Delivering a Greener Data Center

Extreme Networks is committed to reducing energy and promoting a greener world both with our networks and our business practices. The company takes a lifecycle approach to power efficiency and management, from product inception, through deployment and operation, and then onto eventual retirement. When we first design a product, we select components with a view to minimizing the amount of hardware required, while leveraging power-optimized network processors, switch fabrics, and physical interfaces. This is reflected in lower operational costs, where we also provide tools for IT to more effectively monitor and manage power utilization in real-time.

Our high-density 1 GbE, 10 GbE, and 40 GbE Ethernet further allow operators to reduce tiers and switches in their network infrastructure, allowing for simpler, faster, and greener solutions. And, they promote longevity through stacking flexibility, ExtremeXOS® upgradability, and a shared hardware architecture across the portfolio that minimizes obsolescence. Then, we promote maximum recyclability when a product reaches end-of-life. Our supply chain is developed with a view to minimizing environmental impact while streamlining our operations, with Extreme Networks new global distribution center built in Hong Kong located in the same region as our manufacturing. Ultimately, delivering a green data center is about preserving the environment while saving costs across the device lifecycle.

Make Your Network Mobile

The Drive Toward Green Data Centers

The migration to cloud and virtualized network and converged storage infrastructures is driving the centralization of both IT operations and network equipment. As with any technology, the industry has developed specific terminology to describe energy efficiency and utilization for data centers. When engineering a data center, two key metrics include Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCIE).

Overall data center efficiency includes networking, servers, the ability to use 'free cooling', and other determinants of Total Cost of Ownership (TCO). Enterprises are increasingly called to task to demonstrate their Corporate Social Responsibility (CSR), one aspect of which is reducing their carbon footprint. Each country has its own standards. The efficiency of the networking equipment deployed therefore directly impacts a data center's 'green' credentials.

As an example, in the US, Leadership in Energy & Environmental Design (LEED), developed by the U.S. Green Building Council (USGBC), provides building owners and operators a concise framework for identifying and implementing practical and measurable green building design, construction, operations and maintenance solutions. LEED, originally focused on office buildings, is now an element in data center design. An equivalent in Australia would be the '5 Star Rating.'

At the device level, the Telecommunications Energy Efficiency Ratio (TEER) from the Alliance for Telecommunications Industry Solutions (ATIS) is key. This measures the switch's data throughput versus the energy (wattage) consumed.

Circling back to Ethernet switching, in the past, the focus on energy utilization was campus equipment and devices, while today it is the IT infrastructure and the 'cloud' data center. Regarding the cloud, we estimate that at current server and networking gear installation rates, within a decade cloud data centers operated by Amazon, Google, Apple, and others will consume more power than any other business sector. Today:

- Cloud data centers consume 3% of all US electricity, growing at over 12% year-to-year.[1]
- A single data center, such as that recently completed by Apple in North Carolina, may draw the equivalent of 80,000 US or 250,000 EU homes.[2]
- On a global basis, in 2010 data center demand topped 330 Billion kWh, close to the complete power requirement of the UK.[2]
- This grows to 662 Billion kWh when including other aspects of the cloud such as connectivity (Figure 1).[3]
- This number is expected to triple or quadruple by 2020.[2]

Clearly, each networking vendor must do its part.

The IT industry worldwide is supposed to have a bigger carbon-footprint than the airline industry by the end of 2010 330BKwh = 227.5Mt at .00068956mT per Kwh; airline industry is 170Mt

Sources epa.gov and EQ2 Insight white paper

[1] Koomey, Worldwide Electricity Used in Data Centers, 2008
[2] Greenpeace, How Dirty is your Data?, 2010
[3] "Make IT Green: Cloud Computing and its Contribution to Climate Change," Greenpeace, 2010

