

# Upgrade from 7 to 8

---

## Contents

---

Introduction

New installation

Breaking Changes

In-place upgrade

- Prerequisites

- Testing the Upgrade

- Actions step-by-step

  - Continuously use the **pve7to8** checklist script

  - Move important Virtual Machines and Containers

  - Update the configured APT repositories

    - Update Debian Base Repositories to Bookworm

    - Add the Proxmox VE 8 Package Repository

    - Update the Ceph Package Repository

    - Refresh Package Index

  - Upgrade the system to Debian Bookworm and Proxmox VE 8.0

  - Check Result & Reboot Into Updated Kernel

- After the Proxmox VE upgrade

  - For Clusters

- Checklist issues

  - proxmox-ve package is too old

- Known Upgrade Issues

  - General

    - Upgrade wants to remove package 'proxmox-ve'

  - Third-party Storage Plugins

  - Older Hardware and New 6.2 Kernel and Other Software

  - 6.2 Kernels regressed KSM performance on multi-socket NUMA systems

  - GRUB Might Fail To Boot From LVM in UEFI Mode

  - VM Live-Migration

    - VM Live-Migration with different host CPUs

    - VM Live-Migration with Intel Skylake (or newer) CPUs

  - Network

    - Network Interface Name Change

    - Network Setup Hangs on Boot Due to NTPsec Hook

  - cgroup V1 Deprecation

  - NVIDIA vGPU Compatibility

  - Systemd-boot (for ZFS on root and UEFI systems only)

- Troubleshooting

Failing upgrade to "bookworm"  
Unable to boot due to grub failure  
External links

## Introduction

Proxmox VE 8.x introduces several new major features. You should plan the upgrade carefully, make and verify backups before beginning, and test extensively. Depending on the existing configuration, several manual steps—including some downtime—may be required.

**Note:** A valid and tested backup is *always* required before starting the upgrade process. Test the backup beforehand in a test lab setup.

In case the system is customized and/or uses additional packages or any other third party repositories/packages, ensure those packages are also upgraded to and compatible with Debian Bookworm.

In general, there are two ways to upgrade a Proxmox VE 7.x system to Proxmox VE 8.x:

- A new installation on new hardware (restoring VMs from the backup)
- An in-place upgrade via apt (step-by-step)

## New installation

- Backup all VMs and containers to an external storage (see [Backup and Restore](#)).
- Backup all files in /etc
  - required:** files in /etc/pve, as well as /etc/passwd, /etc/network/interfaces, /etc/resolv.conf, and anything that deviates from a default installation.
- Install latest Proxmox VE 8.x from the ISO (this will delete all data on the existing host).
- Empty the browser cache and/or force-reload (**CTRL** + **SHIFT** + **R**, or for MacOS **⌘** + **Alt** + **R**) the Web UI.
- Rebuild your cluster, if applicable.
- Restore the file /etc/pve/storage.cfg (this will make the external storage used for backup available).
- Restore firewall configs /etc/pve/firewall/ and /etc/pve/nodes/<node>/host.fw (if applicable).
- Restore all VMs from backups (see [Backup and Restore](#)).

Administrators comfortable with the command line can follow the procedure [Bypassing backup and restore when upgrading](#), if all VMs/CTs are on a single shared storage.

## Breaking Changes

See the release notes for breaking (API) changes: <https://pve.proxmox.com/wiki/Roadmap#8.0-known-issues>

# In-place upgrade

In-place upgrades are carried out via apt. Familiarity with apt is required to proceed with this upgrade method.

## Prerequisites

---

- Upgraded to the latest version of Proxmox VE 7.4 on all nodes.  
Ensure your node(s) have correct package repository configuration (web UI, Node -> Repositories) if your pve-manager version isn't at least 7.4-15.
- Hyper-converged Ceph: upgrade any Ceph Octopus or Ceph Pacific cluster to Ceph 17.2 Quincy **before** you start the Proxmox VE upgrade to 8.0.  
Follow the guide [Ceph Octopus to Pacific](#) and [Ceph Pacific to Quincy](#), respectively.
- Co-installed Proxmox Backup Server: see the [Proxmox Backup Server 2 to 3 upgrade how-to](https://pbs.proxmox.com/wiki/index.php/Upgrade_from_2_to_3) ([https://pbs.proxmox.com/wiki/index.php/Upgrade\\_from\\_2\\_to\\_3](https://pbs.proxmox.com/wiki/index.php/Upgrade_from_2_to_3))
- Reliable access to the node. It's recommended to have access over a host independent channel like iKVM/IPMI or physical access.  
If only SSH is available we recommend testing the upgrade on an identical, but non-production machine first.
- A healthy cluster
- Valid and tested backup of all VMs and CTs (in case something goes wrong)
- At least 5 GB free disk space on the root mount point.
- Check [known upgrade issues](#)

