HIGH-PORT WHITEPAPER SERIES

#1: How Many Disk Drives Do I Need?

Large Drive Counts and High Port RAID Controllers

Today, you are faced with a large number of problems that have to be solved when creating storage platforms. In the past, you selected a CPU performance point, added some memory and a SCSI controller to connect to your disk drives and you were good to go. About the only choice on the I/O side was whether to go with single- or dual- channel SCSI.

But the world has changed. Motherboards have multiple CPU sockets and CPUs have multiple cores. Memory capacities are increasing quickly, as are the demands of your customers' operating systems. I/O got more complicated with the advent of blindingly fast (but expensive) Serial Attached SCSI (SAS) drives and low cost, high-capacity (slower) Serial ATA (SATA) drives. And then, to top it all off, virtualization entered the picture, so you might not even know whether your system will be host to one operating system or ten.

Designing a system (or a range of systems) that can support these diverse needs is difficult. But when it comes to the I/O side, a Unified Serial™ (SATA/SAS) controller can help you handle high speeds and high port counts.

RAID controllers built with Unified Serial Architecture can connect to both SATA hard drives and SAS hard drives. Today, you should probably not even consider purchasing SATA-only controllers anymore. Unified Serial controllers provide greater flexibility and much higher performance at the same or lower cost, so even if you plan to only use SATA drives, an Adaptec Unified Serial is a better choice. It also lets you cover the possibility that a customer will suddenly ask for SAS drives for some particular application, without having to re-qualify your system to accommodate them.

Having made that I/O decision there are really three more questions that you ought to consider:

- How many drives should I use?
- What type of drive should I use?
- How should I do the "plumbing" to connect up the drives?

We'll discuss each in turn in each whitepaper of this series.

How many drives should I use?

The simplest way to figure out the number of drives that you need is to talk to your favorite drive vendor and identify their highest capacity drive. Then, take the space that you think your customers need, and divide one into the other to get the number of drives.

For example, with 1.5TB SATA drives available today, you can build a 3TB server with two disk drives. Of course, you'll want RAID protection, and RAID 5 will require an extra disk, so it looks as though three drives will do the job.

But, thinking this way is likely to work only in the simplest of applications with the lowest requirements. Large hard drives are great… they give you lots of space. However, they don't help performance in the slightest. And, when you factor in RAID data protection, which is a must have for most server customers, performance tends to relate directly to the number of drive spindles you have in your system. Most often, the truth is that the more drives you have working, the quicker things get done.

So, to get both capacity AND real performance from a server, use more disks with smaller capacity. From this you get two huge benefits:

- Better performance for a wider class of applications
- Better chance of hitting the "sweet spot" in terms of price-performance with your drive vendor

Don't underestimate the latter point. The highest capacity drives are not usually the cheapest. The price-capacity attributes of different drives tend to move around a lot, but the most cost-effective drives are usually those down a few notches from the highest capacity. Combine this with the better performance that you can get from using more drives and you have a significant win.

So instead of using three 1.5TB drives to get your 3TB of application space, think about using seven 500GB drives in a RAID 5. You'll get the same capacity, but much greater performance, simply because the reads and writes are now spread across seven spindles rather than three. Both transactional performance and streaming performance will be more than twice as good. While past generations may not have required that much performance, the onset of virtualization is changing everything. With your customers running many different operating systems on top of your I/O system, it is more important than ever that you over-provision your I/O.

You'll probably also pay less for the drives. [1]

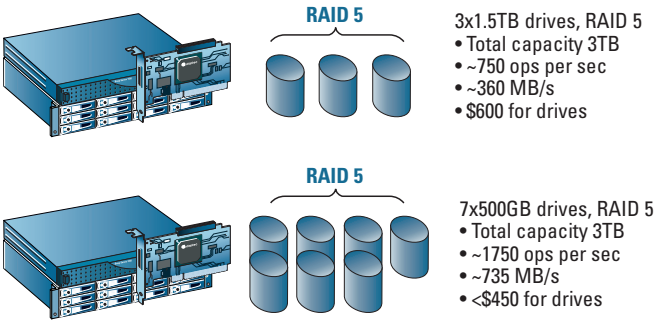

Of course, you could take this even further. Some applications today are looking for very large disk spaces. Video surveillance systems, organizations specializing in processing rich media, disk-

1. Google search indicates that 1.5TB drives are selling for approximately $200 each, whereas 500GB drives are $60-$70 each.

to-disk backup solutions, and many more are looking for massive arrays that offer large capacity and also higher performance and greater redundancy.