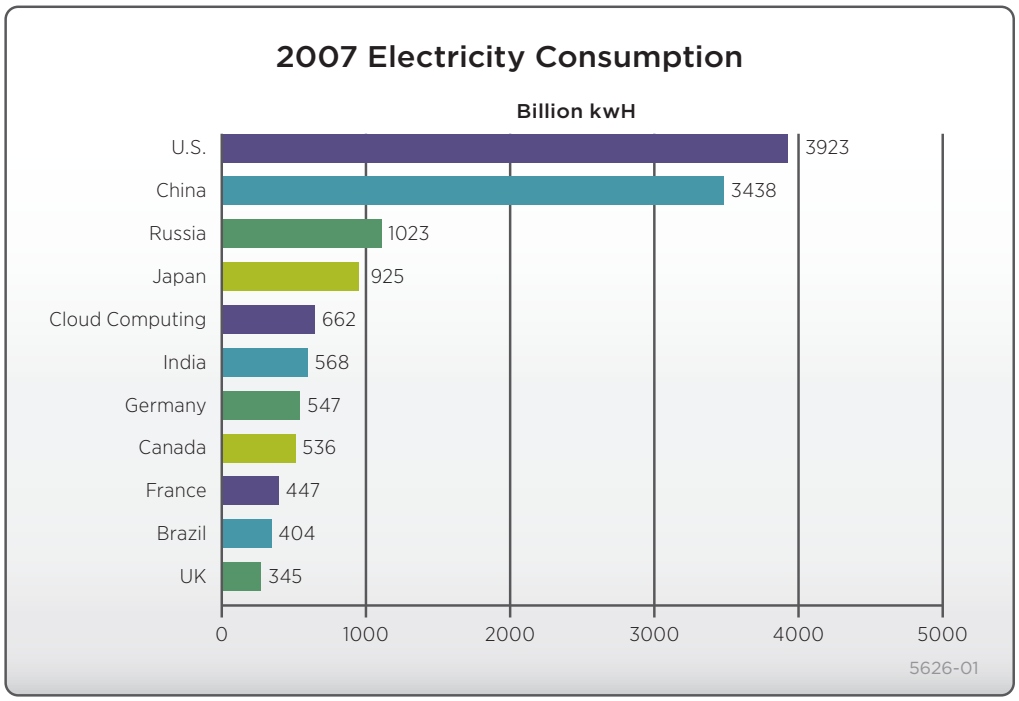

Figure 1: Global Cloud Power Consumption[4]

And it gets worse. In April of 2011 Greenpeace released a seminal study on data center power entitled 'How dirty is your data?' Though not focused on absolute hardware power requirements, the report did highlight the sources of energy for most of the major cloud players.

Reading between the lines, even if different operators are limited in their choice of power sources, it is better if they can deploy energy efficient hardware. The report called to task companies such as Apple and HP, major consumers of 'dirty' coal-generated energy, while giving high marks to Yahoo, Akamai, and Google. The analogy is the electric car, clean in concept and possibly in manufacture, but dependent upon power generated from coal and even diesel to recharge its batteries.

One reason for growth in the cloud is the acceleration of new applications and use cases. For example, media-hungry smartphones and tablets have appeared on the scene only over the last three years, and Intel has stated that a server is required for every 600 smartphones or every 122 tablets.[4]

- In a future of two billion smartphones and 500 million tablets, 3.3m+4.1m = 7.4m servers are required just to service the mobile users.
- If a server draws 427W on average, it draws 3741 kWh annually.
- These 7.4m servers therefore draw 27.6B kWh.

Now, using virtualization, the overall server requirement will of course be reduced, but the point is clear. Everywhere one turns; cloud demand is increasing.

[4] San Jose Mercury News

In-Service Power Utilization, Thermals, and Monitoring

Once the switch is powered-on, top considerations include average wattage per port, power supply efficiency, the ability to dynamically adjust power on a port-by-port basis, as well as real-time monitoring.

Building on our fifteen years of switching expertise, Extreme Networks Summit® and BlackDiamond® series product families are leaders in power efficiency in the industry for a given port speed in their switch class (Table 1). Though port-by-port savings may seem modest, they add up when deploying an enterprise backbone or large data center.

Extreme Networks Product	Power Consumption	Closest Competition	Power Consumption	Competitive References
Summit X450e	121W	Catalyst 3750E-48PD	152W; measured Cisco provided	100% load; no PoE out
Summit X460-48t	129W	Catalyst 3750X-48t	140W	Data sheet
Summit X480-48t	182W	Juniper EX4200-48t	199W	Data sheet
Summit X670-48X	120W (48 ports)	Arista 7050	128W (64 ports)	Lippis test
Summit X670V-48x with 40G uplinks	290W	Arista 7050	384W	Configuration guide
BlackDiamond 8000 Series	490W	Cisco Catalyst 6500	1665W	Tolly test; equivalent configurations
BlackDiamond X*	< 5W/10G	Arista DCS-7508	~ 10W/10G	Data sheet

Table 1: Comparative Power Consumption – like-for-like models

*Future availability.

