

## The persistent importance of traditional storage

RAID is first and foremost – and will be for years to come

Joseph Martins Wed 14 Oct 2020 // 06:00 UTC

---

It's clear that how your organisation uses and exploits data is fundamental to its success if not its very survival.

The democratisation of analytics, machine learning and AI means that sophisticated data analysis is now in the hands of a broader range of companies and other organisations. Meanwhile, the cloud native revolution – whether on-premises or in the actual cloud – means large and small companies can develop and deploy software and features at a dizzying pace. This inevitably produces more data to analyse and, hopefully, monetize.

Some might question whether storage hardware has matched this pace of development – though it has clearly not prevented it. Nevertheless, the last decade has seen undeniable storage innovations at the micro level, with flash becoming more affordable. At the same time, material scientists have pushed the limits of physics with continued innovation in hard drive technology, meaning this decades-old format still has a role to play even in the rarified area of HPC storage. This has been matched at the macro level with new architectures, such as software-defined storage (SDS) and hyperconverged infrastructure (HCI), and the continuing evolution of storage management software, resulting in more robust and secure systems. One thing that hasn't changed though is that the traditional RAID system often remains the centrepiece of most enterprise storage designs.

RAID stands for redundant array of inexpensive disks, and both the concept and acronym have been around since the second half of the 1980s, when scientists at Berkeley posited the concept of replacing the very expensive single disk architectures used for mainframe storage with large numbers of the (relatively) cheaper drives used in PCs, delivering improved performance, and reliability.

Whereas RAID and enterprise class storage were once synonymous, there are now other options such as the aforementioned HCI and SDN. Each has similarities to classic RAID in terms of offering a virtualized storage pool, but the implementations, and potential applications, are different.

There are solid reasons why RAID maintains its place at the core of many organisations' storage infrastructure. One obvious reason is that many still rely, to a greater or lesser degree, on legacy applications, which were, and are, built with the presumption of RAID at the back end. This is not to imply that they – or the storage supporting them – is in anyway archaic, or that the only reason they are still being used is that no one knows what would happen should someone unplug them. A bank will hardly look at its OLTP system in this way. Such applications may also greatly improve performance using a state-of-the-art RAID system.

## **RAID performance**

Storage Performance underlines one of the reasons for RAID's continued preeminence. If you need to dedicate "absolute performance" for specific applications or environments, RAID will likely be your preferred choice, whether for a legacy environment or not.

Your organisation's own attitude to risk or innovation is a more nebulous consideration. For instance, a more conservative business might prefer to stick with the proven reliability of traditional RAID, rejecting more fluid architectures such as HCI or software defined networking. Customers may simply feel that RAID will not let them down – and they have over 30 years of experience to base that supposition on.

That said, when considering today's RAID systems, predictability and reliability should not be confused with slow and steady. Although they might appear monolithic, modern RAID systems encompass a range of technologies – though newer systems overwhelmingly use flash drives to a greater or lesser degree.

## **Flash – not quite master of the universe**

Not so long ago, flash for enterprise storage was considered rather exotic. And today flash drives remain considerably more expensive than traditional hard disk drives.

So, the entry level RAID system these days is likely to be a hybrid storage setup, at least when maximum storage consolidation is the objective. These systems will use a tier of SSDs together with HDDs to deliver a balance of cost, capacity and performance, depending on the application. The hard drives, for example, offer a potential advantage for sequential workloads, making them a sensible choice for applications such as media streaming. But such systems also have the issue of tiering to contend with. Ultimately, the choice comes down to budget.

All flash arrays can now squarely be considered mainstream. If a tier of flash delivers a performance boost in a hybrid system, an all flash array delivers the full range of benefits of SSDs. First up is a clear performance advantage, with such systems able to offer sub-millisecond latency, and up to 50 times more IOPs than hard drives. So, you can expect a boost across all operations, including backup, recover and snapshotting.

