GREEN POWER

Adaptec Intelligent Power Management

Reducing IT costs and storage power consumption by up to 70%

Introduction

Ten years ago, IT organisations purchased equipment with an eye to maximising the amount of computing power that could be obtained for a particular price. Maximising the GHz-per-pound ratio was the name of the game. Faster processors meant that more could be accomplished.

As time passed, the cost to manage a piece of IT equipment also became important, as companies struggling to control expenditures realised that over the lifetime of a system, management costs could substantially exceed its original purchase price. Now, system manageability became as important to choosing a system as the purchase price. Ease-of-use became an essential attribute, and automation came to be perceived as a way of scaling management and operational costs.

After absorbing the cost reductions achieved by streamlining management, companies are still faced with trying to cut the fixed operating costs while maintaining an effective amount of computing power. In the past few years, this focus has settled on power management, the balancing act between the power demands of IT equipment and the availability and cost of the necessary power. Today, every company faces these two factors:

1. IT equipment requires more and more power to operate and cool.
2. The amount of power that can be provided into data centres is limited by the power distribution grid.

Power needs can double the cost of a system

The power and cooling requirements of disk drives are a primary cost in operating a server. In most servers, full power is maintained to every drive at all times, even when it is inactive. Naturally, this constant full power state also increases the costs of cooling the system.

At current rates, the cost of operating a server over a typical four-year life is about the same as its initial purchase price. Estimates vary, but it is likely that more than US $10B will be spent worldwide in 2008 operating servers and their associated cooling.[1] Of this, as much as US $2B to $4B will be spent to provide the energy needed to rotate and cool disk drives inside servers.

As a result of these figures, many companies have started to look at different metrics when purchasing IT equipment. Instead of concentrating on brute computational power, the ratio of computation power per Watt of electrical power is becoming an important number. Particularly in larger data centres where the supply of electricity is limited, organisations are more and more focused on power optimisation instead of on raw processing power. However, even in smaller organisations, optimising power usage can lead to substantial cost savings.

A new generation of RAID controllers can provide power savings

But how can these power savings be realised?

In theory, you could focus on each subsystem individually, figuring out how to use a slower CPU, or fewer disk drives, or changing algorithms or applications to cut down on requirements, but this would be prohibitively difficult to manage and would require constant care and attention – i.e., expense. Industry-wide, there are some emerging themes that can help, such as server consolidation and virtualisation, but little has been available at the level of general-purpose storage.

Plus, storage is a complex puzzle involving multiple disk drive manufacturers, multiple models of disk drives and disk interface protocols, different external storage enclosures, different backplanes, etc. Fortunately, there is a way to address this problem without having to interact with each vendor independently because there's a common component connecting all your applications and servers to their storage – the HBA or RAID controller that you're using.

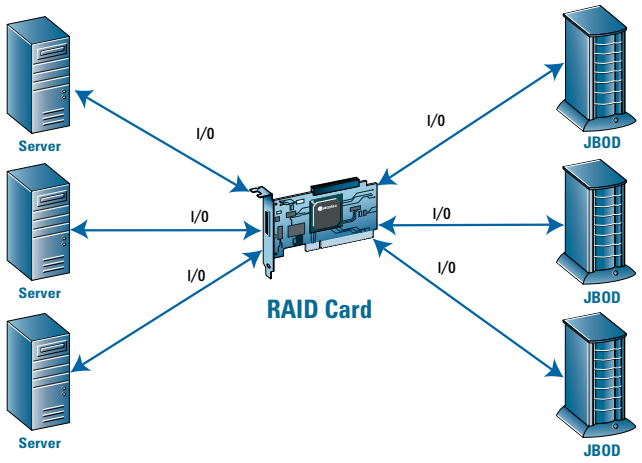

Figure 1. All data in a virtualised network storage environment passes through the RAID controller or HBA.

The controller occupies a uniquely valuable position in this regard, because it can monitor all the I/O to and from your storage, detecting usage patterns. It also has access to the underlying disk drives and enclosures and can issue commands to them that can affect their power consumption – e.g., to spin a disk more slowly, or even to spin down a drive. As a result, with sufficiently intelligent software on the controller, you have a single point where you can set policy for your entire storage infrastructure.

