DOCUMENT – JANUARY 2021

PRINTED 30 SEPTEMBER 2021

ESXI SYSTEM STORAGE FAQ

VMwareLifecycle Management

vmware

Table of Contents

ESXi System Storage FAQ

Overview

This article details frequently asked questions and answers about vSphere ESXi core storage, including Auto Deploy and vLCM.

ESXi System Storage

What has changed for ESXi system storage?

We have increased the boot bank sizes, and consolidated the system partitions and made them expandable. This article details these changes introduced with vSphere 7 and how that reflects on the boot media requirements to run vSphere 7.

What happens to the system storage layout when I upgrade my ESXi hosts to vSphere 7?

The ESXi system storage requirements are also applicable for ESXi upgrades (6.x > 7). However, the boot media requirements differ between a new vSphere 7 install, and an upgrade to vSphere 7. There's a requirement for boot media to run a 4 GB storage device at minimum, when upgrading to vSphere 7. Even though 4 GB boot media devices are supported, it is **strongly recommended** to adhere to the boot media requirements for a fresh vSphere 7 installation. Check out this article for detailed information.

What is supported for ESXi system storage?

ESXi 7.0 requires a boot disk of at least 32 GB of persistent storage such as HDD, SSD, or NVMe. Use USB, SD, and non-USB flash media devices only for ESXi boot bank partitions. A boot device must not be shared between ESXi hosts.

What is recommended for ESXi system storage?

Looking ahead, there will be new features that require ESXi hosts with high-quality persistent storage (local hdd/ssd or boot from SAN). The recommended ESXi 7.0 install options are the following:

- A disk with a minimum of 32 GB. The disk contains the boot partitions and ESX-OSData volume.
- A disk of 142 GB or larger. The disk contains the boot partitions, ESX-OSData volume, and VMFS datastore.

Are USB/SD boot devices being deprecated?

Legacy SD and USB devices are supported with some limitations listed below. To chose a proper SD or USB boot device, see Knowledge Base article 82515.

- You must provide an additional VMFS volume of at least 32 GB to store the ESX-OSData volume and required VMFS datastore. If the boot device is larger than 138 GB, the ESXi installer creates a VMFS volume automatically. Delete the VMFS datastore on USB and SD devices immediately after installation to prevent data corruption. For more information how to configure a persistent scratch partition, see Knowledge Base article 1033696.
- If the VMware Tools partition is stored locally, you must redirect it to the RAM disk. For more information, see Knowledge Base article 83376.
- You must use an SD flash device that is approved by the server vendor for the particular server model on which you want to install ESXi on an SD flash storage device. You can find a list of validated devices on partnerweb.vmware.com.
- Check Out Knowledge Base Article 85685

I see *bootbank cannot be found at path '/bootbank'* errors being seen after upgrading

to ESXi 7.0 Update 2. What's happening?

This is often because low-endurance SD or USB boot media is used. SD or USB devices have a small queue depth and due to a race condition in the ESXi storage stack, some I/O operations might not get to the device. Such I/Os queue in the ESXi storage stack and ultimately time out. For more information, see Knowledge Base article 83963.

- Upgrade to vSphere 7 Update 2c. Check out the release notes for more details.

I'm experiencing *VMFS-L Locker partition corruption* on ESXi hosts. What's happening?

This is likely be the result of using low-endurance SD or USB boot media for ESXi hosts. Files, for example vmtools, which are read from time to time by virtual machines could cause corruption of the SD or USB boot media due to things like read-disturb symptoms where neighbouring flash cells get affected from the read activity. This means that even though you might be reading from one sector, it could affect other sectors that have not been accessed in the flash device - hence we see filesystem metadata sectors getting corrupted.

We can reduce heavy access to SD or USB boot media by the steps explained in Knowledge Base article 83376. This article details how the `ToolsRamdisk` advanced options is used to copy vmtools, which exists on the SD or USB boot media, to a RAMdisk. This behaviour is of mainly seen on vSAN or VDI environments.

I'm using SD or USB boot media to boot my ESXi hosts, what are immediate actions that I can take to prevent such issues?

- SD cards or USB drives should only be considered to store the system boot and boot bank partitions. The ESX-OSData partition must be created on persistent storage.
- In any case, if a SD or USB device is used for boot media (only SD/USB or SD/USB + local hdd/ssd disk), follow the below instructions to reduce the amount of IO sent to the SD or USB boot media:
  - Ensure the **`/scratch`** partition is configured on persistent storage like a local hdd/ssd or boot from SAN device. The ESXi 7.0 installer will not create the **`/scratch`** partition on a SD card or USB drive. It tries to find persistent storage and attempts to create **`/scratch`** on that.
  - If there is no supported persistent storage available, a 250 MB **`/tmp`** partition is created on a ramDisk. This may limit the performance of ESXi hosts if this partition runs out of space.
  - Do not create the /scratch partition on the SD or USB boot media in any case. Running the **`/scratch`** partition on the SD or USB boot media is not supported by VMware. Always configure /scratch on locally attached disk (HDD or SSD). If the local disk is not available, then configure it on a SAN.
  - It is good practice to connfigure ESXi Dump Collegator to offload core dumps.
  - Enable the toolsRAM disk option to offload VMtools install/upgrade IO requests on the RAM disk.

Degraded Mode

Why do I see my ESXi hosts are running in 'Degraded Mode'?

When ESXi hosts are booted from a USB or SD device, and are not equiped with an additional local HDD or Flash media (NVMe or SSD) to store the OSData partition, ESXi goes into degraded mode because the USB/SD device isn't used for logs or persistent storage. Partitioning for hosts that are upgraded to ESXi 7 from earlier versions differs significantly from partitioning for new installations of ESXi. The size of bootbank partitions is different and autoconfiguration might not configure a coredump partition on the boot disk due to size limitations.

When in degraded mode, a sysalert is shown: *"ALERT: No persistent storage available for system logs and data. ESX is operating with limited system storage space, logs and system data will be lost on reboot."*

What does Degraded Mode mean?

Degraded mode, in summary, is a state where logs and state might not be persistent (get lost when the host is rebooted), with a side effect that it can cause boot up to be slower.

With new features and capabilities like vSphere Lifecycle Manager, vSAN, and NSX-T, the way VMware approaches ESXi host installations have changed. More and more features and capabilities will rely on host local storage for additional kernel modules, but also for state information like PCIe mappings and SSL certificates.

Can I add local storage after I upgrade to ESXi 7?

Yes. It is good practice to equip ESXi hosts with recommended local storage media. When adding local storage after upgrading to ESXi 7, enable the advanced setting *autoPartition=TRUE* and reboot the ESXi hosts as described in KB article 77009 to mitigate degrade mode.

I can't retrofit local storage media, how do I mitigate?

If local storage is not present and the ESXi host is running in Degraded Mode, there's still ways to make sure logs are stored, by using external Syslog solutions, NetDump Collector or Core Dump Partition. See this article how to enable Syslog on ESXi hosts.

Auto Deploy

I’m using Auto Deploy, what do I do?

vSphere Autodeploy is fully supported. However, with the recommendation to use ESXi host local storage for new features and capabilities, it is recommended to (re)configure the Host Profiles with Auto Deploy to use **Stateful Installs**. Doing so does require local storage for the ESXi hosts.

vSphere Lifecycle Manager (vLCM) is a new feature that relies on local storage. When vLCM is used in an environment with Auto Deploy, customers should change to statefull installations, see this article for more details.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com