## Testing the Upgrade

---

An upgrade test can be easily performed using a standalone server. Install the Proxmox VE 7.4 ISO on some test hardware, then upgrade this installation to the latest minor version of Proxmox VE 7.4 (see [Package repositories](#)). To replicate the production setup as closely as possible, copy or create all relevant configurations to the test machine, then start the upgrade. It is also possible to install Proxmox VE 7.4 in a VM and test the upgrade in this environment.

## Actions step-by-step

---

The following actions need to be carried out from the command line of each Proxmox VE node in your cluster

**Perform the actions via console or ssh; preferably via console to avoid interrupted ssh connections. Do not carry out the upgrade when connected via the virtual console offered by the GUI; as this will get interrupted during the upgrade.**

Remember to ensure that a valid backup of all VMs and CTs has been created before proceeding.

### Continuously use the **pve7to8** checklist script

A small checklist program named **pve7to8** is included in the latest Proxmox VE 7.4 packages. The program will provide hints and warnings about potential issues before, during and after the upgrade process. You can call it by executing:

```
pve7to8
```

To run it with **all** checks enabled, execute:

```
pve7to8 --full
```

Make sure to run the full checks at least once before the upgrade.

This script only **checks** and reports things. By default, no changes to the system are made and thus, none of the issues will be automatically fixed. You should keep in mind that Proxmox VE can be heavily customized, so the script may not recognize all the possible problems with a particular setup!

It is recommended to re-run the script after each attempt to fix an issue. This ensures that the actions taken actually fixed the respective warning.

## Move important Virtual Machines and Containers

If any VMs and CTs need to keep running for the duration of the upgrade, migrate them away from the node that is being upgraded.

Migration compatibility rules to keep in mind when planning your cluster upgrade:

- A migration of a VM or CT from an older version of Proxmox VE to a newer version will always work.
- A migration from a newer Proxmox VE version to an older version may work, but is generally not supported.

## Update the configured APT repositories

First, make sure that the system is using the latest Proxmox VE 7.4 packages:

```
apt update
apt dist-upgrade
pveversion
```

The last command should report at least 7.4-15 or newer.

## Update Debian Base Repositories to Bookworm

Update all Debian and Proxmox VE repository entries to Bookworm.

```
sed -i 's/bullseye/bookworm/g' /etc/apt/sources.list
```

Ensure that there are no remaining Debian Bullseye specific repositories left, if you can use the # symbol at the start of the respective line to comment these repositories out. Check all files in the `/etc/apt/sources.list.d/pve-enterprise.list` and `/etc/apt/sources.list` and see [Package\\_Repositories](#) for the correct Proxmox VE 8 / Debian Bookworm repositories.

## Add the Proxmox VE 8 Package Repository

```
echo "deb https://enterprise.proxmox.com/debian/pve bookworm pve-enterprise" > /etc/apt/sources.list.d/pve-enterprise.list
```

For the no-subscription repository, see [Package Repositories](#). Rather than commenting out/removing the PVE 7.x repositories, as was previously mentioned, you could also run the following command to update to the Proxmox VE 8 repositories:

```
sed -i -e 's/bullseye/bookworm/g' /etc/apt/sources.list.d/pve-install-repo.list
```

## Update the Ceph Package Repository

**Note:** For hyper-converged ceph setups only, check the ceph panel and configured repositories in the Web UI of this node, if unsure.

Replace any ceph.com repositories with proxmox.com ceph repositories.

NOTE: At this point a hyper-converged Ceph cluster installed directly in Proxmox VE **must run Ceph 17.2 Quincy**, if not you need to upgrade Ceph first before upgrading to Proxmox VE 8 on Debian 12 Bookworm! You can check the current ceph version in the Ceph panel of each node in the Web UI of Proxmox VE.

With Proxmox VE 8 there also exists an enterprise repository for ceph, providing the best choice for production setups.

```
echo "deb https://enterprise.proxmox.com/debian/ceph-quincy bookworm enterprise" > /etc/apt/sources.list.d/ceph.list
```

If updating fails with a 401 error, you might need to refresh the subscription first to ensure new access to ceph is granted, do this via the Web UI or `pvesubscription update --force`.

