Taming the Power Hungry Data Center

Extraordinary power savings are achieved by integrating the world's highest performance storage.

FUSiON-iO

The Revolutionary Change

There have been a number of unique, evolutionary technology advances applied in today's data center to improve system application performance, scalability, and, more recently, to reduce total power consumption. Many such changes are examples of an evolutionary change driven by the disparity between the performance of the central processor unit (CPU) and the supporting storage subsystems. The result is a complex tangle of power-hungry hardware that fails to keep the processor occupied and, worse yet, adds more servers to make up the inefficiency. The ever-widening disparity between CPU performance and that of storage subsystems has spawned numerous, complex "solutions" that do not fundamentally fix the problem. This is the dark and dirty secret exacerbating the exponential growth in demand for power in the world's data centers. The goal of this white paper is to expose this fundamental flaw and introduce the means available today to substantively reduce power consumption while improving performance: Fusion-io's revolutionary flash memory-based storage technology.

The Fusion-io Advantage

Installing two ioDrives let one company disconnect a 100-hard-disk storage array, remove 3 of 4 CPUs, and reduce memory by half while increasing performance.

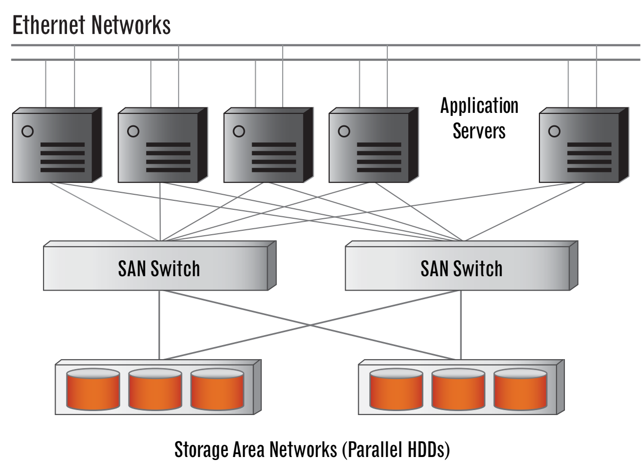

This idea may sound counterintuitive. How could underutilization of the CPU lead to increased power consumption? This white paper will address this question and demonstrate how Fusion-io's state-of-the-art, solid state storage solution, the ioDrive, stops the increasing power demands caused by the current data center path. The ioDrive's secret sauce is its ability to utilize and extract performance from common flash memory. Thumb drives, digital cameras, hand-held media players, and MP3 (music) players all contain and are viable because of flash memory. This type of storage memory is unique in that it can store data even when the device's power is switched off, requires very little power when in use, is very durable, because it has no moving parts, and is much faster than hard disk drive storage. The ioDrive incorporates all the advantages of flash memory, overcomes its limitations, and delivers mind-boggling storage performance and reliability. The ioDrive heralds a change in data center architectures that can provide orders-of-magnitude more computing capability at significantly lower power, thus dramatically increasing the data center's delivered performance versus the amount of power consumed. But more importantly, this revolutionary change not only significantly reduces power consumption, it drastically reduces other data center resources, including equipment, floor space, and temperature and humidity conditioning.

The Power and Cooling Pandemic

The Environmental Protection Agency (EPA) took a closer look at the United States' data center energy issue in 2005 after operators presented their power, cooling, and space issues to the agency. A joint study by Jonathan Koomey, Lawrence Berkeley National Laboratory, and Stanford University found that energy use by corporate data centers doubled from 2000 to 2005, and was expected to increase by another 75 percent by 2010 [1]. The study also estimates that power used by servers, cooling and ancillary infrastructure in 2005 accounted for about 1.2 percent of the electrical usage in the United States—enough energy to fuel the entire state of Utah for a year.

The Web, e-commerce, and the growth of corporate databases have contributed to the remarkable increase in the need for computational and storage capacity and the corresponding increases in power use. Compounding the problem is widespread adoption of new high-density bladed servers, which require up to 15 times more power than the last generation of systems, and place a significant strain on data center power and cooling demands. The Uptime Institute presented a white paper at the Uptime Institute Symposium in 2007 estimating the three-year cost of powering and cooling servers is currently 1.5 times the cost of purchasing server hardware and equals 44 percent of an average data center's operating budget [2].

The Fusion-io Advantage

One ioDrive produces 1/1000th the heat of similar performing disk drive arrays.

