:

```
lsscsi -g      # confirm the right /dev/sgX first
sg_format --format --size=4096 --fmtinfo=0 /dev/sgX
```

When the sector size changes, sg_format sets the block count to the drive's maximum on its own, so --count isn't needed.

Verify:

```
sg_readcap --long /dev/sgX
```

A full 12TB HC520 at 4Kn shows: Logical block length=4096 bytes, Number of logical blocks=2929721344, prot_en=0, Device size 12.00 TB (the same total size as 512e, just one-eighth the number of blocks).

SATA drives: sg_format doesn't apply. If the drive's firmware supports changing the sector size, use:

```
hdparm -l /dev/sdX | grep -i 'sector size'
hdparm --set-sector-size 4096 --please-destroy-my-drive /dev/sdX
```

If hdparm says it isn't supported, the drive's sector size can't be changed this way. Then do step 5 again (power-cycle, clear foreign config, confirm the size in the PERC).

7. Check for and handle bad sectors (SAS)

A full sg_format already certifies every sector and remaps bad ones to spare sectors (reallocation). Remapped sectors go on the drive's grown defect list. They can't be recovered, but they stop causing errors.

Keep certification on during the format (it's the default, so don't add --ffmt or --dcrt):

```
sg_format --format --size=4096 --fmtinfo=0 /dev/sgX
```

After the format, check the grown defect list and error counters:

```
sg_reassign --grown /dev/sgX  
smartctl -a /dev/sgX
```

In the smartctl output, check 'Elements in grown defect list' and the read/write 'Total uncorrected errors' values.

Optional full read test (more than 20 hours on a 12TB drive):

```
smartctl -t long /dev/sgX  
smartctl -l selftest /dev/sgX # check the result after it finishes
```

To remap one specific bad LBA by hand (this erases the data in that block):

```
sg_reassign --address=<LBA> /dev/sgX
```

Rule of thumb: a few dozen grown defects that stay the same over time is fine. If the count keeps going up, or there are any uncorrected errors, the drive is failing and should be replaced instead of reused.

Once reformatted, the Dell PERC controller will detect `0:1:3` at the full **11175.5 GB** size.

DELL - XPS Killer Wifi

I was only getting 1.5mb transfer and 60mb web download

WARNING -- Changing the registry can break your machine if done wrong!!! Do so at your own risk!

I had (sounds like) the same WiFi issues you are having. I was able to "fix" it with a registry hack. Go to this

key Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Power\PowerSettings**19cbb8fa-5279-450e-9fac-8a3d5fedd0c1** and create a value (if not there) Attributes with a DWORD value of 2.

Go a little further down, under the above key

to Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Power\PowerSettings\19cbb8fa-5279-450e-9fac-8a3d5fedd0c1\12bbebe6-58d6-4636-95bb-3217ef867c1a\

DefaultPowerSchemeValues\381b4222-f694-41f0-9685-ff5bb260df2e and add the same Attributes value here or set it to 2 if it's already there. Under this key, change the DCSettingIndex to 0 (zero) to match the ACSettingIndex.

This makes the WiFi card use full power on battery under the Balanced Power Plan which is the default plan.

Also mentioned in my previous posts was to remove all the Killer software and use the Intel drivers instead. The current version works fine. Of course download the Intel drivers before removing the Killer stuff. This did not fix the issue by itself but trust the Intel drivers more.