

PAYING TWICE TO STORE DATA ON DISKS



To store large amounts of data, we need to exceed the capacity of any one hard disk drive. The only way to do that is to have the data span multiple drives, but nevertheless look and behave as if it was one, really big, imaginary drive.

Computer data storage is an increasingly lucrative business. The amount of data that needs to be kept — reliably, persistently, and with easy access — is growing constantly. Vendors who sell hard disk drives and the related “storage solutions” are keeping pace with this demand. All at a price, of course!

There are four factors that must be considered when anyone is thinking about how they are going to store computer data on disk drives. These are size, reliability, performance, and cost. It is possible to store your data on one hard drive, and that’s the end of that. The amount of data that you can store will be limited to the size of the physical disk, you will be able to read and write your data until the physical disk fails, the performance will be slow but acceptable for a casual user, and the cost will be very low. That might be very well you’re at home with your own laptop or desktop computer, but it is inadequate for an enterprise user that has substantially greater demands.

To store large amounts of data, we need to exceed the capacity of any one hard disk drive. The only way to do that is to have the data span multiple drives, but nevertheless look and behave as if it was one, really big, imaginary drive. Virtual volumes, sometimes called metadisks, can take care of spanning. They can also take care of the second concern, which is reliability. The virtual

volume can physically store data in two places at once, but present it as only one object. That's called mirroring, and there are other strategies that can enhance reliability. We want to avoid the "single point of failure." Performance can be addressed in the virtual volume by stretching the data across the multiple physical drives in such a way that the disks can work in parallel to satisfy the demands for reads and writes that are coming from the application.

"Many hands make light work," is the idea here. Striping is the main strategy to achieve this improved performance. Finally, there is cost. It is going to cost a lot more money to have multiple physical disk drives store large amounts of data, with reasonable reliability and with good performance.

No matter what, if you want virtual volumes you are going to pay more money to buy more physical disks. But what about the storage solution of configuring and managing virtual volumes? Do we have to pay for that too? Major IT vendors say yes.

There are two approaches to managing virtual volumes. Do it on your computer, and send the reads and writes of data to "Just a Bunch Of Disks" (the JBOD). This is usually called host-based RAID, where RAID is an older term that stands for "Redundant Array of Independent Disks." The other approach is to do nothing on your computer, but send reads and writes of data to the disk array that has a hardware processing component built in to deal with the virtual volume effects. This is usually called controller-based RAID or array-based RAID.

There are vendors selling licenses for software to support a host-based RAID solution. Other vendors sell licenses plus disk arrays to support a controller-based RAID solution. Both types of storage companies are aggressively targeting enterprise customers like banks, telcos and government departments. The surprising thing is that many of these customers opt for and pay for both. The result is quite costly, is of no greater benefit to them, and they end up not fully exploiting the potential of either approach in dealing with their data storage requirements. For example, a controller-based disk array already consists of collections of disks presented as single "logical units."

But then the customer turns around and uses a virtual volume management layer of software to concatenate those logical units into an even larger virtual volume. It all could have been done at the hardware level using the smart disk array, or it all could have been done at the software level using virtual volume management software taking care of Just a Bunch of Disks. Why do people do something that is inefficient and costly to their business? The only explanation I can come up with is simple inertia. Their company has always bought licenses from a host-based virtual volume management provider, and they unreflectively renew them each time they expire. Their company has always bought arrays of disks, and then the controller and its related software and licenses just comes along with that. Money goes out the door to two supporting IT vendors to achieve the same result, and no one bothers to ask why.

A reflective approach would lead to the knowledge that not only would it be better to definitively choose one or the other vendor, but indeed it might be better to go with a modern solution that is embedded right in the operating system, and is free. ZFS has been a fully-supported part of Solaris for years, and it has the ingenious feature of integrating virtual volume management and a standards-supporting file system, all in one package. Simply toss disks into the “pool” and the space on them is configured for use right away. Buy a simple array of disks — which you needed to do anyway — and then manage how you store data on them easily and without separately purchased licenses. ZFS is open-source software, licensed under the Common Development and Distribution License.

Too many people put the solution first (“should we renew our license?”) and don’t take the time to rationally analyze what it is they need to do. They need to store large amounts of data reliably and with good access, and they might have to pay more money to do that. They can buy Just a Bunch of Disks and perhaps licenses from a host-based virtual volume management provider, and do it that way. They can buy a controller-based disk array, and related licenses. Or they can use ZFS that’s ready for use as an integral part of Solaris. In any event, it doesn’t make any sense to pay twice to store data on disks.

Source : <https://www.exitcertified.com/blog/michael/2012/paying-twice-to-store-data-on-disks/>