If you do not have any subscription you can use the no-subscription repository:

```
echo "deb http://download.proxmox.com/debian/ceph-quincy bookworm no-subscription" > /etc/apt/sources.list.d/ceph.list
```

If there is a backports line, remove it - the upgrade has not been tested with packages from the backports repository installed.

## Refresh Package Index

Update the repositories' package index:

```
apt update
```

## Upgrade the system to Debian Bookworm and Proxmox VE 8.0

Note that the time required for finishing this step heavily depends on the system's performance, especially the root filesystem's IOPS and bandwidth. A slow spinner can take up to 60 minutes or

more, while for a high-performance server with SSD storage, the dist-upgrade can be finished in under 5 minutes.

Start with this step, to get the initial set of upgraded packages:

```
apt dist-upgrade
```

During the above step, you will be asked to approve changes to configuration files, where the default config has been updated by their respective package.

It's suggested to check the difference for each file in question and choose the answer accordingly to what's most appropriate for your setup.

Common configuration files with changes, and the recommended choices are:

- `/etc/issue` -> Proxmox VE will auto-generate this file on boot, and it has only cosmetic effects on the login console.  
Using the default "No" (keep your currently-installed version) is safe here.
- `/etc/lvm/lvm.conf` -> Changes relevant for Proxmox VE will be updated, and a newer config version might be useful.  
If you did not make extra changes yourself and are unsure it's suggested to choose "Yes" (install the package maintainer's version) here.
- `/etc/ssh/sshd_config` -> If you have not changed this file manually, the only differences should be a replacement of `ChallengeResponseAuthentication no` with `KbdInteractiveAuthentication no` and some irrelevant changes in comments (lines starting with #).  
If this is the case, both options are safe, though we would recommend installing the package maintainer's version in order to move away from the deprecated `ChallengeResponseAuthentication` option. If there are other changes, we suggest to inspect them closely and decide accordingly.
- `/etc/default/grub` -> Here you may want to take special care, as this is normally only asked for if you changed it manually, e.g., for adding some kernel command line option.  
It's recommended to check the difference for any relevant change, note that changes in comments (lines starting with #) are not relevant.  
If unsure, we suggested to selected "No" (keep your currently-installed version)

## Check Result & Reboot Into Updated Kernel

If the dist-upgrade command exits successfully, you can re-check the pve7to8 checker script and reboot the system in order to use the new Proxmox VE kernel.

Please note that you should reboot even if you already used the 6.2 kernel previously, through the opt-in package on Proxmox VE 7. This is required to guarantee the best compatibility with the rest of the system, as the updated kernel was (re-)build with the newer Proxmox VE 8 compiler and ABI versions.

## After the Proxmox VE upgrade

---

Empty the browser cache and/or force-reload (`CTRL` + `SHIFT` + `R`), or for MacOS `⌘` + `Alt` + `R`) the Web UI.

## For Clusters

- Check that all nodes are up and running on the latest package versions.  
If not, continue the upgrade on the next node, start over at [#Preconditions](#)

## Checklist issues

---

### proxmox-ve package is too old

Check the configured package repository entries; they still need to be for Proxmox VE 7.x and Bullseye at this step (see [Package Repositories](#)). Then run

```
apt update
```

followed by

```
apt dist-upgrade
```

to get the latest Proxmox VE 7.x packages **before** upgrading to PVE 8.x

## Known Upgrade Issues

---

### General

As a Debian based distribution, Proxmox VE is affected by most issues and changes affecting Debian. Thus, ensure that you read the [upgrade specific issues for Debian Bookworm](https://www.debian.org/releases/bookworm/amd64/release-notes/ch-information.en.html) (<https://www.debian.org/releases/bookworm/amd64/release-notes/ch-information.en.html>), for example the [transition from classic NTP to NTPsec](https://www.debian.org/releases/bookworm/amd64/release-notes/ch-information.en.html#changes-to-packages-that-set-the-system-clock) (<https://www.debian.org/releases/bookworm/amd64/release-notes/ch-information.en.html#changes-to-packages-that-set-the-system-clock>)

Please also check the known issue list from the Proxmox VE 8.0 changelog: <https://pve.proxmox.com/wiki/Roadmap#8.0-known-issues>

### Upgrade wants to remove package 'proxmox-ve'

If you have installed Proxmox VE on top of a plain Debian Bullseye (without using the Proxmox VE ISO), you may have installed the package 'linux-image-amd64', which conflicts with current 7.x setups. To solve this, you have to remove this package with

```
apt remove linux-image-amd64
```

before the dist-upgrade.

