

Q: Which is best, hardware RAID, vs Software RAID, vs ZFS ?

A: Short answer: If you can use ZFS, use that, its among the best file formats around. *But, no matter what you decide to use, remember to backup your data!*

Still have questions? Read on:

If you have been a system administrator for any reasonable length of time you've had a server with a hardware RAID controller report a failed disk. When replacing the failed disk with a new one, another drive fails during the rebuild, the system crashes and you have no choice but to restore your data from backup. This will happen due to the controller being bad or more likely the silent corruption of data on one of the other disks in the array. Silent corruption of data, colloquially referred to as **"bit rot"**, is more frequent these days as hard drives store larger amounts of data. I have personally seen this happen on several occasions over the last 20 years, most frequently within the last 5, so it is natural to wonder if the pros of hardware RAID really outweigh the cons for a general purpose server or NAS device.

What follows is not meant to be the last word on this subject nor a scholarly article. My goal is merely to help XigmaNAS users make a fast, informed decision by collecting the most important information from the last few years in one place. When you are done reading everything here you should know exactly why you have chosen one solution over the others. Different people have different needs and different budgets, there is no single solution which is right for everyone, but it is possible to arrive at some general recommendations that will be best for most people. *Those recommendations are the short answer above.* If you are looking for "How To" type answers or have specific questions about configuring or managing SoftRAID or ZFS arrays you should look elsewhere in the Wiki.

Pros of Hardware RAID

- Improves server up-time statistics and availability which might be negatively impacted by common, predictable disk failures when configured and monitored properly.
- Easy to set up - Some controllers have a menu driven wizard to guide you through building your array or automatically set up right out of the box.
- Easy to replace disks - If a disk fails just pull it out and replace it with a new one assuming all other hardware is hot-plug compatible.
- Performance improvements (sometimes) - If you are using an underpowered CPU, a HW RAID card with a decent cache and BBU (write backs enabled of course) will generate significantly better throughput. However this will only be a benefit if you are running tremendously intense workloads (transferring data faster than 1Gbps or hosting apps doing constant reads/write), the typical NAS4Free server is not used in such a way, though it could be.
- Some controllers have battery backups that help reduce the possibility of data damage/loss during a power outage.

Cons of Hardware RAID

- Proprietary - Minimal or complete lack of detailed hardware and software specifications/instructions. You are stuck with this vendor.
- On-Disk meta data can make it near impossible to recover data without a compatible RAID card - If your controller dies you'll have to find a compatible model to replace it with, your disks won't be useful without the controller. This is especially bad if working with a discontinued model that has failed after years of operation.
- Monitoring applications are all over the road - Depending on the vendor and model the ability and interface to monitor the health and performance of your array varies greatly. Often you are tied to specific piece of software that the vendor provides.

- Additional cost – Hardware RAID cards cost more than standard disk controllers. High end models can be very expensive.
- Lack of standardization between controllers (configuration, management, etc)- The keystrokes, nomenclature and software that you use to manage and monitor one card likely won't work on another.
- Inflexible – Your ability to create, mix and match different types of arrays and disks and perform other maintenance tasks varies tremendously with each card.
- **Hardware RAID is not backup**, if you **disagree** then you need to **read the definition again**. Even the most expensive Hardware RAID is so much junk without a complete and secure backup of your data which would be necessary following a catastrophic, unpredictable failure.

Pros of Software RAID

- ZFS and SoftRAID can improve server up-time statistics and availability which might be negatively impacted by common, predictable disk failures when configured and monitored properly.
- ZFS and similar (like BtrFS) were deliberately developed as software systems and for the most part do not benefit from dedicated hardware controllers.
- Development and innovation in ZFS, BtrvFS and others is active and ongoing.
- Hardware independent – RAID is implemented in the OS which keeps things consistent regardless of the hardware manufacturer.
- Easy recovery in case of MB or controller failure - With software RAID you can swap the drives to a different server and read the data. There is no vendor lock-in with a software RAID solution.
- Standardized RAID configuration (for each OS) – The management toolkit is OS specific rather than specific to each individual type of RAID card.
- Standardized monitoring (for each OS) – Since the toolkit is consistent you can expect to monitor the health of each of your servers using the same methods.
- Removing and replacing a failed disk is typically very simple and you may not even have to reboot the server if all your hardware is hot-plug compatible.
- Good performance – CPUs just keep getting faster and faster. Unless you are performing tremendous amounts of IO the extra cost just doesn't seem worthwhile. Even then, it takes milliseconds to read/write data, but only nanoseconds to compute parity, so until drives can read/write much faster, the overhead is minimal for a modern processor.
- Very flexible – Software RAID allows you to reconfigure your arrays in ways that are generally not possible with hardware controllers. Of course, not all configurations are useful, some can be downright dumb.

