




Wikibon.org

Networks Go GrEEN

David Floyer, Stuart Miniman



The IT Power Problem

From one perspective, IT trends are moving in the right direction; the energy used per unit of processing is decreasing by 20% per year, and networking products are on the same power improvement curve. However, the elastic nature of IT is such that even with flat IT budgets, the overall energy consumed is growing by 18% per year. Even more significant, IT equipment is becoming denser, and the energy density (watts/sq. ft.) is growing by more than 20% per year, leading to higher data center costs for providing power and cooling to IT equipment.

Figure 1 shows the percentage of power consumption by data center equipment. It shows that processors and storage account for the lion's share of power consumption (86%), and networking accounts for about 9%. It is worth noting that some of the intelligent network power management solutions discussed in this article can allow server and storage devices to power down when inactive.

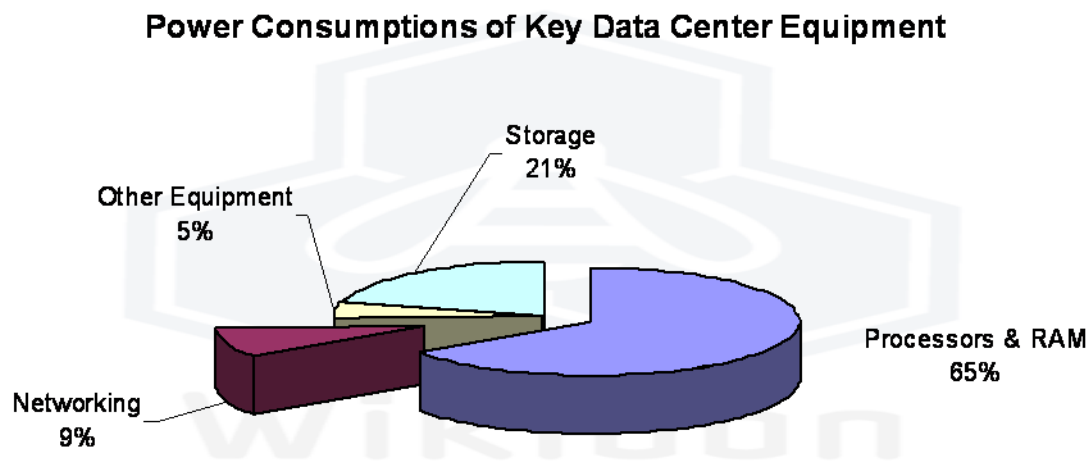


Figure 1 – Percentage Contribution of Power Consumption of Key Data Center Equipment
Source: Wikibon 2011, base on detailed figures in the EPA Final Report to Congress 2007, reallocated to extract all the Networking components

Both processors and storage have been subject to significant attention to reduce power consumption. This has focused on four main initiatives:

- **Turning off power when not required:**
 - Within processors this has been achieved mainly at the chip level through smart power design;
 - Within storage this has been achieved by parking the heads, slowing down disk rotation and stopping disks according to the I/O demands;
 - For both processors and storage, power management software can coordinate power reduction and ensure that unwanted side-effects (e.g., timeouts) do not occur throughout the data center.
- **Leveraging new technologies that increase performance while reducing power:**
 - Within the high-speed network controllers at the server, the use of hardware protocol offloads (such as iSCSI) decreases server power by up to 90W when

- enabled on a single 10G port, lowers CPU utilization, and increases throughput;
 - Within storage by creating a new tier with SSD drive technology, which provides significant gains in performance (especially for many modern applications which require IOPS more than capacity) at a much lower power per GB.
- **Improved data center and equipment design to reduce power distribution and cooling costs:**
 - For instance using improved power supplies (e.g., the recent [Facebook specification](#) including power supplies with 93% efficiency), and improved data center power and cooling management. For example, by using techniques including external air sources, much hotter air in the data center and managed air flows in the data center, Facebook have been able to reduce the power overhead of data center facilities from an industry standard of 100% of the power directly used by the IT equipment to about 10%.
- **Virtualization of processor and storage resources, which has led to significant improvements in utilizations and reductions in overall power consumption:**
 - The average utilization of stand-alone (or bare-metal) servers has been found to be 8% or less. By virtualizing servers and increasing the utilization (i.e., running many low-utilization virtual servers in a single physical server), significant savings can be realized.
 - Virtualization of storage has allowed smart storage controllers to use the optimum placement of storage on the disks as well as techniques such as compression and de-duplication to reduce the size of the objects stored.