---

## Third-party Storage Plugins

If you use any external storage plugin you need to wait until the plugin author adapted it for Proxmox VE 8.0.

---

## Older Hardware and New 6.2 Kernel and Other Software

Compatibility of old hardware (released  $\geq$  10 years ago) is not as thoroughly tested as more recent hardware. For old hardware we highly recommend testing compatibility of Proxmox VE 8 with identical (or at least similar) hardware before upgrading any production machines.

Ceph has been reported to run into "illegal instruction" errors with at least AMD Opteron 2427 (released in 2009) and AMD Turion II Neo N54L (released in 2010) CPUs.

We will expand this section with potential pitfalls and workarounds once they arise.

---

## 6.2 Kernels regressed KSM performance on multi-socket NUMA systems

Kernels based on 6.2 have a degraded Kernel Samepage Merging (KSM) performance on multi-socket NUMA systems, depending on the workload this can result in a significant amount of memory that is not deduplicated anymore. This issue went unnoticed for a few kernel releases, making a clean backport of the fixes made for 6.5 hard to do without some general fall-out.

Until a targeted fix for the upstream LTS 6.1 kernel is found, the current recommendation is to keep your multi-socket NUMA systems that rely on KSM on Proxmox VE 7 with its 5.15 based kernel. We plan to change the default kernel to a 6.5 based kernel in 2023'Q4, which will also resolve this issue.

---

## GRUB Might Fail To Boot From LVM in UEFI Mode

Due to a [bug in grub](https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=987008) (https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=987008) in PVE 7 and before, grub may fail to boot from LVM with an error message `disk `lvmid/...' not found`. When booting in UEFI mode, you need to ensure that the new grub version containing the fix is indeed used for booting the system.

Systems with Root on ZFS and systems booting in legacy mode are not affected.

On systems booting in EFI mode with root on LVM, install the correct grub meta-package with:

```
[ -d /sys/firmware/efi ] && apt install grub-efi-amd64
```

For more details see [the relevant wiki page](#).

---

## VM Live-Migration

## VM Live-Migration with different host CPUs

Live migration between nodes with different CPU models and especially different vendors can cause problems, such as VMs becoming unresponsive and causing high CPU utilization.

We recommend testing live migration with a non-production VM first when upgrading. For this reason, we highly encourage using homogenous setups in clusters that use live migration.

## VM Live-Migration with Intel Skylake (or newer) CPUs

Previous 6.2 kernels had problems with incoming live migrations when all of the following were true:

- VM has a restricted CPU type (e.g., qemu64) – using CPU type host or Skylake-Server is ok.
- the source host uses an Intel CPU from Skylake Server, Tiger Lake Desktop, or equivalent newer generation.
- the source host is booted with a kernel version 5.15 (or older) (e.g. when upgrading from Proxmox VE 7.4)

In this case, the VM could hang after migration and use 100% of one or more vCPUs.

This was fixed with pve-kernel-6.2.16-4-pve in version 6.2.16-5. So make sure your target host is booted with this (or a newer) kernel version if the above points apply to your setup.

---

## Network

### Network Interface Name Change

Due to the new kernel recognizing more features of some hardware, like for example virtual functions, and interface naming often derives from the PCI(e) address, some NICs may change their name, in which case the network configuration needs to be adapted.

In general, it's recommended to either have an independent remote connection to the Proxmox VE's host console, for example, through IPMI or iKVM, or physical access for managing the server even when its own network doesn't come up after a major upgrade or network change.

### Network Setup Hangs on Boot Due to NTPsec Hook

If both ntpsec and ntpsec-ntpdate are installed, the network will fail to come up cleanly on boot and hang, but will work if triggered manually (e.g., using `ifreload -a`). Even if the two packages are not already present before the upgrade, they will be installed during upgrade if both ntp and ntpdate are present before the upgrade.

Since the chrony NTP daemon is used as default for new installations since Proxmox VE 7.0 the simplest solution might be switching to that via `apt install chrony`. If this is not possible, it suffices to keep ntpsec but uninstall ntpsec-ntpdate (according to [its package description \(http://packages.debian.org/bookworm/ntpsec-ntpdate\)](http://packages.debian.org/bookworm/ntpsec-ntpdate), that package is not necessary if ntpsec is installed). If the host is already hanging during boot, one quick workaround is to boot into recovery mode, enter the root password, run `chmod -x /etc/network/if-up.d/ntpsec-ntpdate` and reboot.