Cons of Software RAID

- Need to learn the software RAID tool set for each OS. – These tools are often well documented but often more difficult to learn/understand as their hardware counterparts. Flexibility does have it's price.
- Possibly slower performance than dedicated hardware – A high end dedicated RAID card might match or outperform software RAID if you are willing to spend over \$500.00 for it.
- Development of SoftRAID0 through 6 has pretty much stopped as focus moves to newer technologies like ZFS and BtrvFS.
- Additional load on CPU – RAID operations have to be calculated somewhere and in software this will run on your CPU instead of dedicated hardware. While I have yet to see real-world performance degradation it is possible some servers/applications may be more or less affected. As you add more drives/arrays you should keep in mind the minimum processor/RAM needs for the server as a whole. You need to test your particular use case to ensure everything works well.

- Additional RAM requirements - particularly if using ZFS and some of its high end features, but not for plain old SoftRAID.
- Susceptible to power outages - A UPS is strongly recommended when using SoftRAID as data loss or damage could occur if power is lost. ZFS is less susceptible, but I would still recommend a UPS.
- Disk replacement sometimes requires prep work - You typically should tell the software RAID system to stop using a disk before you yank it out of the system. I've seen systems panic when a disk was physically removed before being logically removed. This is of course minimized if all your hardware is hot-plug compatible.
- **ZFS and Software RAID are not backup**, if you **disagree** then you need to **read the definition again**. Even the largest and cheapest SoftRAID is so much junk without a complete and secure backup of your data which would be necessary following a catastrophic, unpredictable failure.

It will be pretty obvious to most people after weighing the pros and cons above that SoftRAID is the better choice for most XigmaNAS users. This is true based on price/performance even before we consider stability and longevity of your data.

Now that we've established Software RAID is the way to go, what do you choose? SoftRAID1, SoftRAID5, 6? Maybe a hybrid SoftRAID 10 or 50? What about ZFS?

As drive capacities increase so does silent data corruption or "**bit rot**". You must understand what this is to properly assess the odds of a catastrophic loss of data happening to you. Silent data corruption is not a new concept, it has been well discussed for years that RAID5 is pretty useless and RAID6 will soon be in [Why RAID 5 stops working in 2009](#), also [archived here](#). Read the article, it's about 2 pages and easy to understand even though you have to use powers of 10 to do the math. The take away is don't waste your money or your time on RAID5 or RAID6 if you really care about your data. Personally I have never trusted my data to either.

Since you now understand data integrity can no longer be improved with RAID5 or 6 you are left with just a few choices.

1. Hybrid SoftRAID - Not really a good option due to cost and management issues for most people, but still valuable in certain specific scenarios.
2. SoftRAID1 - I have always used SoftRAID1 with a good backup and it has always been cheap and easy to manage.
3. ZFS - Has a big advantage over plain old SoftRAID, data integrity is built in from the start while SoftRAID has none. Data integrity is the big advantage of ZFS and to learn more about it you should read [A Conversation with Jeff Bonwick and Bill Moore - The future of file systems](#) also [archived here](#).

Now that we are starting to use 6TB (or even more) drives on a regular basis, bit rot becomes an even bigger issue.

How big? Well, the science is not yet fully settled because testing this stuff is so boring. You should read [Keeping Bits Safe: How Hard Can It Be?](#) by David S.H. Rosenthal, published in the monthly magazine of the Association for Computing Machinery (ACM). I will summarize best as I can: A NetApp study encompassing data collected over 41 months from NetApp's filers in the field, covering more than 1.5 million drives found more than 400,000 silent corruption incidents. A CERN study writing and reading over 88,000TB of data **within a 6 month period** found .00017TB was silently corrupted. On average, it is not unreasonable to expect roughly 3 corrupt files on a 1TB drive according to another CERN study; [CERN's data corruption research](#) also [archived here](#). Yes, the numbers appear small, but these studies were conducted in 2007/08, before we had 3TB hard drives. The numbers have not

gotten better, they get worse with increased capacities and longer time frames.

Stop and think about it, when you had 400GB of data spread among a bunch of 500GB drives, how likely was it you would lose everything? Now that you have 2.5TB of data spread among a bunch of 3TB drives, how likely is it you could lose everything? Knowing that **MTBFs** and error rates **have not improved**, yet **capacities have increased 6X** or more, you've got to ask yourself one question, "Do I feel lucky?".

The bigger and older your drives are the more likely you are to suffer a failure or silent corruption. So save your time and money, data integrity is only one of the many benefits of ZFS, choose it from the start and avoid problems later. Don't forget you still need a backup even if you are using ZFS.

Basic	Basic storage concept	⇒ All versions
-------	-----------------------	----------------

From:
<https://www.xigmanas.com/wiki/> - **XigmaNAS WIKI**

Permanent link:
<https://www.xigmanas.com/wiki/doku.php?id=faq:0144>

Last update: **2018/08/10 15:42**

