
Boot from RAID of an ASUS PRO WS X570-ACE

It should basically work

Onkobu Tanaake

2022-08-28T10:00:00

Mounting was easy. The previous case Antek NSK4100 had very little space left for the CPU cooler. I replaced it with a Nanoxia Deep Silence 5 Rev. B. I also exchanged RAM with ECC modules. A hardware raid controller – LSI MegaRAID – was added, too. To make the mainboard boot I also plugged in the smallest passively cooled graphics card I could find. Powering up and going through the BIOS options I was quite optimistic. All extensions as well as the MegaRAID controller could be configured. Also the mainboard's RAID controller showed up.

I focused on the mainboard RAID controller. The MegaRAID will take some time, not only to get configured but especially to prepare the RAID 6. So I had to prepare a RAID 1 for two SSDs. Each drive is connected to one SATA port. The controller correctly recognized both drives. I chose the RAID type and selected the drives. That was all.

Sometimes such a configuration happens in two parts separated by a restart. Therefore I restarted the computer but nothing changed. To save much time experimenting I opened the manual to read about the RAID controller configuration. But there was not much to learn. The manual mentioned a few easy steps. But nothing was written about formatting.

Next source was the Internet. The server runs Gentoo Linux so there are sufficient degrees of freedom but there are no drivers anymore for an AMD RAID controller. This was already written down in 2017 in many forums. I am familiar with this situation for many different types of hardware. Either it is supported right away take the first search result. Or it runs only with limitations or not at all. So I had to abandon the idea of a hardware RAID 1 and instead switched to a classic software variant.

I already know that this makes booting a bit more challenging. Doing this without UEFI is not an option. A signed kernel is very important. But luckily UEFI cannot boot from a software RAID especially not the linux variant. A hardware RAID is offered as a single drive by the BIOS. Instead a structure of two independant disks doesn't mean anything to UEFI. Thus a software RAID is something completely different:

- RAID member marker is written to the disk itself
- depending on the formatting strategy this is stored at the beginning or at the end of the drive
- boot partition must be FAT32 with UEFI boot directory
- Grub does support RAID root partition (through initramfs, see #1 [#it_linux])

Boot from RAID of an ASUS PRO WS X570-ACE

My first attempt was to partition both drives for complete usage. A single partition spans the entire drive. This has to be broken into logic volumes using LVM. This renders the UEFI system partition unusable. The ASUS board does not recognize it. With my experience in partitioning Linux systems I tried the following with no success:

- Software RAID with superblock at the beginning, next a single partition
- Software RAID with superblock at the end (format 1.0), followed by a single partition
- Like before but format 1.2

It is very likely a general issue with ASUS boards no matter the partition layout. A software RAID with a single partition containing logical volumes cannot be booted. The ESP will not be recognized. The only option left is to break the ESP out of the logical volumes. It has to be before the remaining partition(s). This leads to the final structure of two identical ESPs. One on each drive. The remaining space is a LVM structure containing logical volumes forming the RAID array. Now the expert concludes: booting from the hardware RAID failed. All others keep in mind *the boot partition must be synchronized with a script*. If this is not done regularly the loss of one disk risks being able to boot the entire system.

Finally `initramfs` must be packed with modules and helpers for LVM, software RAID (`mdadm`) and LUKS. Kernel has to be parameterized with `do_lvm domdadm` and matching parameters `keymap=de cryptroot=UUID=...`. The `keymap` is essential to correctly type the password with special characters. I leave the `framebuffer/ VGA` module untouched. Since it is a server it can happily live forever taking over UEFI's `framebuffer`.

- Linux is an Alternative [<https://onkobutanaake.de/it/linux.html>]