climate
savers
computing SM

Considerations for Selecting Power Supplies for Networking Equipment and Evaluating Power Conversion Efficiency

Table of contents

Introduction and Background

As business sectors and consumers shift toward smarter, more connected economies and lifestyles respectively, the demand for computing and networking – in the data center and across enterprise networks to all connected devices -- will continue to grow. With this growth in capacity comes a potential for unsustainable energy use growth, as well; the information and communications technology (ICT) industry must therefore address energy consumption though increased energy efficiency of networking equipment and connected devices. Already, one example of positive technology impact can be found in computing energy use in datacenters, where newer components and systems plus consolidation of workloads through server virtualization has driven projected energy growth from 100% (EPA, 2007) down to 1-2% over the last five years (Koomey, 2011). Today, energy consumption in telecommunications and networking faces a similar challenge. Without explicit focus on energy efficiency, the increasingly connected world may face greenhouse gas (GHG) emissions growth by as much as 6% per year (LBNL, 2010). On the operational side of business, unless today's CFOs and CIOs move to adopt best practices for energy efficiency in networks and across connected devices, they will continue to see the negative effect of unneeded energy use on operating costs and productivity.

Launched in 2007, Climate Savers Computing Initiative's (CSCI) original mission was to reduce greenhouse gas emissions from PCs and servers by promoting aggressive improvements in their energy efficiency. In 2010, several leading companies in the networking industry (Broadcom, Cisco, Emerson Network Power, F5, and Juniper Networks) joined CSCI to address network equipment energy consumption and associated greenhouse gas emissions and to establish best practices for the enterprise and telecommunication networking equipment segments. Our approach to reducing energy consumption for network devices is two-fold - first, accelerate industry's design and product offerings of energy efficient network equipment. Second, educate IT purchasers and network operators regarding best practices and energy efficient considerations when developing and operating network systems for their IT infrastructures. We make a strong case for establishing an open, energy-proportional network device ecosystem that will improve the sustainability of the world telecommunication infrastructure.

To address different facets of the energy consumption opportunities within network equipment, CSCI is launching a series of white papers. These white papers are divided into three distinct but interconnected topics:

1. Considerations for Selecting Power Supplies for Networking Equipment and Evaluating Power Conversion Efficiency

2. Power Management for Networking Devices

3. Energy Efficiency Guide for Networking Devices

Together these three white papers will direct system designers, IT managers, and IT procurement professionals to design, purchase, install, and manage energy efficient networks.

This white paper focuses on unique usage considerations found in networking equipment, and the resulting complications that affect power supply unit selection as well as operational challenges associated with system demand requirements. This paper contains information which should be of interest to a broad cross section of the IT community, but it is particularly relevant to vendors who provide power subsystems to networking equipment suppliers and system designers in networking companies who design and specify power subsystems and system power architectures for network devices.

Networking Equipment Differs From Computing Equipment

Computing Equipment Example

CSCI, in concert with other organizations such as 80plus.org, has done substantial work on the efficiency of power supplies designed for use in computing systems intended for data centers. Computer OEMs have widely and successfully implemented the results of this work, resulting in significant baseline energy efficiency improvements for server platforms over the past several years. It is tempting to simply adopt this previous work when approaching networking equipment by using existing “efficient power supplies” and then declare that the resulting networking equipment is energy efficient. However, as is described below, there can be significant differences in the design and operation of computers or high volume rack-mounted servers and that of networking equipment. The design methodologies and parameters optimized for server systems can yield suboptimal power supply designs if applied to the creation of networking equipment for two reasons.

First, most networking products are modular systems and are deployed with the intention to handle network growth over several years. The result is most systems are deployed in configurations that use a fraction of their full power (or “rated” loading), requiring the power supply to operate at loading of 30% of its rating or less for a significant portion of its operational life; the efficiency of the power supply is significantly reduced at this low level of loading. Second, fixed configuration products or products that are deployed in full configurations will likely not have all the ports connected or may have a mix of high power and low power attached terminal equipment, again forcing the power supply to function at a relatively low load. Additionally, Power Over Ethernet (POE) equipment may power off ports during non-business hours or be provisioned with connected equipment that is rated at a fraction of the rated maximum power supported by the dictated POE specification. In comparison, the predominant power architecture for rack mount servers in data centers is a 1+1 redundant architecture, which provides good protection against single point failure at a reasonable cost in terms of dollars and “consumed system real estate”. A simplified diagram of a 750W server using a 1+1 power architecture is shown in Figure 1.