In 2008 the Tolly Group released a report comparing the BlackDiamond 8000 series to both the Cisco 6509 and the Foundry (now Brocade) BigIron RX-16 with the following results:

- Loaded with 8x10G ports and 48xGbE ports, the BlackDiamond 8000 series drew 3.4x less power than Cisco's 6509 and 2x less than the Foundry RX-16.
- Total power consumed was 535 watts.
- At full chassis loading the BlackDiamond 8000 series consumed 3.3x and 1.7x less power versus the two platforms.
- Under no load, with only the management modules and fabric connected, the BlackDiamond 8000 series drew 219 watts versus 617 watts for the Catalyst and 452 for the BigIron.

What is important to note is that although these results are three years old, the lifespan of switching platforms within the data center is substantially greater. In fact, Extreme Networks platforms have demonstrated exceptionally long in-service lifespans, with uptimes of over ten years common.

Our recently announced BlackDiamond X* series carries this efficiency forward, which based on lab testing, has the lowest 10G per-port power consumption of any product in its class (Figure 2).

Vendor References

- Juniper: Power reference for QFX3500: Data sheet http://www.juniper.net/us/en/local/pdf/datasheets/1000361-en.pdf
- Juniper: Power for QFX fabric/interconnect power not available; estimates based on Arista
- Juniper: Pricing for QFX3500 http://www.juniper.net/us/en/local/pdf/datasheets/1000361-en.pdf
- Juniper: Pricing for QFX fabric/interconnect power not available; estimates based on Arista
- Cisco: http://www.cisco.com/en/US/products/ps9402/products_data_sheets_list.html
- Cisco: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps11541/data_sheet_c78-651097.html
- Cisco: Pricing from published pricelist, January 2010
- Arista: Power reference for Arista 7050 fixed-configuration switch and 7500 modular switch: Arista product quick reference guide
- Arista: Pricing reference for Arista 7050 fixed-configuration switch http://www.aristanetworks.com/en/news/pressrelease/323-pr-20110328-01
- Arista:Pricing reference for the Arista 7500 modular switch: http://www.theregister.co.uk/2010/04/19/arista_7500_modular_switch/
- For the 7500, fully loaded pricing supports Arista's claimed $1200/10GE port price. This value was used to calculate 7504 cost.
- Force 10: http://www.force10networks.com/products/ZettaScale.asp
- Force 10: http://www.force10networks.com/products/pdf/Force10_S4810_DS.pdf
- Force 10: http://www.force10networks.com/ResourceLibrary/pdf/F10Report_custom_4810_2011-01-24.pdf

*Future availability.

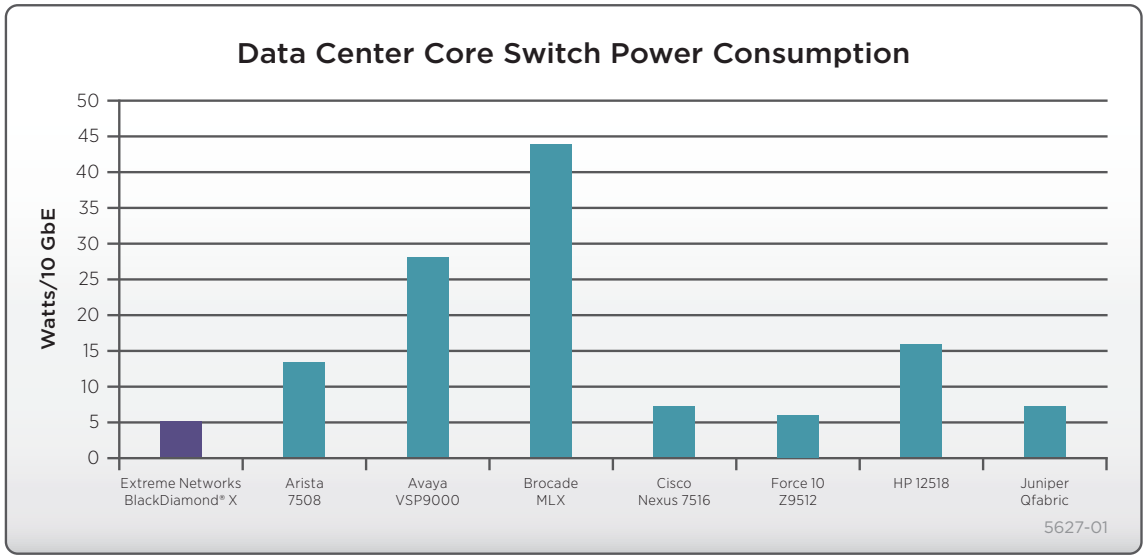

Figure 2: End-of-Row Switch Power Consumption

For comparison, consider the BlackDiamond X series versus our major competitors in a fully loaded configuration. When tested in a lab environment, the nearest shipping competitor, Arista, requires almost twice the power per port and offers only 50% the density. Table 2 compares this power consumption per 10 GbE interfaces for a typically configured EoR chassis as reported by the individual vendors.