1. Extrapolation based on "Estimating Total Power Consumption by Server in the US and the World", Jonathan G. Koomey, Ph.D., Lawrence Berkeley National Laboratory and Consulting Professor, Stanford University, Feb 2007

Potential savings depend on usage patterns

To see how this might help, consider a typical SATA II or SAS disk. When running at full speed, these drives can consume between 8W and 15W of electrical power. Studies show that the power required to cool these components is roughly the same, for a total power draw of 16W-30W per disk drive. When spun down, such drives typically use a total of 3W-5W – a relative saving of 75%-85% depending on the drive.

To see whether you can benefit from this kind of saving, it is important to look at usage patterns.

Consider a typical fileserver in a small or medium business that operates on a five-day working week. This device is probably used for much of the working day, and is mostly idle during nights and weekends. As a result, the system is in use for around 40-50 hours of the 168-hour week. If you could spin down the drives when the system wasn't active, you could save power during the 130 hours that the system isn't in use – 77% of the time.

Other systems may be in use even less frequently. A server used for a nightly disk-to-disk backup might only need its disks spinning for 2-3 hours a night. The same applies to an archiving system, a print server or fax server. Even general-purpose servers running applications can have usage patterns with significant idle time. An accounting system, for example, might only be in active use during business hours, and a high-end transactional application (such as ERP) can have significant periods of inactivity.

This issue isn't only restricted to general-purpose computing systems. Speciality equipment in other industries such as medical, industrial manufacturing or video processing is likely to have periods when their I/O subsystems are basically idle.

However, one key point about these different applications is that they *are* different. Each one will have its own pattern of usage, which goes back to the role of the RAID controller. While you could try to configure each system individually to maximise your power settings, this would take a lot of effort and would likely lead to a sub-optimal solution. Plus, you'd have to reconfigure things as your usage patterns evolved to meet different business needs.

RAID controller power policies can provide substantial savings

Adaptec has automated the ability to manage power based on usage patterns. Adaptec Intelligent Power Management allows you to choose a low-power mode that spins active disks at a lower RPM and completely spins down idle disks. It is designed to reduce power costs by as much as 70% without compromising performance.

Adaptec Intelligent Power Management allows users to choose from three levels of disk drive power modes:

1. Normal operation: full power, full RPM (revolutions per minute)
2. Standby: low-power mode; spins disks at lower RPM when not in use
3. Power-off: disks not spinning; reduces power and cooling energy consumption by as much as 70% with unmatched performance

Adaptec Series 5 and Series 2 RAID controllers with Intelligent Power Management allow you to set a simple policy that is based on the actual I/O patterns being observed in real-time by the device. Indicate a time period, and if the controller performs no disk I/O for that period, it spins down the drives. Because the controller is in a position to actually see all the I/O to and from the devices, it can perform this task intelligently across all your disk drives regardless of which applications are accessing them.

For example, if the drives for a particular volume have been idle for the timeout period, and are spun down, you are saving power. If a subsequent I/O touches one or two of the disk drives that make up the volume, then only those few drives are spun back up. As a result your power savings can continue even during small amounts of I/O.

As an example, disk-to-disk backup can have up to 80% idle time where the disks can be spun down. With Intelligent Power Management, this can lead to £67 per system yearly savings with a 16-drive configuration.

Number of drives per controller	4	8	16	24	48	96
Savings 1 system	£29	£57	£114	£172	£343	£686
Savings 50 systems	£1,430	£2,859	£5,719	£8,578	£17,156	£34,311
Savings 250 systems	£7,148	£14,296	£28,593	£42,889	£85,778	£171,556
Savings 500 systems	£14,296	£28,593	£57,185	£85,778	£171,556	£343,112

Figure 2. Power and cooling savings in a disk-to-disk backup application using Intelligent Power Management.

Battery-backed cache

The Adaptec Series 5 controllers take this one step further by offering a battery-backed cache on the controller itself that can be used to buffer some of the I/O. Using these controllers, you can delay spinning the drives back up for an extended period by keeping data in the battery-backed cache. This is particularly important when you are using an operating system or tools that make periodic low-frequency access to storage for simple bookkeeping purposes – for example, updating a single timestamp or Windows registry entry. In those environments, the Series 5 controller can maximise your power savings even on your operating system's boot volume without sacrificing data integrity or reliability.