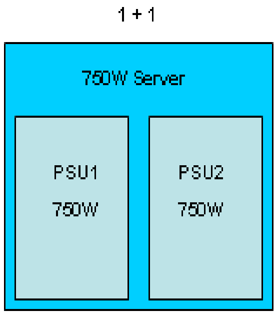

Figure 1: A 750W server using a 1+1 power architecture.

Note that the total, fully configured system power consumption in this example is 750W, and includes two redundant 750W power supplies that in normal operation share load equally. This application allows a single power supply failure without impacting system uptime as the second redundant power supply fully supports the power needs of the system. During normal operation over the life of the system, each power supply operates at peak system load at only 350W, or one half of its loaded rating. However, reality quite often differs from this ideal; It is normal design practice to add “margin” to their power consumption calculations so as not to be caught with insufficient power in the event that a higher performance processor, memory, or disk drive becomes available.

A system with a 750W rating may have an *actual* power consumption of only 700W when fully configured – the difference being an added design margin added. In addition, the actual power consumption of a system depends on its configuration. For example, a system capable of holding four processors and 32 memory DIMMs (Dual Inline Memory Module) may only ship with two processors and 16 memory DIMMs actually installed. Such less than full configurations may actually represent the majority of system shipments.

It is quite common for a system with a 750W fully configured design point to actually only consume something on the order of 460W during actual use. In a case such as this, each of the two 750W power supplies would actually operate at a lower capacity such as 230W, or 31% of capacity for the vast majority of their lifetimes. However, given current state-of-the-art power supply designs as shown in Figure 2, CSCI Gold Power Supplies, achieve near optimal energy efficiency within the operating range of 31% to 50%, falling off substantially at lighter loads.

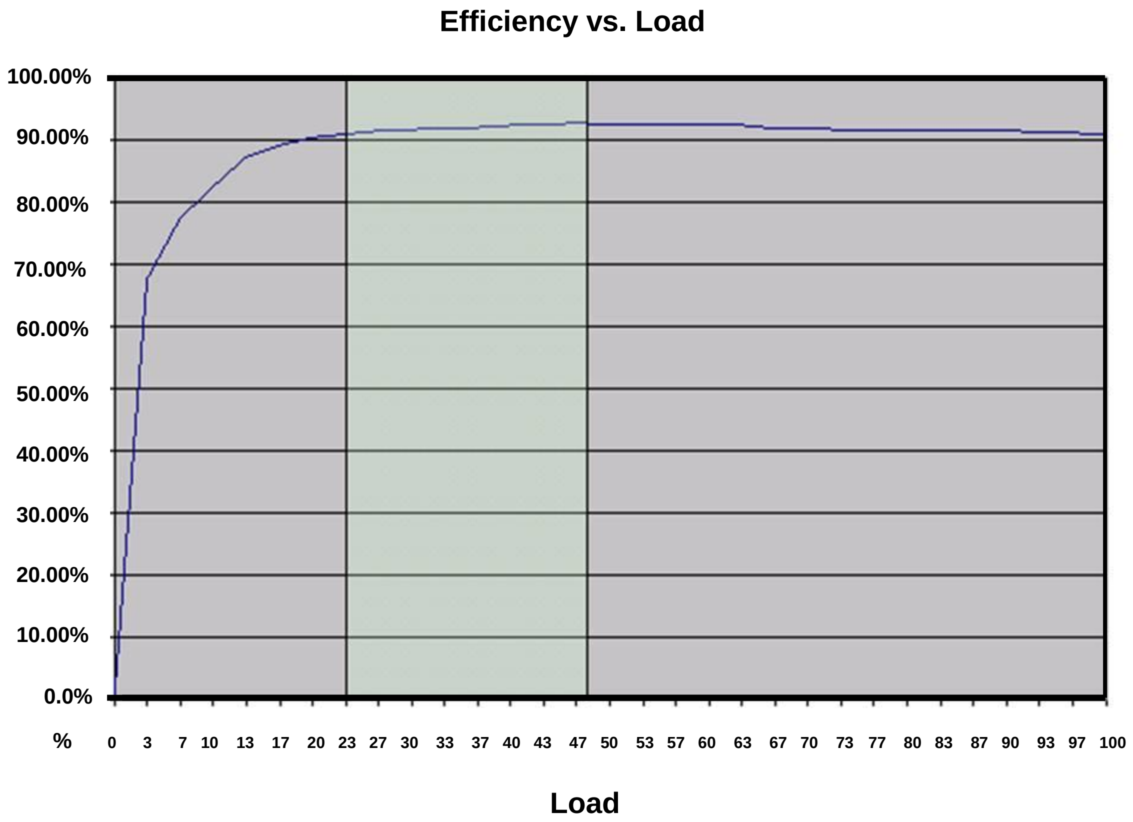

Figure 2: CSCI "Gold" Power Supply, with server typical operating range indicated.