The Networking Power Problem

Networking represents a modest percentage of power consumption within the data center (9%), but its portion of total consumption increases as servers and storage become more efficient. Networking has to be tackled, therefore, using the same ideas and blueprint described above for processors and storage.

Due to the architectural practices of cloud computing, virtualization, and I/O convergence, networking infrastructure has become a strategic IT asset. For example, traditional tape backups were taken once a day, and the tapes driven by truck to an off-site location to ensure that recovery could be made in the case of a disaster at the primary site. The modern method of backing up systems is to take snapshots of changed data only, compressing and de-duplicating these snapshots and transmitting them over the network continuously. The results are better recovery times, less data lost in the event of a disaster, and a significant reduction in energy spent on tape trucks.

Power consumption attributable to networking is drawn by:

- Networking cards within servers and storage arrays,
- Switches, routers, and directors used to transmit data within the data center,
- Equipment within the data center to provide for long-distance transmission of data,
- Equipment within the data center to provide power and cooling for networking equipment,
- Equipment to ensure the security and compliance of data networks.

Figure 2 shows the comparative heat densities of different data center technologies. It shows that extreme density communication has the highest heat load, and that normal data center network equipment has the same power density as 1U compute servers. High heat loads imply that the cost of cooling the technology and of providing power are significantly higher than average.

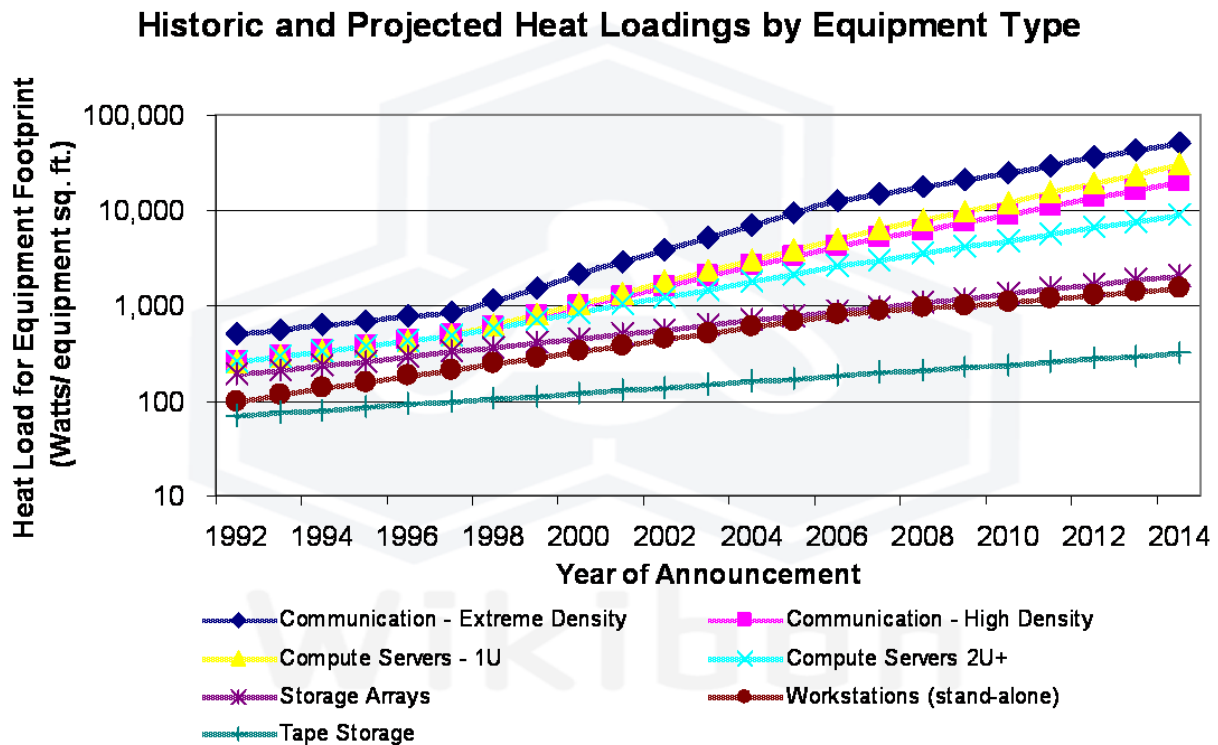


Figure 2 – Historic and Projected Heat Loadings by data center equipment type 1992-2014
 Source: Wikibon 2011, based on historic data from ASHRAE and Wikibon.org projections

Assuming 9% as the portion of consumption attributable to networking and the [EPA estimate](#) of the current national energy consumption by servers and data centers, the power used by data center networking equipment is about 9 billion kWh, representing a \$0.7 billion annual electricity cost. The EPA estimates that with reasonable best practices, these costs could be reduced by about 50%. These savings in electricity use for networking correspond to reductions in nationwide carbon dioxide (CO₂) emissions of about 4 million metric tons (MMT) per year.