Again you can go with ultra-large drives. Why not get eight 1.5TB drives and build a RAID 5? This should provide around 10TB capacity. Sounds great, but what are the consequences of doing this?

Putting a large number of drives in a RAID 5 is the most economical way of gaining storage, and it provides great performance when things are running well, but when a drive dies things can go bad quickly. When that drive dies (and it will) the RAID algorithm in this configuration kicks in to preserve your data, but every hard drive in the system is now involved in the rebuild process, as well as doing its day-to-day activities. Plus, with the process taking around 24 hours to complete, you have an issue on your hands.

So how can you work around this?

Adaptec Unified Serial controllers offer other RAID protection levels; RAID 50 and RAID 60 have major benefits over single-level arrays. These "multi-level" arrays are actually several RAID 5's combined into a RAID 0.

To see how this works, consider an example:

- Ten hard drives in total
- Two 5-drive RAID 5 arrays combined together

From the perspective of the attached server, this yields a single RAID-protected array with the capacity of eight drives, as shown in the following figure.

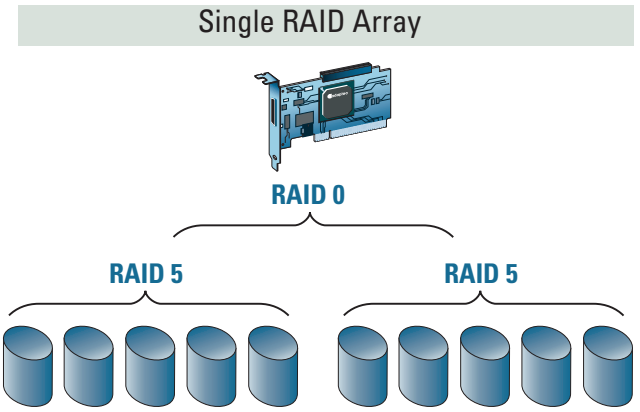

So why is this better than a single-level RAID array such as a RAID 5?

On the performance side, these configurations are even faster for both reads and writes. More drives are involved and the firmware in the Adaptec Unified Serial controller is able to take advantage of this by load balancing. This also provides reliability benefits. When a drive dies, only the disks involved in the particular RAID 5 containing the dead drive are involved in the rebuild. The other half of the array just continues to do normal day-to-day work. This means that performance degradation during the rebuild is almost unnoticeable, there is less stress on the overall array, and the rebuild finishes faster because there are fewer drives involved.

So, back to the question of how to build the I/O system in your servers; using a larger numbers of drives has many advantages over buying the highest capacities, including:

- Allows the use of more cost-effective drive pricing
- More drives adds both transactional and streaming performance
- Less impact when a drive fails using one of the higher level RAID protection schemes
- Higher reliability

To exploit this, you need to be able to attach more than the standard four or eight disk drives to your system. The best and simplest way to do that is with a hardware RAID controller that supports a higher number of drives, such as the Adaptec 16-, 20- and 28-port Unified Serial RAID controllers.

Conclusion

Designing the best storage subsystem for a new server platform isn't as easy as it used to be. SCSI configurations tended to be pretty straightforward and the mathematics of performance, cost and reliability were well-known.

The advent of SATA and SAS drives with widely differing performance and prices makes the problem somewhat more complex. Plus, virtualization brings a completely new problem with its own requirements for I/O.

More than ever, it will pay for general purpose or high-performance servers to support a large number of drives of different types.

Adaptec has designed its Series 5 Unified Serial RAID family with these issues in mind and deliberately supports a wide range of port-count configurations to help you implement these systems. Combined with best-of-breed performance, industry-leading compatibility testing and Intelligent Power Management capability, these RAID controllers are the best fit for all your designs.

The next whitepaper, "*Which Drive Types Do I Need – SATA or SAS?*" in the High-port RAID Series will discuss how to select the type of drive that best meets customer needs, either SATA or SAS.

adaptec

Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, California 95035
Tel: (408) 945-8600
Fax: (408) 262-2533

World Wide Web: www.adaptec.com
Pre-Sales Support: US and Canada: 1 (800) 442-7274 or (408) 957-7274
Pre-Sales Support: Europe: Tel: (44) 1276-854-500 or Fax: (44) 1276-854-505

Part Number: 667108-011 Printed in U.S.A. 10/08 6303_1.1