Networking Equipment Example

Now consider the equivalent scenario in networking equipment. As an example, we'll use an Ethernet switch with POE capability and attached VOIP terminal equipment. To simplify the analysis, we'll assume the switch has 100 Ethernet ports with a POE design capacity of 35W per port, with 80% in use and 20% spare. The power consumption of a VOIP terminal is assumed to be 1.25W when idle, and 12W when active. The power consumption of the switching engine is assumed to range from 300W to 500W depending on the number of ports in simultaneous use.

From a design perspective, the total power and cooling design point of this switch would be 4000W, comprised of 500W in the switching engine and 3500W to support 100 POE ports at 35W each. Using the 1 + 1 power architecture illustrated in Figure 3 would require two, 4000W power supplies to provide single point failure protection, each operating at 50% of its rated load.

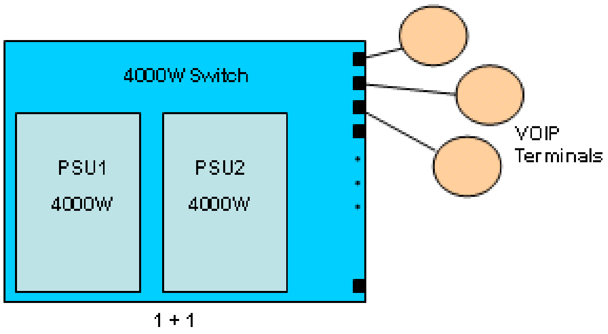

Figure 3: Two 400W power supplies to provide single point failure protection, each operating at 50% of its rated load.

From an operational perspective however, the picture can be radically different. Consider a likely operating point during a typical workday. 80 POE ports are utilized with some percentage of the VOIP terminals active and some idle. As an example, imagine 40 terminals active and 40 idle, which is a reasonably foreseeable condition. Actual power consumption at this instant using our VOIP terminal estimates from above would be 480W consumed by the active terminals, 50W consumed by the idle terminals, and 380W consumed by the switching engine (assuming linear scaling per active port), for a total power consumption of 910W. If each power supply shares the load equally, the operating point for each power supply will be 455W, or approximately 11% of rated capacity. From our typical power supply efficiency curve shown in Figure 4, the power conversion efficiency at 11% load is considerably lower than the peak efficiency, which occurs between 40% and 50% load.