In 2006, the U.S. House of Representatives passed a bill, H.R. 5646, that calls for a six-month EPA study on data center efficiency. The EPA released a report in August, 2007 analyzing the state of the data center's rising power consumption and outlined a number of proposals to reduce waste. Here are some of the findings from that study:

- In 2006, U.S. data centers consumed an estimated 61 billion kilowatt-hours (kWh) of energy, which accounted for about 1.5% of the total electricity consumed in the U.S. that year, up from 1.2% in 2005. The total cost of that energy consumption was $4.5 billion, which is more than the electricity consumed by all color televisions in the country and is equivalent to the electricity consumption of about 5.8 million average U.S. households [3].
- Data centers' cooling infrastructure accounts for about half of that electricity consumption [3].
- If current trends continue, by 2011, data centers will consume 100 billion kWh of energy, at a total annual cost of $7.4 billion and would necessitate the construction of 10 additional power plants [3].

The EPA report to Congress [3] estimated that if state-of-the-art technology were adopted, energy efficiency could be improved by as much as 70 percent. In addition, the EPA estimated the U.S. could save approximately 23 to 74 billion kWh of power by 2011, representing more than $1.6 billion in energy cost. Savings of that magnitude correspond to reductions in nationwide carbon dioxide emissions of 15 to 47 million metric tons in 2011 [2]. The U.S. Department of Energy (DOE) claims saving a modest 10 percent of total energy use would amount to energy savings of 10.7 billion kWh per year—an amount equivalent to the electricity consumed by one million US households and valued at about $740 million [4].

Others in the industry trying to solve this problem also recognize the impact. Modius, a green data center solutions company, reports:

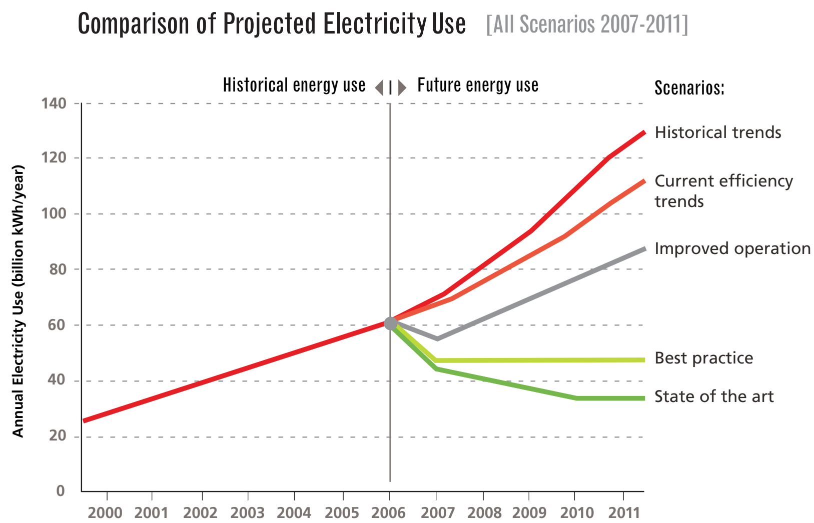

The Fusion-io Advantage

The projected energy costs of the ioDrive are easily 50% lower than an equivalent disk storage array.

"As a result of the overall growth in the mission and capacity of IT data centers, most corporations will see their data center energy costs multiply by three to five times over the next five years, making power and energy the second largest IT cost factor at most data centers worldwide. Beyond the rise in cost, within two years, 50 percent of data centers will simply lack the power and cooling to meet core computing demand. At the same time, regulatory and public pressures are prompting the IT industry to create more efficient data centers. For these reasons, 70 percent of US CIO's surveyed by the Gartner Group in 2006 believe that data center power and cooling is their No. 1 challenge" [6].

Similarly, Michael Bell, research vice president at Gartner Inc., who headed the Data Center Power and Cooling Challenge seminar at the Gartner IT Infrastructure, Operations and Management Summit 2007, calls power and cooling a "pandemic in the world of the data center." He goes on to warn, "By next year, about half the world's data centers will be functionally obsolete due to insufficient power and cooling capacity to meet the demands of high-density equipment" [7].

The Green Grid consortium's whitepaper, "Guidelines for Energy-Efficient Data Center," exposes the painful reality of the power cost to support the IT equipment in a data center [8]. It takes as much as 70% of the power consumed in a typical data center to house, power, cool and protect the servers, appliances, networks and storage equipment. Looking back at the EPA report, that equates to 42.7 billion kilowatt-hours (kWh) of energy or 1.05% of the total electricity consumed in the U.S. in 2006.