Solution to the Problem

Due to the complexities of interoperability specific to networking equipment, there are an unusually high number of networking standards, many of them published by IEEE.

In 2010, the [IEEE approved and published IEEE Std 802.3az-2010TM](#) – Energy Efficient Ethernet (EEE) - which provides a mechanism for reducing power consumption on Ethernet ports and subsystems of networking equipment during periods of no data traffic. By powering down the ports of networking equipment, the energy savings for the physical layer microchip devices can be reduced by up to 70% or more, and overall switch system savings of greater than 33% can be achieved. A good analogy for this is hybrid versus

standard gasoline powered cars. When the hybrid car stops at an intersection, the engine and electric motors shut off, re-engaging only when you press the accelerator. Prior to IEEE Std 802.3az-2010, networking equipment operated similarly to the traditional gasoline powered car, in that the engine continued to run at high RPM's even when the car was stopped. This IEEE standard is a key building block for the networking industry to start realizing energy savings.

Building on the EEE standard, equipment manufacturers can further decrease power consumption by taking advantage of the fact that ports are powered down to cascade that power savings deeper into the networking equipment. By way of example, network interface controllers can power down their PCI bus interfaces and switching equipment can power down portions of their packet processing logic. This type of energy-efficient networking (EEN) can further increase the equipment power savings to around 50%.

“IEEE Std 802.3az-2010, EEE, enhances popular Ethernet interfaces for widely deployed and emerging applications, such as enterprise and the data center, for energy savings. The technologies that the standard provides position an already ubiquitous wired connectivity technology, Ethernet, for green computing and green networking. Broadcom's extensive ING portfolio supports EEE and innovates on top of the standard for further savings.”

~Wael William Diab

Senior Technical Director at Broadcom and Vice-Chair,
IEEE 802.3 Ethernet Working Group

One key area that is being tackled at the standards level right now is the overall management of networking power. The IEEE has standardized a communications mechanism for equipment to exchange power savings capabilities. By coordinating the policies that different pieces of communications equipment use to enable EEE and EEN, more effective power savings can be realized. A simple example would be for a server attached to a switch by two Ethernet ports. During peak compute periods, both ports would be utilized to transfer data to and from the server; EEE/EEN would take effect during the idle periods between transfers. But during off-peak periods, the network management could apply a policy where either side could put the second port into standby and only use one port for transfers. The EEE and EEN operations would realize more aggressive power savings, knowing that the port is in standby. If compute load increased, the equipment would communicate between them to bring up the second port to accommodate the increased traffic.

In addition, the manufacturers of networking chipsets have implemented better power and performance characteristics into an advanced suite of hardware-based stateful and stateless offload services. Running without hardware offloads, the host CPU would consume close to 100% of CPU resources processing the networking stack. In addition to the performance enhancements, full or partial offloads of network protocols significantly reduce the power consumption of a configuration compared with running against the host CPU.

Vendor Implementation Actions - Example

Broadcom provides a complete portfolio of semiconductor communications devices that

can be used to implement Energy Efficient Ethernet equipment.

1. Starting with IEEE Std 802.3az-2010/EEE, Broadcom has implemented this capability across physical layer (PHYs), switch and network controller devices.
2. Furthermore, Broadcom has developed EEN mechanisms in addition to the EEE standard and within its switch and controller devices to provide deeper power savings through the use of more sophisticated policies. Equipment manufacturers such as Dell are using these to develop switch products that can provide enhanced power savings.
3. The IEEE Std 802.3az-2010 was ratified by IEEE in late 2010. To enable equipment manufacturers to adopt EEE as quickly as possible, Broadcom developed its proprietary AutoGrEEEn technology. AutoGrEEEn allows manufacturers to replace existing physical layer devices with EEE-enabled ones without having to re-design the rest of the system. This allows rapid adoption and deployment in existing equipment.
4. Broadcom stateless and stateful offload technologies in 10Gb Ethernet server adapters can provide power savings of approximately 180W/server. The figure below shows the impact of Broadcom's offload technology on CPU use and power consumption.

Network Power Management Hardware Offloads – Protocol Boosters




Figure 3 – Impact of Protocol Offloads on PUE
Source: Broadcom

Future plans include working with Power Management Solutions to improve the management of the complete environment.