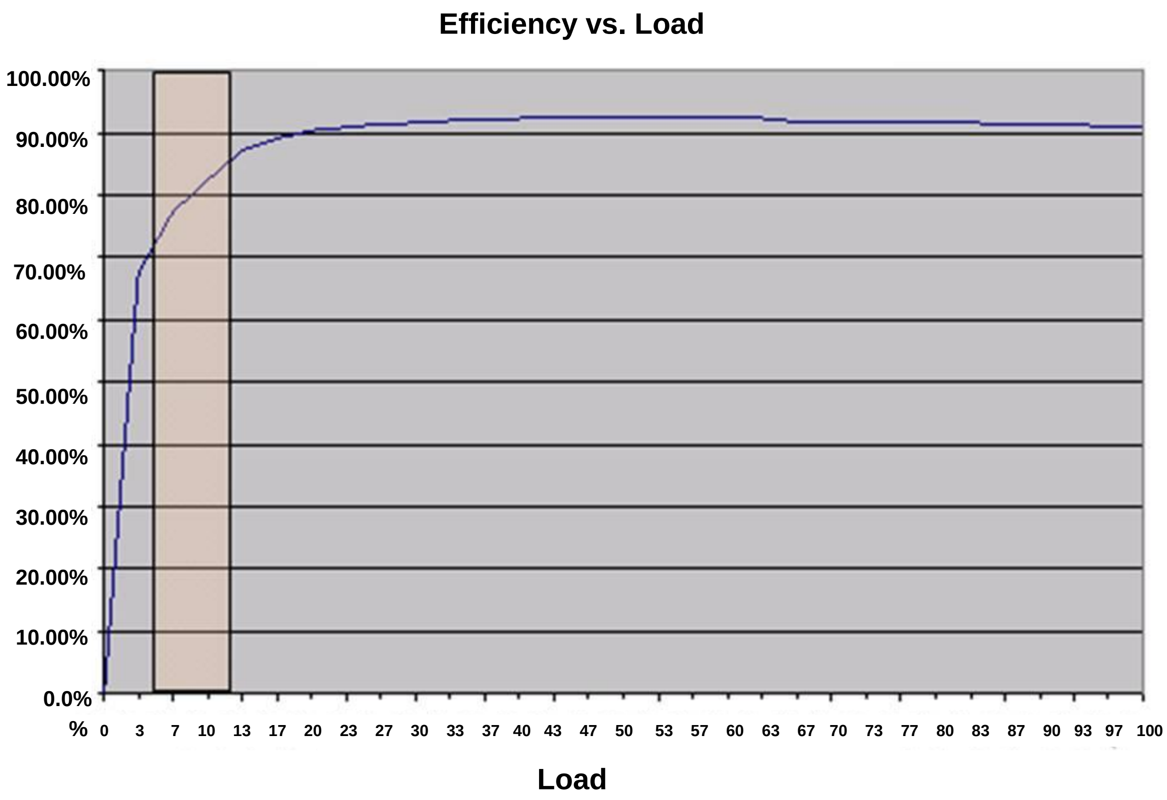

Figure 4: CSCI "Gold" Power Supply, with POE enabled Ethernet switch typical "daytime use" operating range indicated.

This example illustrates that using a 1+1 power architecture and power supplies designed for use in servers and optimized for operating conditions typical of server applications would result in a substantial degradation in the expected efficiency in networking equipment; If usage considerations are taken into account, a 1+1 power architecture is a poor choice for networking equipment applications if the power supplies selected do not maintain high efficiency over the expected variation in operating load.