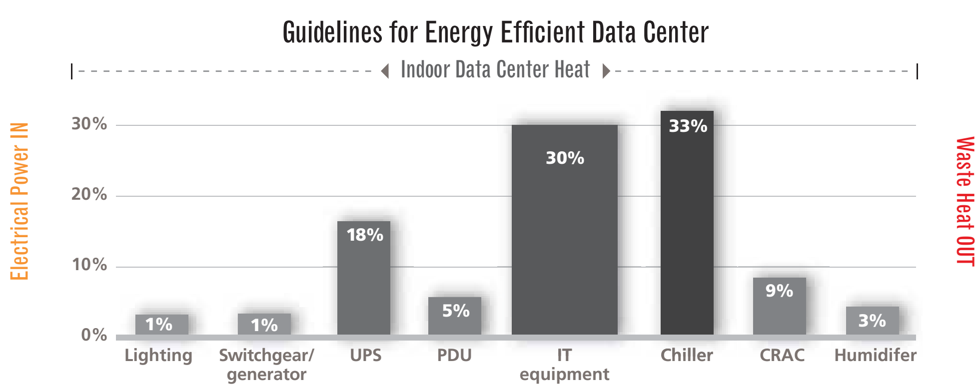

Outside of applying best practices to improve data center efficiency, the clear path to reducing overall power consumption, maintain or increase application performance is to improve hardware utilization and reduce the amount of IT equipment.

How did we get here?

It all starts with the extraordinary improvements in the performance of the central processing unit (CPU, or processor) since its introduction. Intel, the leading supplier of x86 processors, seemingly hit a speed barrier with their Pentium-4 processor at 4GHz back in 2004 and then abandoned the project after declaring victory against their rivals. Why was this so? Why did they then head down a path of developing slower multi-core processors? The answer is simple. Basically, the processor got so fast when compared to the storage subsystem that it was forced into a "hurry up and wait" situation for greater than 90% of the time [16]. As the processor executes a program, it requests information (also known as data) from the storage. When this occurs — and it does very often — the processor must wait a very long time for the data to be delivered. In the meantime, as the processor waits, it is effectively "stalled," performing no other operations. In other words, precious processing cycles are lost as the processor remains idle, though the system continues to burn power. CPU vendors had no choice but to aggregate multiple slower operating processors to extract as much performance out of the lethargic subsystem responsible for storing information.

The Fusion-io Advantage

The ioDrive delivers microsecond – not millisecond – speeds in a new high performance memory tier on the system bus.

It took the computer industry years to get to this point of disparity between the speed of the processor and the supporting memory and storage subsystems. The effect is known as the "Performance Gap" [9]. The processor continued to increase in performance at a substantially faster rate than the supporting memory and storage subsystems, creating a gap. Since 1980, the CPU has improved in performance approximately 60% annually whereas memory subsystem has only increased 10% per year [10] and traditional storage drives even less. Storage drives are particularly bad in that they have doubled in capacity every 24 months but have only improved performance by a factor of ten over the last 15 years [11].

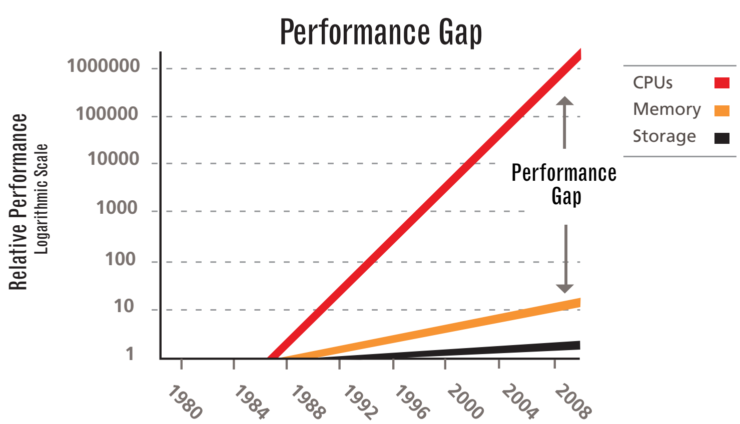

Because of this performance gap, the world's first and fastest 4GHz x86 CPU was no more powerful than previous (slower operating) predecessors. The computer and storage industry has been coping with this performance gap for many years. Consequently, data center architectures have also been constrained by the performance gap, resulting in complex and sophisticated solutions focused solely on storage performance, neglecting its ravenous appetite for power. Unfortunately this translated into literally millions of underutilized processors connected to even more hard disk drives in a feeble attempt to scale performance at the expense of power. Again, the downside is measurable as a decreasing amount of performance relative to the quantity of energy consumed.