Standby – reduced disk spin

The trade-off involved in this process is that when you do need to access the disk drive, it must be spun back up to full speed before data access can occur. The time taken for the transition from fully spun-down to full operation depends on the particular disk drive, but can take as long as 30 seconds. Obviously, the storage controller hides this from you as best it can, although you may see some delay when a system wakes up from its spun-down state. To alleviate this issue, some drives – particularly newer SATA II models – offer an intermediate option in which the drive can be spun at a reduced

speed. While this state doesn't provide the full power savings of spinning down the disk, you can still achieve 30-40% of the possible savings. Furthermore, it takes only 5-10 seconds to re-enter the full speed state and perform I/O. Most SATA drives currently support this capability, and SAS drives will begin to support it by the end of 2008.

Intelligent Power Management supports this feature by offering you the option to enter this intermediate state if it is available on your disk drives – Standby mode. As a result, you can have the best of both worlds – full power savings on systems that perform little or no I/O, and intermediate power savings on those systems that perform at slightly higher I/O levels.

Override feature

There may be times when you will want to override the intelligence built into the RAID controller. For example, if you know that your systems will be needed between 8:00 pm and midnight every evening to perform some task, you may want to override this power-saving behaviour. Adaptec allows you to set an override period during which the power management monitoring and behaviour is temporarily disabled.

Another issue that concerns some users is whether or not drives will successfully spin up when I/O starts. For example, some systems may only be used once a month for data mining operations, or to help close out accounts. Such a system might have its drives spun down for 29 or 30 days a month. To help alleviate this concern, Adaptec controllers support an option to periodically spin back up and verify drive integrity. If the drive has been spun down for longer than the indicated period, it can be briefly re-activated and tested, and then spun down again. This technique can be helpful if you actually have a disk failure because it will be detected during the periodic wake-up instead of only being seen when you need the system for its real purpose. Detecting the problem early allows you to repair it before it becomes a critical issue.

Alternative power-saving techniques

Obviously the techniques described here are primarily intended to address power savings by spinning down drives. While this is perhaps the most obvious way to save money, there are alternative approaches that you might need to consider.

One of the most obvious is "consolidation".

Consolidation

For a variety of reasons, many IT organisations have grown up with a lot of different servers providing lots of different functions. It is quite common, for example, to have each server performing only one or two tasks. This is often convenient from a management perspective because, with a one-to-one relationship between applications and hardware, it is easy to understand and coordinate upgrades, downtime, etc. Each application and its hardware is an isolated "processing island".

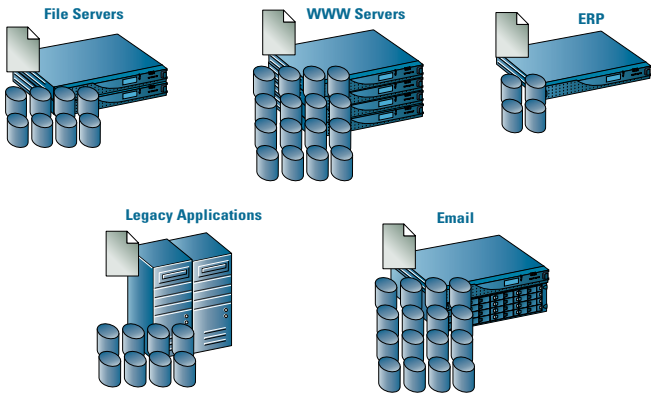

From a power management perspective, however, this is not an ideal solution. Each server has its own power supply, its own fans and its own disks, etc.

A more power-efficient approach would be to consolidate multiple applications onto a single, "bigger" server as shown here:

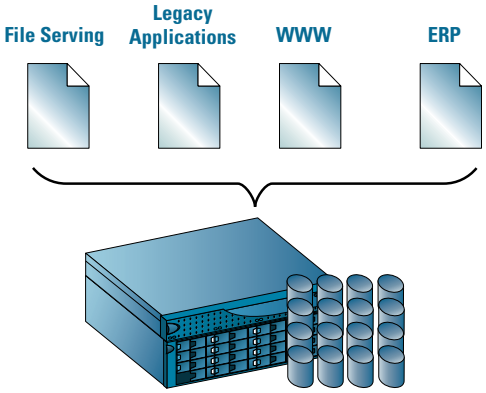

Obviously, you need more capacity in the bigger server, but this is readily available in current designs, which are likely to offer multiple, multi-core processors and large memory capacities that are easily capable of running many applications.