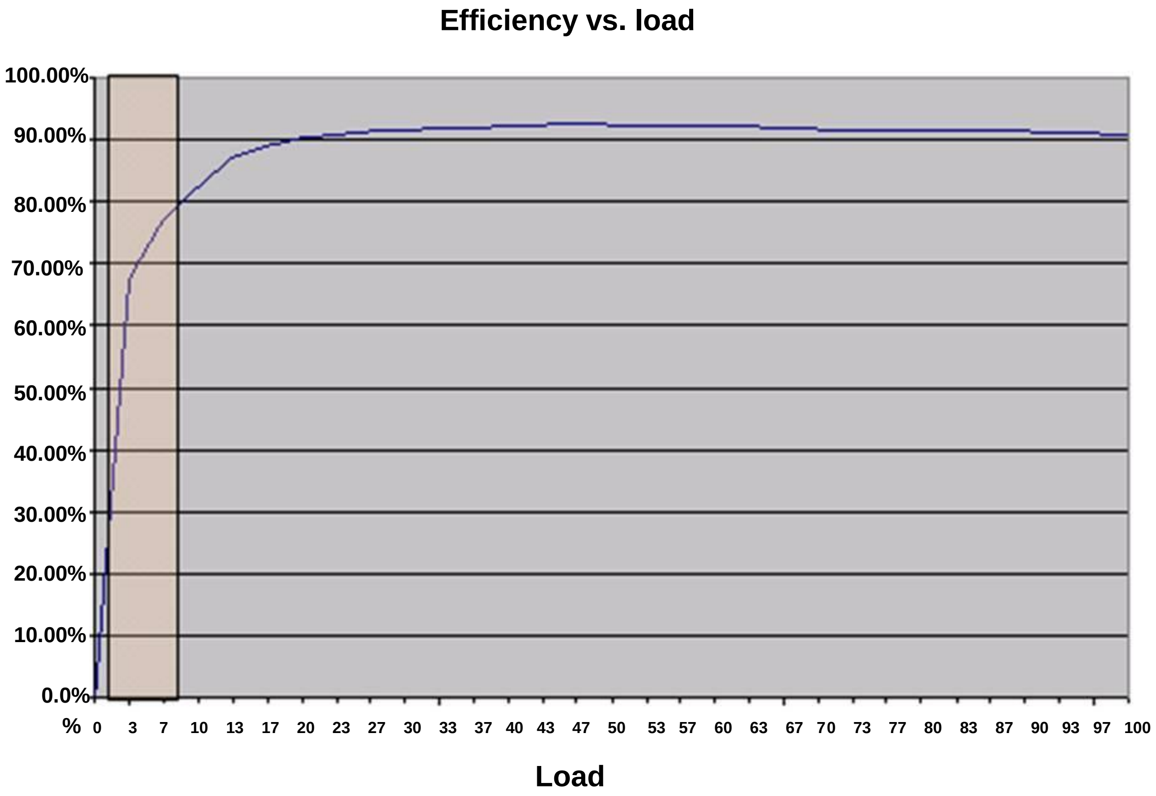

Figure 5: CSCI "Gold" Power Supply, with POE enabled Ethernet switch typical "nighttime use" operating range indicated.

The situation degrades further at night. In a typical office with operating hours of 8AM to 5PM, it is quite reasonable to consider that only 4 VOIP terminals might be active simultaneously late at night, with the other 76 terminals idle. The system power consumption under this condition would be a mere 451W, comprised of 308W in the switching engine, 48W in the active terminals, and 95W in the idle terminals.

As before, assuming equal sharing in the power supplies each power supply unit would be operating at approximately 225W, or only 5.6 percent of its rating. Referring again to our typical power supply efficiency curve shown in Figure 5, the power conversion efficiency at 5.6% load is substantially lower than the peak efficiency.

Challenges that Result

Design Challenges

Networking equipment power supply subsystems face performance challenges that are different than those of their server system counterparts. If traditional server power system architectures such as 1+1 redundancy are simply transferred to networking equipment, the resultant need to operate efficiently across an entire load range of 3% to 100% puts an impractical and unrealistic burden on the power supply designer. Alternate architectures more suited to the unique needs of networking equipment (such as the N+1 approach) allow implementation of advanced system power management features that can idle unneeded power supplies resulting in more nearly optimal power supply load ranges for those power supplies left operating. Adjusting power management to optimize the operating range of the power supplies can also minimize the power factor correction requirements across the expected load ranges, cost, and the constraints presented by Electromagnetic Interference conducted through the power line.

Networking Equipment Power Architecture Decisions

How might a designer best address the challenges discussed above? First, the combination of market requirements and equipment usage model must be addressed in choosing a power architecture to include appropriate modular power design approaches and power management of active modules. This can reduce the wide range of load currents with which the system must deal and enable PSU modules within it to operate at or near their most efficient load points. For example, if the networking system discussed above were powered by a 4+1 redundant power subsystem each PSU module would be rated at 1000 Watts. Without power management implemented, when the expected lightest system load of 5% (200W) is applied, five PSU units each will see a load of 1% (40W) of the total. However, with the capability to implement system level power management, three of the five PSU units could be powered OFF forcing the total load current to be carried by two units resulting in each unit being loaded to 10% (100W) of their respective ratings, which results in a substantial increase in conversion efficiency resulting from nothing more than a change in operating point.

Efficiency Testing

Networking equipment must be evaluated for energy efficiency in typically deployed configurations while managing intended workloads and functions, but the first part of system sizing should be measuring power supply efficiency and power factor as standalone units. If power architecture decisions have been made correctly, system designers can use the same testing protocol that has been established for desktop and server PSUs. This protocol has been developed by the Electric Power Research Institute (EPRI) with contributions from Climate Savers Computing Initiative and other groups and is used by ENERGY STAR™ as the basis for server and client specifications. The test protocol is available at this link. http://efficientpowersupplies.epri.com/pages/Latest_Protocol/Generalized_Internal_Power_Supply_Efficiency_Test_Protocol_R6.5.pdf

To address a usage model and appropriate power architecture for extreme light load (described above) the testing protocol required to assess performance under these conditions does represent a change from the existing EPRI Server test protocol and performance standards which only measure performance down to 10% of the rated load. Note that since extreme light load efficiency and power factor are difficult to measure, special care must to be taken in the selection and configuration of the measurement equipment and in the measurement methodology.