Impact of the Stalled Processor

From the processor's point of view, data and programs are located in a hierarchical memory infrastructure. The processor accesses memory to execute programs, store and read information. The memory hierarchy has several layers. The top layers, those closest to the processor, are small and fast; the lower memory layers, however, have been growing in capacity but operate much more slowly, and coincidently consume more power. Computer and server designers use these layers to optimize the system's price and performance but at the same time are limited by the size and speed of each layer.

Memory Hierarchy

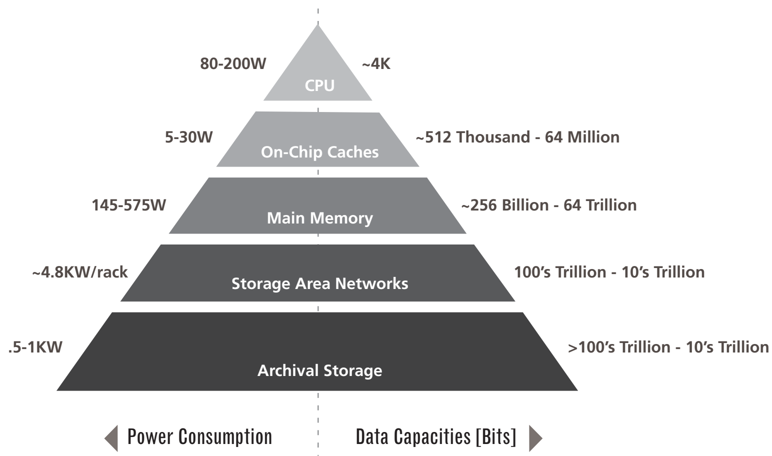

The Fusion-io Advantage

The ioDrive bridges the performance gap between RAM and disks.

Therein lies the problem. For example, modern processors contain small but very fast memory caches. These caches are designed to deliver data requests nearly as fast as the processor's demand, but they are very small compared to the size of the programs or working datasets (by several orders of magnitude), forcing the processor to wait as the memory requests are passed on to the next memory layer, the RAM. The main memory, RAM, is many times larger than caches but several orders slower in speed and greater in power consumption. Main memory is one layer that has been exploited to avoid having to access the slower lower layers. A great deal of power is consumed on large memory configurations. The idea is to try and hold as much of the requested data without having to access the lower layers. Unfortunately, main memory technologies do have a size limit and adding more only increases the amount of energy each server wastes. Both caches and bulk memory are temporary storage devices that are supported by the subsequent layer, non-volatile storage devices such hard disk drives (HDDs).

-300 W on average for each server

Layer	CPU	L1 Cache	L2/L3 Cache	Main Memory	Direct Attached HDDs	SANs
Speed	2-3GHz	2-3GHz	1-1.5GHz	800MHz	>3.5mS or <285Hz	>3.5mS or <285Hz
Size	64 Registers	32-64KB	2-8MB	32-128GB	500MB-5TB	100TB- >1PB
Power	80-200W	Inc. w/CPU	Inc. w/CPU	~4.5W/GB	~10W/TB (5-15W/Drive)	~5.1KW/System bay ~4.8KW/Storage bay

Since, in contrast to its capacity, HDD performance has barely changed for decades (both in terms of access speeds and transactions per second), system designers have tried to compensate by using Storage Area Networks (SANs), which connect 100s or even 1000s of HDDs in parallel, in an attempt to extract better performance from HDDs. In addition, SANs provide additional levels of service, availability and some power savings by aggregating the application servers' direct attached HDDs in to a single shared location. This storage layer is also tiered, using an approach similar to that of CPU caches in the sense that higher speed HDDs are used sparingly at the top of the tier, progressing to slower larger HDDs at the bottom of the tier. For example, "Tier-1 Storage" is built using today's fastest HDDs operating at 15,000 RPM supporting ~300 operations per second or with 10,000 RPM, ~200 operations per second-based HDDs and "Tier-2 Storage" uses lower-speed, lower-cost 7200 RPM HDDs. Again, the approach is to layer from faster to slower technology to optimize performance and cost. It should be noted that 'very fast' in the figure shown below is when compared to other HDDs in the storage hierarchy, but as mentioned above, is several orders of magnitudes slower than main memory, caches and the CPU needs.