The root cause for the hang is that the script `/etc/network/if-up.d/ntpsec-ntpdate` installed by `ntpsec-ntpdate` causes `ifupdown2` to hang during boot if `ntpsec` is installed. For more information, see [bug #5009](https://bugzilla.proxmox.com/show_bug.cgi?id=5009) ([https://bugzilla.proxmox.com/show\\_bug.cgi?id=5009](https://bugzilla.proxmox.com/show_bug.cgi?id=5009)).

---

## **cgroup V1 Deprecation**

Reminder, since the previous major release Proxmox VE 7.0, the default is a pure `cgroupv2` environment. While Proxmox VE 8 did not change in this regard, we'd like to note that Proxmox VE 8 will be the last release series that supports booting into the old "hybrid" `cgroup` system ([http://pve.proxmox.com/pve-docs/chapter-pct.html#pct\\_cgroup\\_compat](http://pve.proxmox.com/pve-docs/chapter-pct.html#pct_cgroup_compat)), e.g. for compatibility with ancient Container OS.

That means that Containers running `systemd` version 230 (released in 2016) or older won't be supported at all in the next major release (Proxmox VE 9, ~ 2025 Q2/Q3). If you still run such container (e.g., CentOS 7 or Ubuntu 16.04), please use the Proxmox VE 8 release cycle as time window to migrate to newer, still supported versions of the respective Container OS.

---

## **NVIDIA vGPU Compatibility**

If you are using NVIDIA's GRID/vGPU technology, its driver must be compatible with the kernel you are using. Make sure you use at least GRID version 16.0 (driver version 535.54.06 - current as of July 2023) on the host before upgrading, since older versions (e.g. 15.x) are not compatible with kernel versions  $\geq 6.0$  and Proxmox VE 8.0 ships with at least 6.2.

---

## **Systemd-boot (for ZFS on root and UEFI systems only)**

Systems booting via UEFI from a ZFS on root setup should install the `systemd-boot` package after the upgrade. You will get a Warning from the `pve7to8` script after the upgrade if your system is affected - in all other cases you can safely ignore this point.

The `systemd-boot` was split out from the `systemd` package for Debian Bookworm based releases. It won't get installed automatically upon upgrade from Proxmox VE 7.4 as it can cause trouble on systems not booting from UEFI with ZFS on root setup by the Proxmox VE installer.

Systems which have ZFS on root and boot in UEFI mode will need to manually install it if they need to initialize a new ESP (see the output of `proxmox-boot-tool status` and the [relevant documentation](https://pve.proxmox.com/pve-docs/chapter-sysadmin.html#sysboot_determine_bootloader_used) ([https://pve.proxmox.com/pve-docs/chapter-sysadmin.html#sysboot\\_determine\\_bootloader\\_used](https://pve.proxmox.com/pve-docs/chapter-sysadmin.html#sysboot_determine_bootloader_used))).

Note that the system remains bootable even without the package installed.

It is not recommended installing `systemd-boot` on systems which don't need it, as it would replace `grub` as bootloader in its `postinst` script.

# Troubleshooting

---

## Failing upgrade to "bookworm"

Make sure that the repository configuration for Bookworm is correct.

If there was a network failure and the upgrade was only partially completed, try to repair the situation with

```
apt -f install
```

If you see the following message:

```
W: (pve-apt-hook) You are attempting to remove the meta-package 'proxmox-ve'!
```

then one or more of the currently existing packages cannot be upgraded since the proper Bookworm repository is not configured.

Check which of the previously used repositories (i.e. for Bullseye) do not exist for Bookworm or have not been upgraded to Bullseye ones.

If a corresponding Bookworm repository exists, upgrade the configuration (see also special remark for Ceph).

If an upgrade is not possible, configure all repositories as they were before the upgrade attempt, then run:

```
apt update
```

again. Then remove **all** packages which are currently installed from that repository. Following this, start the upgrade procedure again.

## Unable to boot due to grub failure

See Recover From Grub Failure

If your system was installed on ZFS using legacy BIOS boot before the Proxmox VE 6.4 ISO, incompatibilities between the ZFS implementation in grub and newer ZFS versions can lead to a broken boot. Check the article on switching to proxmox-boot-tool ZFS: Switch Legacy-Boot to Proxmox Boot Tool for more details.

## External links

---

Release Notes for Debian 12.0 (bookworm), 64-bit PC (<https://www.debian.org/releases/bookworm/amd64/release-notes/>)

---

This page was last edited on 23 October 2023, at 08:41.