Platform	Power/10 GbE	10/40 GbE Capacity	Loaded Power
Extreme Networks BlackDiamond X	5W	768/192	3,840W
Arista 7508/7516	13.2W	384/na	2,550W/5,750W
Avaya VSP9000	28W	240/na	6,720W
Brocade MLX	44W	256/na	11,391W
Cisco Nexus 7516/7508	7.59W	512/na	3,886W/9,719W
Force 10 Z9512	6.25W (est)	480/120	3,000W
HP 12518	16W	128/na	2,048W
Juniper QFabric	7.5W (est)	512/128	5,000W

*Table 2: End-of-Row (EoR) Power Comparison**

NOTE
In some cases power requirements are an approximation due to differences in vendor reporting which varies from 30% load, 30% weighted average load, and 50% load.

*Actual results may differ.

Table 3 provides this same comparison for Top-of-Rack (ToR) switches with either 10 GbE or 40 GbE uplinks depending upon system capabilities.

Platform	Capacity	Loaded Power
Extreme Networks Summit X670V	48x10+4x40	257W
Arista 7050	64x10	350W
Avaya VSP 7024	24x10+2x40	200W
Brocade VDX 6720	60x10	360W
Cisco Nexus 3064	64x10	345W
Force 10 S4810	48x10+4x40	220W
HP 6600	24x10	405W
Juniper QFX 3500	48x10+4x40	365W

*Table 3: Top-of-Rack (ToR) Power Comparison**

Continuing the analysis, Table 4 details total hardware and power costs for a data center supporting 2,256 10G servers. This topology utilizes the maximum density of the BlackDiamond X series in the core with a standard Multi-chassis Link Aggregation (M-LAG) topology to connect the core switches in order to support VMware vMotion, and redundant 2x40 GbE (or 8x10 GbE) links to each core switch. Although the analysis assumes 10 GbE, due to the BlackDiamond X series 10 GbE density, savings are even greater for a mix of 1 GbE and 10 GbE servers or where the data center can be consolidated into a pair of BlackDiamond X series switches, bypassing the ToR layer, if density and existing cabling permit. Tables 2 and 3 provide sufficient data to conduct an overall data center comparison vs. Arista, Cisco, Juniper, and Force 10.

Vendor	EoR Switches	Core Oversubscription	ToR Switches	Total Power Requirement	Total 3-Year DC Cost
Extreme Networks	2	3:1	94	31,838W	$87,519
Arista	8 (6 x aggregation + 2 x core)	16:1	94	58,800W	$161,634
Cisco	6 (4 x aggregation + 2 x core)	16:1	94	75,192W	$217,376
Force 10	8 (4 x aggregation + 2 x core)	16:1	94	44,680W	$122,820
Juniper	4	3:1	96	55,440W	$152,398

*Table 4: Overall Data Center Power Comparison**

NOTES
0.1046 per KwH = U.S. average retail cost (in US$) of commercial grade power as of June 2010 as per Dept. of Energy Electric Power Monthly (US Energy Information Administration) Loaded power = 1.33 x base to account for cooling costs.

*Actual results may differ.

On high-end platforms such as the new BlackDiamond X, additional efficiency is a result of new digital power supplies, where a higher proportion of the incoming power is usable by the linecards due to a higher efficiency Power Supply Unit (PSU) design.

Lower power consumption also results in cooler-running chassis, creating a double-win for the data center operator. Our products have demonstrated optimized airflow design in comparison to the competition (Figure 3), design expertise reflected in our successful BlackDiamond 8000 series as well as our new BlackDiamond X that runs cooler than any comparable chassis (Figure 4).