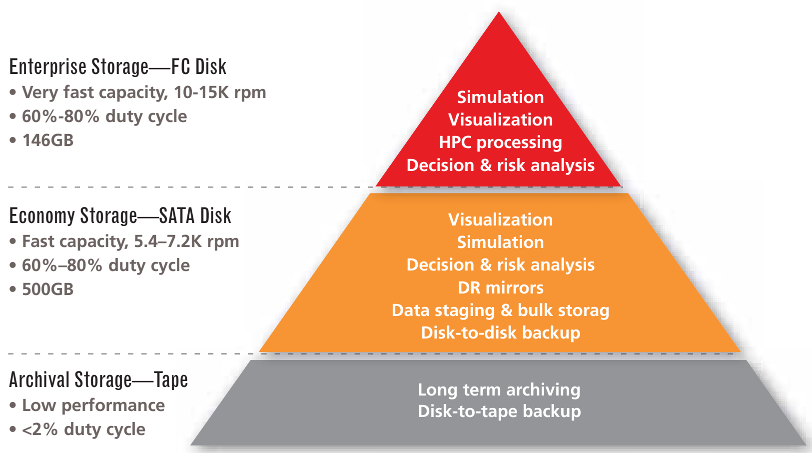

The Fusion-io Advantage

The ioDrive can reduce power consumption by over 73%.

What's not taken into account, however, is the power requirements of this arrangement. Unfortunately, the SAN-based storage layer consumes many kilowatts of power for a capacity hundreds of times larger than but thousands of times slower in performance than the adjacent RAM layer. Making the problem even worse, to maximize performance, Tier-1 storage HDDs are sometimes configured using a technique known as "short stroking" which improves a drives performance by around 15% but restricts the useable storage to as little as 10% of its capacity. Another approach makes use of the network connectivity between layers with the addition of a network appliance devised to cache or buffer data for accelerated storage performance. Once again, the slight improvement in performance is accomplished at the expense of power consumption and cost.

Question of Utilization, Efficiency or Reduction?

Over time, many system architects have made incremental changes to the traditional memory hierarchy in an attempt to improve system performance, reduce overall power consumption, or both. Others even try to justify existing layers. Some argue that the power-hungry SANs consolidate and eliminate the need for direct-attach HDDs normally found in individual servers, but this solution can burden customers with the huge, and ever expanding, costs of cooling and space requirements.

More recently, server virtualization technology, which has the ability to aggregate parallel executing applications on a single physical server, provides the ability to access the untapped performance of a stalled processor. In this scenario, when one application is held up while the processor waits for data, another switches in and makes use of the same processor during this time, thus slightly increasing the processor's utilization. Although this saves in the deployment of additional servers, the same power-hungry data center infrastructure is still required. In point of fact, it must be correspondingly expanded.

Mark Blackburn commented, in a paper published by The Green Grid consortium, that the processor can be placed in a low power state (called "P-State") during idle time to save a substantial amount of power [13]. He shows that the average server burns about 300W, but that when stalled, this amount drops to only around 200W. Placing the processor in to the P-State can achieve an overall system saving of about 20%. This can be an effective approach to reducing wasted power for idle equipment, but it does not address the performance or throughput of a server nor the power-hungry memory storage subsystem.

Other hardware vendors wedge devices between layers, attempting to shadow the functionality of lower layers while trying to improve performance through caching techniques, in some cases consuming up to 450W of additional power for each appliance. For example, one appliance containing large amounts of DRAM memory (one-half of a terabyte) is placed between the servers and the SAN. These approaches all have one thing in common, however: they continue to involve the same basic data center architecture while making only slight improvements to the processor's efficiency, frequently at the cost of higher power consumption.

The Fusion-io Advantage

The ioDrive matches the performance of 600 parallel hard-disk drives in a storage array.

The issue is a question of utilization, efficiency and reduction; partial solutions are not sufficient. What we are really talking about is the amount of power required for a given amount of system performance or application throughput. This metric of performance per watt is the real measurement of success. To improve processing utilization and efficiency, to reduce application stalls, and to eliminate power-inefficient memory layers, the processor demands a very large and very fast storage subsystem. This is the best way to improve the application performance per amount of energy consumed.