Expected Performance Levels

The energy efficiency of existing networking systems has not been well documented although it may be well known to system designers. Data from testing should focus on the efficiency of networking power systems below 10% of their power rating worst case or, even better, focus on the expected operating load. If a load point down to 3% is possible and the system is expected to operate for significant periods of time at this operating point, then 3% should be the lowest level evaluated or the system design should be changed to operate at a higher percent of load to prevent operating at a poor efficiency. To complete the picture, testing should also determine performance for the remaining portion of the rated load rating with special emphasis on the range below 10% of the rating. The results from such testing can inform designers to possible low efficiency operating points that they may want to avoid by changing their system architecture, selecting a different power supply, or better defining the operating expectations during the specification and sourcing of the power subsystems. The 80 PLUS website has a large database of PSU modules that are characterized for operation from 20% to 100% of rating that is available and downloadable in Excel format that can be referenced at http://www.plugloadsolutions.com/80PlusPowerSuppliesDetail.aspx?id=0&type=2.

PSU modules, whether used in Networking, Storage or Servers, use similar if not identical power components and topologies. The performance of such systems may vary from design to design, however the expectation of performance above 10% loading is in alignment with existing levels established for Climate Savers, 80 PLUS and the ENERGY STAR™ standards.

Testing Setup

There are unique challenges in testing for a load range of 3% to 10%, both for the PSU and for the test equipment. The test setup and methodology must address potential interference from noise generated in the AC source providing power to the PSU and also the susceptibility of the PSU to control loop interaction. The test setup of the Generalized Internal Power Supply test protocol (listed above from EPRI) provides the recommended equipment and needed accuracies. In particular, testers must pay attention to the inclusion of the isolation impedance provided by the Line Impedance Stabilization Network placed in series with the AC source in the testing setup. This impedance is required when measuring power factor below 20% of the unit rating.

Issues with Harmonics and Power Factor

Contribution of Power Factor and Harmonic Generation to Energy Use

Given that the typical operating point of a networking power supply is well below 20% loading and often below 10% loading in POE applications, designers must understand how power factor and harmonic generation contribute to higher energy use. Graph 1 plots Real Power [Watt], Apparent Power [Volt Amperes], and Reactive Power [Volt Amperes Reactive] input vs. % Load for a 1750W power supply. Table 1 displays the lower rate of increase in reactive power [VA-r] to real power [Watt] when the operating point greater than 10% of full load. There is a clear leveling off in the increase of reactive power to real power when operating above 10% load.

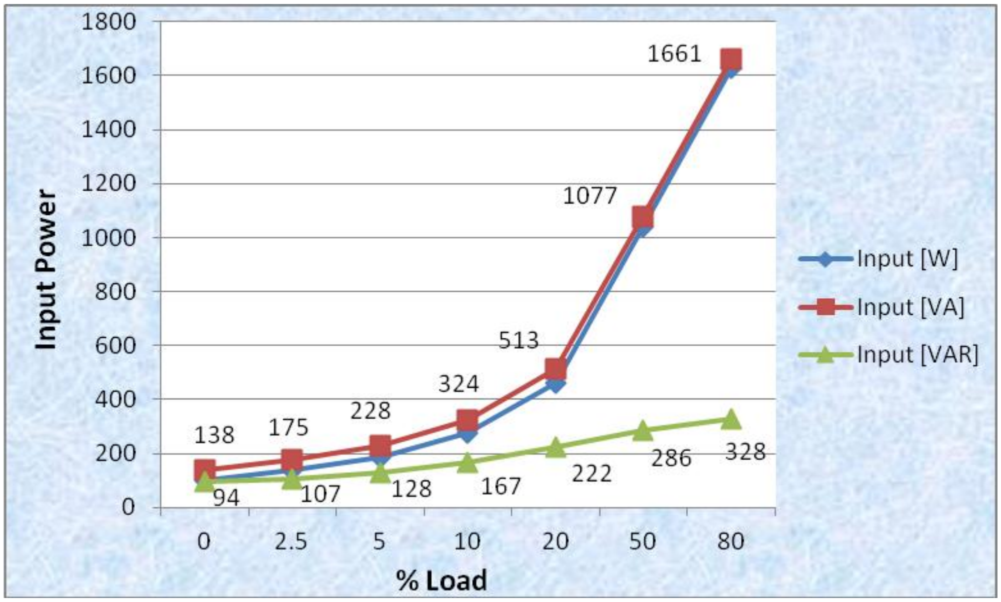

Graph 1: Input vs. Load.

% Load	Load	Input [W]	Input [VA]	Input [VA]	PF
0	0	102	138	94	0.73
2.5	44	138	175	107	0.79
5	88	188	228	128	0.83
10	175	278	324	167	0.86
20	350	462	513	222	0.90
50	875	1038	1077	286	0.96
80	1400	1628	1661	328	0.98

Table 1: Input Power and Power Factor vs. Load.