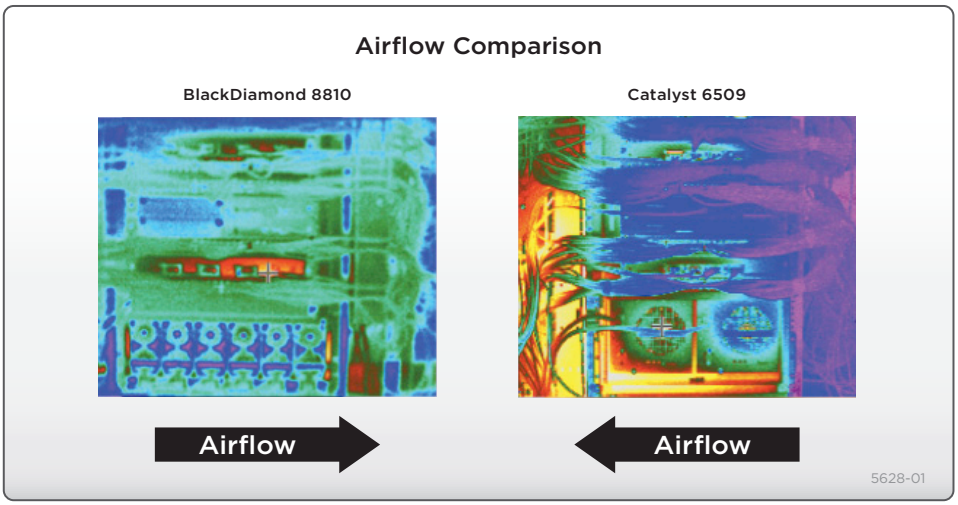

Figure 3: Airflow Comparison

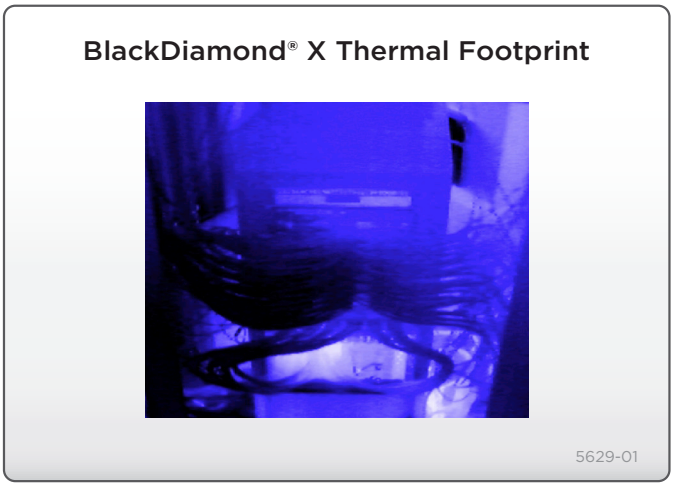

Figure 4: BlackDiamond X Series Thermal Footprint – chassis running cool (blue) in upper part of picture, with 'hot' chassis, a test platform, (white) in lower part.

Power efficiency also continues to grow with each new generation of merchant silicon technology. In the same way that Intel has stated their desire to lower utilization across their product portfolio, from 35 to 15 watts over the next few years, network processors follow the same curve, with today's 10G interfaces consuming the same power as 1G interfaces from only a few years back. Sounds a bit like average MPG from the auto companies, doesn't it? Companies with deep merchant silicon expertise, able to compress their new product development cycle, are best placed to take advantage of this trend.

Within the core, where a BlackDiamond 8000 series or BlackDiamond X typically aggregate multiple edge or ToR switches, interfaces are always operational, and therefore at full power. The situation is very different at the edge where devices such as PCs, VoIP phones, or even servers may have differing power requirements throughout the course of a day. Here, intelligent power management is required.

Intelligent Power Management

One key ExtremeXOS capability is Intelligent Power Management. Building on our ExtremeXOS Universal Port capability of identifying the type of device connected to an interface and modifying the port parameters accordingly, ExtremeXOS Universal Port Manager continually samples the link for activity. If there is no activity, power for the link is shut down, and the power is sampled at intervals for any new activity. The other option is to turn off power to certain interfaces based on policy.

For example, the IT administrator can place all ports serving VoIP phones in a group and apply time-of-day policies. This leverages the ExtremeXOS extensibility framework. By creating this policy based on time-of-day, overall power costs may be reduced by up to 70%, achieving savings of up to 91% versus competitive chassis-based solutions based on Extreme Networks lab testing (Figure 5). The late-night employee may of course override this policy.