Dell has also released PowerConnect GbE switches supporting EEE.

Impact on a Typical Data Center

Wikibon has developed detailed financial models of IT organizations. Table 1 shows the IT budget for a “standard” data center for an organization with \$10 billion in revenue, about 40,000 employees and spending about 4% on IT.

IT Budget (4% of \$10B Organization)	Percentage	\$Budget
Data Center Facilities	7.0%	28,000,000
Network Infrastructure	4.0%	16,000,000
Storage Hw & Sw	5.0%	20,000,000
Server Infrastructure	8.0%	32,000,000
PCs with Basic Software	9.0%	36,000,000
Operations	15.0%	60,000,000
Development	22.0%	88,000,000
Software	25.0%	100,000,000
Management	5.0%	20,000,000
Total	100.0%	\$ 400,000,000

Table 1 – IT Budget breakdown for Wikibon “Standard” Organization with \$10B in revenue
Source: Wikibon 2011

The figures show that the network infrastructure costs about \$16 million per year, and that the cost of environmental for the network is about \$2.55 million. The potential savings to this organization over five years is $\$2.5\text{M} \times 5 \times 50\%$, or **\$6.4M**. Together with other savings in data center power, the potential savings are **\$70M** over 5 years.

Adoption Issues

Networking is an integral part of the data center, and in reality savings from servers, storage and networking should be treated as a whole. In particular, power management software should span all equipment in the data center.

An overall coordinated approach that includes servers, storage, and networking together with data center power management software will be the most effective implementation strategy for most organizations. A clear strategy to adopt the IEEE Std 802.3az-2010 for all networking equipment is an essential component for success, as the networking equipment has to work together and be part of the overall management framework.

Wikibon has looked in depth at a number of organizations that have focused on reducing power. One important management step that successful organizations have taken is to ensure that IT is paying for the power used by IT. All IT capital requests include an estimate of total power budget consumed over the lifetime of the equipment. On the credit side, IT should benefit when energy-saving projects are implemented.

Conclusions & Recommendations

The networking industry has provided an important step forward with the adoption of the IEEE standards that allow data-idle networking equipment to be turned off. Wikibon believes that these standards should and will get steady adoption.

The adoption of these standards should be part of an overall data center power management program that emphasizes the contribution that IT applications are making to carbon reduction within and outside the business. Organizations should calculate and publicize the net carbon footprint of IT now and projected forward (as part of—or in advance of—an overall corporate green strategy).

Reducing the tape trucks and their energy consumption by using networking to assist back-up or enabling telecommuting with IT services are projects that can save significant amounts of energy for an organization. These savings should be captured as part of the initial IT business cases. Success in achieving the objectives should be monitored and promoted as part of an overall IT green strategy, as part of the positive side of the IT green ledger.

At the moment, less than 10% of IT organizations are directly responsible for facilities costs and at best have only an indirect apportionment of facilities costs as a budget line item. Organizations should plan so that IT energy costs are separately monitored, and IT facilities costs are separately accounted for. These should be reflected back to the business as IT budget line items in every project business case.

Future data centers should be planned in line with the power and cooling projections of storage, networking, and server technologies. The heat density of IT equipment is projected to increase, and it is essential that the planning of the data center anticipates these trends and has sufficient capability to deliver cost-effective power and cooling. Consideration should be given to aggressively using hot and cold aisles together with heat monitoring to minimize cooling and air movement and raise data center temperatures.

The use of outside air wherever possible to reduce the cost of cooling, and where possible site (or outsource to) data centers to optimize green/low cost power and air cooling. Using outside air to cool the data center significantly reduces a facility's energy costs, especially if the temperature of the data center has been raised as discussed above. With simple changes and planning, data centers can be cooled with outside air from between 2,000 and 6,500 hours/year, depending on the location of the data center and the temperature of the data center.

Action Item: CIOs should include networking as an integral part of a green data center strategy, and should ensure adoption of the IEEE Std 802.3az-2010 and efficient hardware-based protocol offloads for all networking equipment. CIOs should take full responsibility for environmental costs for IT equipment, and full credit for environmental savings that IT projects can achieve. To stay competitive, organizations should set targets to halve the projected power consumption of the data center over a five to eight year period.

This article was downloaded from Wikibon from http://wikibon.org/wiki/v/Networks_Go_GrEEN on 5/8/2011 and follows the guidelines of the GNU Free Documentation License (<http://www.gnu.org/copyleft/fdl.html>).