The take away is that energy use is measured in VA over time [KVA-hrs] or [KW-hrs] when using a true watt meter. Both measurements are based on the real power and reactive power used by the power supply, and are related by the equation: $P = \sqrt{W^2 + VAR^2}$. A simple example demonstrates the importance of minimizing reactive power. For the case of a power supply requiring 1 W and 1 VA-r in equal proportions, then the actual energy is 1.41 times the energy required for a unit only requiring watts. The example above corresponds to a power factor of 0.707, which is not unusual for low cost power supplies running at light loads. Table 1 shows Power Factor vs. load. The above example applies to the fundamental current component or the power source, but in practice switch mode power supplies draw current in high frequency pulses that can attempt to match the frequency of the input voltage source. This allows for high efficiency and small volume power supplies, but it leads to harmonics being consumed by the power supply as a tradeoff for these benefits.

There are industry standards regulating the harmonics distortion produced by power supplies and for the most part such harmonics are not considered to impact energy use. However they are an important factor as shown in Graph 2 that plots the 3rd through 9th harmonics. The 3rd harmonic increases proportionally with the percent of loading. The 5th through 9th harmonics do not increase in any significant way with load and represent a fixed amount of reactive power at these harmonics. The 3rd harmonic increases by over 4 times and presents an increase in reactive power with increasing load. It also represents nearly 10% of the RMS current used by the power supply. For a single 1750W power supply these numbers appear small, but if one scales the same performance to a 1MW data center, then the 3rd harmonic can represent 570 Amps. This current is reflected back to the power delivery system in the form of reactive power and must be supplied by the power plant and ultimately the power grid. This leads to additional heat dissipation in the power distribution system along with increased energy required to deliver the reactive power. The delivery of reactive power reflects back to the power source as an increase in apparent power [VA] and is the main reason a UPS or transformer may need to be over designed when compared to the name plate rating which is usually based on a resistive load or a load drawing power at the same frequency as the power source is capable of operating.

To summarize, both power factor and current harmonics have an impact on the overall energy use of a similarly rated power supply when operating at the same load condition. Specifically higher Volt-Amperes (Energy Use) will be required for lower power factor and also for higher harmonic current content required by the power supplies power factor correction circuit. For server power supplies the performance limits are captured under the Energy Star Program Requirements for Computer Servers. Given that networking power supplies use similar components and topology as computer server supplies it is recommended to meet or exceed this standard.

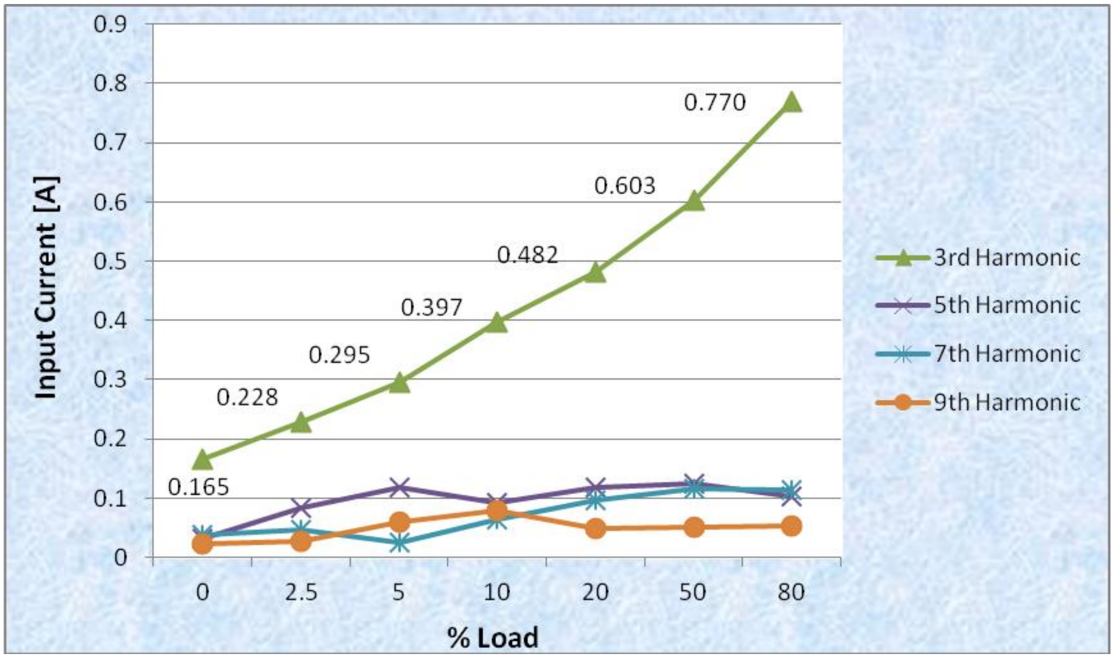

Graph 2: Input Current Harmonics vs. Load.