Driving for Energy Reduction

Fusion-io has brought to the market a non-volatile solid state storage device based on NAND called the ioDrive. The ioDrive is unique in that it offers the performance of a SAN in one compact storage device. Better still, it does so while consuming a minuscule amount of power. In fact, the ioDrive equals the performance of 600 parallel HDDs in a SAN (comprising the HDDs, the redundant power systems; redundant network equipment; HBAs; and more) but only requires the energy of just one of those hard disk drives. This means eliminating around 10,000 Watts of power with the use of a single ioDrive. This is analogous to impact of Compact Florescent Lights (CFLs) versus incandescent bulbs. "They (CFLs) can use less than one-third the electricity of incandescent bulbs of equivalent brightness and last up to nine years…Top-end 24-watt bulbs promise brightness equivalent to that of a 150-watt incandescent. [17]" With this superior storage technology, not only can the performance and throughput of the data center increase but IT managers can reduce the amount of memory installed in a server, and collapse entire storage tiers, thus dramatically reducing, by orders of magnitude, overall energy consumption.

Fusion-io customers have experienced dramatic improvement in application performance when utilizing the power of the ioDrive. In one case, a customer was able to achieve improved MySQL application performance after installing two

160GB ioDrives into a server, which permitted them to disconnect their 100-hard-disk storage array, shut down 3 of the 4 Intel processors, and reduce the main memory from 64GB to 24GB. The number of application transactions per second increased to 11,000 per second from 9000 per second.

Besides the obvious reduction in costs (including software licenses), IT hardware power consumption was reduced by more than 90% for equivalent application performance. Processor utilization dramatically improved resulting in an increase in the performance per Watt because the ioDrive was able to deliver data at a substantially faster rate. But more importantly, the ioDrive did so at a fraction of the power and without the large typical data center infrastructure. The elimination of the power-hungry IT equipment reduces the data center's demand on energy while maintaining or increasing application performance.

FUNCTIONS (LAYER)	POWER	IO DRIVE POWER USE	REDUCTION
Application Server	~450 W	~200W	~250W
Main Memory	~190 W	~120W	~70W
Network Caches	400 W	eliminated	400W
Tier - 1 Storage (per rack)	~4.8KW	eliminated	~4.8KW
Tier - 2 Storage (per rack)	~3.2KW	eliminated	~3.2KW
Tier - 3 Storage (per rack)	~2.8KW	eliminated	~2.8KW
Tape Storage	<1KW	Unchanged	<1KW

Gartner's research vice president Michael Bell, projects that more than 50% of data centers will exceed 6 kW per rack within two years. Bell expects that number to rise to 70% to 80% within four years due to the increased density of IT equipment, and that the ratio of power to cooling will hit 1:1. In addition, electrical costs per rack will increase by a factor of four, he calculates. Previously, the ratio was 0.5:1. "The cost is basically unsustainable," concludes Bell [14]. The trend towards greater and greater infrastructure costs can be reversed with the deployment of Fusion-io ioDrives.

Elimination of data center equipment is just the beginning. Using ioDrives, the application performance per Watt of power consumed also greatly improves. This and similar customer experiences demonstrate that a fraction of the servers is actually necessary to deliver the same level of performance. Conservatively, once reductions are made, at least half or more of the servers can be removed and the applications consolidated. Assuming that 70% of the power consumed in a data center is outside the IT hardware, for every Kilowatt saved in IT hardware, over 3KW can be saved.

For example, examining a data center with 250 servers, supported by four network caching appliances and SANs with a tape archive storage device, ioDrives can reduces power consumption by over 50% and then eliminating 50% of the servers the savings increases another 22%.

The Fusion-io Advantage

The ioDrive reverses the unsustainable trend of increasing data center cooling cost .

FUNCTIONS (LAYER)	TOTAL POWER	BEFORE	AFTER REDUCTION	AFTER ELIMINATION
250 Servers	~ 450 W	112KW	62KW	31KW
(Main Memory)	(~190 W)	(~47KW)	(~30KW)	(15KW)
4xNetwork Caches	400W	1600W	0W	0W
4xTier - 1 Storage	~4.8KW	19.2KW	0W	0W
4xTier - 2 Storage	~3.2KW	12.8KW	0W	0W
4xTier - 3 Storage	~11.2KW	eliminated	0W	0W
Tape Storage	<1KW	<1KW	<1KW	<1KW
Total		~146.6KW	~63KW	~32KW

Using Fusion-io's technology, IBM Corporation announced Project Quicksilver that combines ioDrive with IBM's storage virtualization technology. It improved performance by 250% at less than 1/20th the response time, took up 1/5th the floor space and required only 55% of the power and cooling compared to the fastest HDD based solutions [15]. In addition, Project Quicksilver is capable of delivering over 1,000,000 operations per second. This is about two and half times faster than the industry's fastest HDD-based storage.