Storage can be a trickier issue, but Adaptec offers a complete family of RAID controllers that handle anywhere from 4 to 24 direct-connect disk drives with the same software features and scalable performance, so you can move from, for example, a 1U rackmount server with 4 drives to a 5U server with 24 drives without any risk. And, while the larger system obviously requires more power than any single smaller system, the total system power savings can be significant because many components are shared between multiple applications. At the storage level, Adaptec Intelligent Power Management applies to each volume independently so you can still derive the benefit of disk drive power management even if you have multiple applications sharing a single server. Some of the drives may be running at full speed because their applications are active, while others are spinning slowly, or are spun down because their applications are idle.

Obviously, the performance of your storage subsystem could become an issue in this kind of configuration because you have multiple applications accessing drives through a single storage controller. But, with the highest performance in their respective classes, the Adaptec Series 5 and Series 2 are designed to minimise the impact of complex demands on system performance.

Availability could also be a concern since this kind of configuration has multiple applications dependent on a single hardware platform. At the storage subsystem level, you can use features such as Adaptec RAID 5 and RAID 6 data protection with optional hot spares to protect against drive failures and minimise your exposure to data loss.

Virtualisation

While consolidating your operation to this kind of configuration can optimise your power savings, you might also want to take advantage of virtualisation to simplify overall management. Using virtualisation software such as VMware or Microsoft's Hyper-V adds an additional level of isolation between the applications running on the consolidated server by having them run in their own virtual machines.

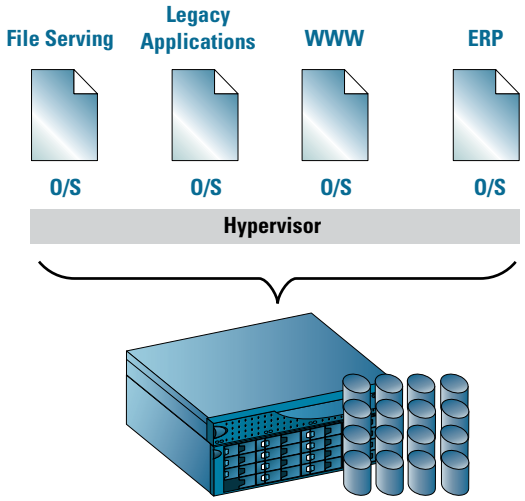

Optimising performance is even more important in this configuration since the virtualisation software itself requires significant resources. Adaptec Series 5 RAID controllers are an ideal solution since they offer not only the industry's highest I/O performance, but also offload all RAID processing from the host processor, allowing it to be used to run your applications instead of a software RAID stack. These Adaptec products have also been certified for use in virtualised environments.

Conclusion

The challenge of saving power while maintaining server performance has become a primary concern within the IT environment. But the latest generation of Adaptec RAID controllers provides Intelligent Power Management software and supports a variety of techniques that, together, enable significant savings without compromising performance.

Intelligent Power Management allows you to save up to 70% of your energy costs by matching power consumption to your true usage requirements.

Intelligent Power Management is a key capability offered on all Adaptec Series 5 and Series 2 RAID controllers. In addition to Intelligent Power Management, these controllers offer the industry's richest feature set, including:

- The industry's highest I/O performance
- RAID 6 for surviving two disk failures before data is lost
- RAID on Chip (ROC) architecture that maximises host processor performance by offloading RAID operations
- Adaptec Storage Manager™ software that reduces cost by centralising and automating management of all Adaptec RAID in the storage infrastructure

Quite simply, no one offers you more in a RAID controller than Adaptec.

Figure 3. Intelligent Power Management is available on all Adaptec Series 5 and Series 2 RAID controllers.

adaptec

Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, California 95035, USA
Tel: (+1) 408 945-8600
Fax: (+1) 408 262-2533

Adaptec – UK
Tel: +44 (0)1276 854 528
Web: www.adaptec.co.uk
Email: uk_sales@adaptec.com

Part Number: 667096-011UK Printed in Europe 08/08 6200_1.4