% Load	I RMS [A]	1^{st} [A]	3rd [A]	5th [A]	7th [A]	9th [A]
0	0.576	0.546	0.165	0.032	0.039	0.022
2.5	0.732	0.684	0.228	0.083	0.047	0.028
5	0.950	0.888	0.295	0.119	0.026	0.060
10	1.352	1.278	0.397	0.093	0.065	0.079
20	2.138	2.074	0.482	0.119	0.097	0.049
50	4.487	4.443	0.603	0.124	0.117	0.052
80	6.922	6.875	0.770	0.102	0.114	0.053

Table 2: Input Current Harmonics vs. Load.

Conclusion

Networking equipment represents a significant opportunity for reducing the total consumed energy in both Data Centers and office environments. The efficiency of individual power conversion components, especially AC-DC power supplies, is an important avenue for reducing total energy consumption within this equipment type, especially when done with other sub-system optimizations. Equally important is for networking equipment designers to fully comprehend how their equipment will be provisioned and deployed so as to be able to adopt the most advantageous power architecture possible and ensure that the power supply characteristics for conversion efficiency and harmonic distortion are properly matched to the application. In addition, a close working relationship between the designers of networking equipment and their power conversion equipment suppliers can help to ensure quick adoption of advances in the state of the art and incorporation of appropriate system power management protocol by the power conversion subsystems.

This publication has been developed by Climate Savers Computing Initiative, the international member-driven consortium working to reduce the energy consumption of ICT (information and communication technologies) by increasing the energy efficiency of computing equipment, increasing the adoption of power management, and providing education about the environmental and financial benefits of energy efficient computing. Since its launch in 2007 the non-profit organization has grown to almost 700 members. Nearly 11,000 people have joined as individuals by taking the CSCI pledge to use power management and to purchase energy efficient computing products. The organization offers an online catalog of PCs, laptops, servers, power supplies and power management software that meet CSCI technical specifications for energy efficiency at http://bit.ly/ba9PDy. CSCI is led by Cisco, Emerson Network Power, F5 Networks, Google, Intel, Juniper Networks, Microsoft, Samsung and World Wildlife Fund.

Contributors to this paper include the following members of the Climate Savers Computing Networking Workgroup. Views expressed in this paper represent the opinions of the authors and do not necessarily represent the views of their affiliated organizations or companies.

Bill Bain	LiteOn
Mike Bax	Emerson Network Power
Akella Chakravarthy	Intel
Wael William Diab	Broadcom Corporation
Doug McIlvoy	Ecova, formerly known as Ecos
Bob Myers	Emerson Network Power
David Owen	Juniper Networks

We recommend the two other CSCI publications in this series on “Energy Efficiency” of Networking Equipment, on the topics of power management (PM BKM) and network-level energy savings (Network BKM).

References

U.S. Environmental Protection Agency, Energy Star Program. 2007. *Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431.* August 2. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf

Koomey, Jonathan. 2011. *Growth in Data center electricity use 2005 to 2010*. Oakland, CA: Analytics Press. August 1. http://www.analyticspress.com/datacenters.html

Lanzisera, Steven. 2010. *Data Network Equipment Energy Use and Savings Potential in Buildings*. September 1. 2010. ACEEE Summer Study on Energy Efficiency in Buildings. Lawrence Berkley National Laboratory. LBNL-3510E http://escholarship.org/uc/item/2998x42q#page-14

EPRI (Electric Power Research Institute) and Ecos Consulting. Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.5 http://efficientpowersupplies.epri.com/pages/Latest_Protocol/Generalized_Internal_Power_Supply_Efficiency_Test_Protocol_R6.5.pdf

80 PLUS® Certified Power Supplies and Manufacturers http://www.plugloadsolutions.com/80PlusPowerSuppliesDetail.aspx?id=0&type=2

.