The Fusion-io Advantage

Fusion-io delivers magnitudes of increased performance with a zero floorspace footprint.

How is this possible? Are there hidden costs?

The quick answer is no. In fact, using Fusion-io's technology will produce other benefits as well, the two largest being the big increase in performance and a big drop in cost. Following is an explanation of why Fusion-io can make these extraordinary claims.

As mentioned before, over the last 20 years, advancements in computer processor speed have been following Moore's Law by doubling every 18 months. Also, network speed has been increasing by an order of magnitude about every five years. Mechanical disks, on the other hand, have experienced lackluster performance improvements. This led to the disparity between the disk speed compared to that of CPUs and networks increasing dramatically (almost by three orders of magnitude) leaving a significant performance gap. By some estimates, in today's environment the CPU (including multi-core CPUs) can wait for about 600,000 missed instructions during I/O access for the data to become available. The differences in access delay time of various storage levels is shown in the figure.

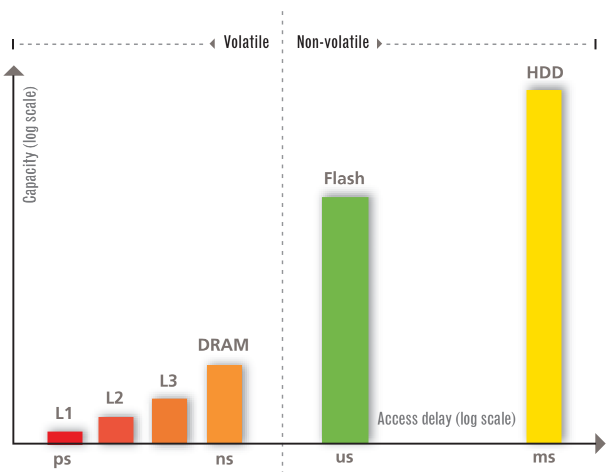

It should be noted that both axes are log scales. With that in mind, we can see the severity in performance between DRAM and HDD.

To drive this point home, the following table shows an interesting analogy:

STORAGE	FOOD	REALATIVE ACCESS TIME
L1 cache	Food in the mouth	Fractions of a second
L2 cache	Get food from the plate	1 second
L3 cache	Get food from the table	Few seconds
DRAM	Get food from kitchen	Few minutes
FLASH	Get food from the neighborhood store	Few hours
HDD	Get food from Mars!	3-5 years

Notice that the difference between Flash and HDD access time is 3 orders of magnitude, which is similar to the food analogy (getting the food from the neighborhood store vs. getting it from Mars!) This helps visualize how much of an effect the proper use of Flash storage has on retrieving data from memory.

As you can see, a huge amount of CPU wait cycles can be eliminated just by using solid state memory instead of disk drives or as a caching device for the much slower hard disks. This can cause a dramatic reduction in power without compromising performance as other solutions do. In fact, there is a dramatic increase in performance as well.

It has already been pointed out that the solid state storage uses much less power than a disk of similar capacity. This is due to a number attributes:

- No moving parts (rotating disks, head servo motors, etc.) to access data
- No continuously rotating disks even when they are not being used and just waiting for the next data access request
- Accesses data electronically instead of mechanically

The Fusion-io Advantage

The ioDrive is up to 1000 times faster at delivering data to the CPU than an enterprise hard-disk drive.

How Fusion-io's Technology increases the Performance per Power ratio

The dramatic power savings with substantially enhanced application performance is due to many factors:

1. No need for huge arrays of power-hungry paralleled disk farms with hopes of reducing the CPU wait time. Using solid state memory for storage instead of (or in conjunction with) hard disks saves a huge amount of power.

2. No need for power-hungry fast disks (15,000 RPM, etc.) since the speed of the disks would no longer be the limiting factor. Using slower, less expensive disks without affecting the system performance will substantially cut back on power consumption.

3. No need to use short-stroking techniques to improve system performance resulting in using only 30% of the disk space (in some cases much less) and hence increasing the number of power-hungry disks needed to handle the required capacity.

4. No need for huge amounts of power-hungry system main memory DRAM since the performance penalty of not finding the information (a.k.a. data) in the main memory would no longer be a huge performance disaster, so long as we have it stored in the solid state storage instead of the much slower disk drives.

5. With the huge savings in power consumption due to these factors, the amount of heat generated by the system is dramatically reduced, which results in another dramatic reduction in power-hungry air conditioning needed to dissipate all of that heat.

6. The amount of space required in the data center can be reduced as well, which also reduces the amount of air conditioning needed, as well as other related facility costs.