But the nature of flash drives means other advantages. With no internal parts – the read/write head, the spindle, the media – there is less to fail, and the consequences are less potentially catastrophic. Simultaneously, SSDs themselves are smaller and require up to 95 per cent less power and up to 95 per cent less space than their spinning counterparts. That means less space and less need for cooling within the data centre.

As well as simple real estate benefits, this means less of an admin overhead, which might be worth considering when doing the overall budget equation. An all flash system should require 40 per cent less operational effort.

So, it's no surprise that flash-based arrays now account for around 80 per cent of the array market by value, according to IDC, evenly split between all flash and hybrid versions.

## NVMe

The all-flash adoption balance is split between arrays that use the same SAS/SATA protocols as classic RAID, and arrays that use the NVMe storage networking protocol. The latter – for Non Volatile Memory Express – is a relatively new technology, and is currently very expensive. NVMe exploits PCIe, giving the CPU a mainline connection to the data it craves. This affords a drastic increase in parallelism. While SAS and SATA SSDs support a single queue with 32 or 256 commands respectively, NVMe supports 64,000 queues, each of which can support up to 64,000 commands. The effect on storage bottlenecks is immediately obvious.

Cost apart, whether you want to plump for NVMe will depend on your applications. That massively increased parallelism might be extremely alluring if your RAID system is to support OLTP, databases and datamining or real time analytics. Those workloads will require constant reading and writing of data, so NVMe will really come into its own there. For something like CRM, the benefits are less clear cut. It should be remembered that the management software layer is a crucial part of the mix. Compression and deduplication software can have a big impact on the amount of raw storage actually required. Likewise, the level of data protection your chosen vendor offers will have a dramatic impact on issues like replication, mirroring and failover. Together with reliable hardware, this will all influence whether your system offers you minimal disruption when upgrading, and zero downtime in operation.

And neither RAID – hybrid, all flash, or NVMe-enabled – will exist in isolation. You might be installing or upgrading a RAID system but have an eye on the cloud because that's where your applications increasingly live, or maybe that's where your data is backed up or replicated. And even if that's not the case today, it could well be in the not too distant future.

This raises the broader point of choosing a partner for all your storage needs, because you need to be mindful of their partners too. For many enterprises, RAID is the workhorse that keeps things going, but increasingly, it will co-exist alongside other use cases. SDN or HCI might not make sense for your general storage workloads, but could absolutely fit for specific applications or environments. Or you might be following a hybrid cloud mode and want a partner that can actively supply systems optimized for such environments.

Whatever your ultimate choice, a few things are clear. With the speed and performance brought by flash, RAID increasingly doesn't actually rely on disks alone. Price performance continuously improves, but no one would pretend it is inexpensive. But over 30 years after its birth, RAID storage certainly isn't redundant, and is set to remain a constant in the storage lifecycle for the foreseeable future. It is still the foremost choice when considering a data storage purchase, and this is set to remain a constant in the storage lifecycle for the foreseeable future.

*Sponsored by Fujitsu*

Published on "The Register", October 14, 2020

[https://www.theregister.com/2020/10/14/raid\\_is\\_first\\_and\\_foremost/](https://www.theregister.com/2020/10/14/raid_is_first_and_foremost/)

If you want to know more about different storage solutions, check out the other articles published on "The Register" or explore the storage data jungle: [www.fujitsu.com/data-jungle/storage/](http://www.fujitsu.com/data-jungle/storage/)

The key role of storage in building a datacentric organisation

<https://docs.ts.fujitsu.com/dl.aspx?id=57cc1295-ba0b-46dc-8882-40db4fe8d188>

SDS or HCI? How to figure the right fit for the right workloads

<https://docs.ts.fujitsu.com/dl.aspx?id=981e167e-2f7e-4cd9-baa6-2f73a47ec2f9>

How backup fits into your data-protection strategy

<https://docs.ts.fujitsu.com/dl.aspx?id=c078abcf-2c33-4f4c-9f22-2e9cd1ed48da>

Your archive contains a treasure trove of information. Use it

<https://docs.ts.fujitsu.com/dl.aspx?id=6ed0abd2-3d22-43d6-b073-3386dd2a0077>