7. No need for large battery backup systems to preserve data. Much smaller ones would be good enough to protect the data in the CPU registers, caches and memory. All of these can be saved in the solid state storage (which is non-volatile) and reloaded quickly when power is restored.

When all of these factors are considered, there is a huge power savings in the whole system. That gets multiplied over and over in a data center that has a big number of such systems and servers.

The other impressive effect is while saving all of this power, there is a dramatic improvement in performance as well for the following reasons:

- Access speed is much higher in solid state storage compared to hard disks (electronic vs. mechanical which translates into microseconds vs. milliseconds)
- A big increase in throughput capacity, especially when solid state storage is used in an efficient design that is able to capitalize on its performance capabilities and not throttle it.
- Substantial improvement in response time for I/O intensive applications which is a big percentage of the applications used in the commercial market
- Extended sharing of storage between users is possible without a big impact on response time due to the huge increase in bandwidth capacity offered by the use of solid state storage (which is a problem for disk drives due to their limited bandwidth capacity).

These impressive performance improvements are accompanied with a significant reduction in footprint area needed in the data center.

The philosophy behind Fusion-io's solid state architecture is to unlock a world of possibilities for performance-starved I/O intensive applications running in power-hungry data centers. It provides several orders of magnitudes breakthrough in performance compared to today's traditional disk-based storage systems, while drastically reducing the amount of power consumption.

Ideally, the performance disparity between the processor and the bulk storage could have been limited all along. Unfortunately, this was not the case. The unfortunate net result is excessive demand for power in our data centers—all because the appetite for application performance drove the need for faster processors while the storage subsystems continued to lag behind. Fortunately, Fusion-io's ioDrive is the solution that imposes an inflection point in data center evolution and starts the storage revolution.

References

1. Jonathan G. Koomey Ph.D. staff Scientist, Lawrence Berkeley National laboratory and Consulting Professor, Stanford University, “Estimating total power consumption by servers in the U.S. and the world,” February 15, 2007.
2. Kenneth G. Brill, “Data Center energy Efficiency and Productivity,” The Uptime Institute, March, 2007.
3. Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431, U.S. Environmental Protection Agency, ENERGY STAR Program, August 2, 2007.
4. U.S. Department of Energy (DOE). 2007. Annual Energy Outlook 2007. Energy Information Administration. Report DOE/EIA-0383(2007).
5. Andrew Fanara, “Global Trends: Government Initiatives to Reduce energy Use in Data Centers”, Slide 3, Environmental Protection Agency, Climate Protection Partnership Program, The Green Grid Technical Forum, 2008.
6. Modius, Inc., “The Data Center Energy Problem”, http://www.modius.com/datacenterenergy.htm, 2008.
7. Bridget Botelho, “Gartner predicts data center power and cooling crisis”, SearchDataCenter.com, June 14, 2007.
8. Mark Blackburn, 1E, “Five ways to reduce data center server power consumption”, White Paper, The Green Grid Consortium, Rev 2008-0, 2008.
9. John L Hennessy and David A Patterson, “Computer Architecture: A Quantitative Approach”, Morgan Kaufman, CA, 1996.
10. David Patterson, Thomas Anderson et al., “A Case for Intelligent RAM: IRAM”, IEEE Micro, April 1997.
11. Patrick Schmidt, “15 years of hard drive history: Capacities Outran Performance”, http://www.tomshardware.com/reviews/15-years-of-hard-drive-history,1368.html, November 2006.
12. Robert C. Woolery, “Optimized Storage Solution”, DataDirect Networks, Inc., July 2006.
13. The Green Grid Consortium, “Guidelines for Energy-Efficient DataCenters”, White Paper, February 16, 2007.
14. “Hot storage -- Power and cooling concerns”, Storage Magazine, March 15, 2007.
15. IBM Corporation, “IBM Breaks Performance Records Through Systems Innovation”, http://www-03.ibm.com/press/us/en/pressrelease/24996.wss, August 28, 2008.
16. Zanfard, IBM Corporation, “Cut the Cost in I/T Infrastructure ”, CTS6-12 Reduce Costs, http://www.ibm.com/developerworks/cn/onlinedownload/temp/event/websphere_roadshow/CTS6-12%20Reduce%20Costs.pdf, October 20, 2003.
17. Emily Masamitsu, “The Best Compact Fluorescent Light Bulbs: PM Lab Test”, http://www.popularmechanics.com/home_journal/home_improvement/4215199.html, May 2007.

FUSiON-iO