In the near future, this capability will be enhanced in a standardized way via Energy Efficient Ethernet (EEE) support on forthcoming chipsets. EEE defines new low-power idle modes that permit the link to reduce its power use if there is no data. For example, VoIP phones are normally idle except for control. Instead of powering off the link entirely, which is probably not desirable for phones, it switches to a lower power level, only raising this if there is an active call. EEE of course is a copper-centric technology. The technology will also come in use with 10G servers which draw proportionately more link power, but which are also idle during some parts of the day.

Another element of power management is variable speed fans. Here, when the stackable or chassis switch is operating at less than full capacity, as reflected by temperature sensors, the hardware requires less cooling, and the fan speed is intelligently adjusted. Most of Extreme Networks Summit switches support this, as well as the BlackDiamond 8000 series and BlackDiamond X series. A secondary advantage of variable speed fans is noise reduction.

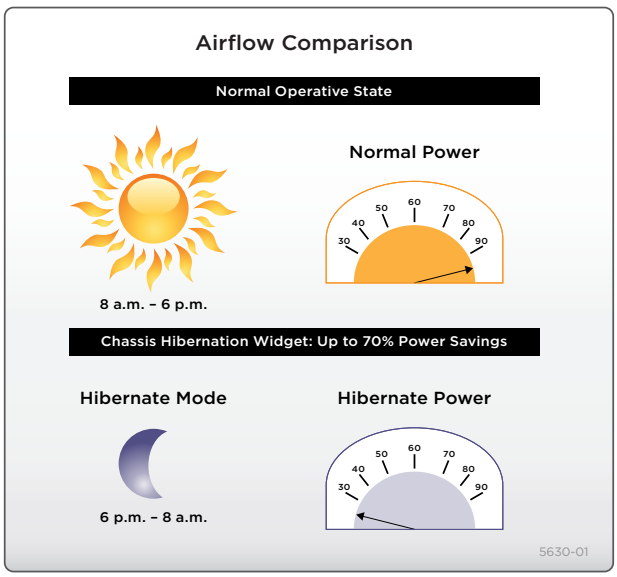

Figure 5: Extreme Networks Intelligent Power Management

Product Retirement

Sometimes neglected, but critical to a platform's environmental impact, is what happens to a product when it reaches the end of its useful lifespan. Beyond the Restriction of Hazardous Substances Directive (ROHS) and other international standards that define material recyclability, we design Extreme Networks switches to use the minimal material necessary for the features and performance delivered. This design approach is reflected across the product line, from the SummitStack™ through the new BlackDiamond X series.

Beyond the Switch

Although Extreme Networks is doing its part in lowering data center power requirements, there are also industry best practices that influence the overall data center efficiency as described earlier. At a recent Uptime Institute Symposium, a number of options were discussed, with the following considered to be the most viable (including the % of respondents currently implementing).

Server virtualization	82%
Hot aisle/cold aisle containment	77%
Power monitoring and benchmarking	67%
Raise inlet air temperatures	57%
Variable frequency drives on chillers and pumps	46%
Modular data center design	34%
Air-side economizers	25%
Power-down features on servers	22%
Water-side economizers	21%
Liquid cooling	18%
Direct current power	7%

Extreme Networks switches support hot/cold aisle containment via front-back or back-front cooling, power monitoring, as well as switch interfaces that adapt to server power-down capabilities. In addition, low power consumption results in a lower operating temperature, directly impacting the acceptable airflow temperature. And, the rich ExtremeXOS feature set facilitates virtualization.

Case Studies

Our energy efficient switching products have been deployed across the globe, and within many vertical markets. The case studies below are only a few examples of our successes.

- Manchester Essex Deploys Eco-friendly network with Extreme Networks
  http://www.extremenetworks.com/libraries/casestudies/CSManchesterEssex_1729.pdf
- Green Hospital: Muskogee gets LEED Certification
  http://www.extremenetworks.com/libraries/casestudies/CS_Muskogee_2209.pdf
- Austrian Ministry of Interior
  http://www.extremenetworks.com/libraries/casestudies/CSAustrianMoI_1751.pdf

Conclusion

From the Summit edge to the BlackDiamond core, Extreme Networks has consistently delivered energy-saving designs to both enterprises and service providers. This expertise is also reflected in the complete product lifecycle, from inception through recycling. And, as our merchant silicon vendor and contract manufacturer introduce new, more efficient designs and manufacturing techniques, we too are quick to incorporate these innovations across our products. Extreme Networks®– purple on the outside; green inside.

References

- Environmental Protection Agency (EPA) Green Power Partnership Website http://www.epa.gov/greenpower/pubs/calcmeth.htm
- EQ2 Insight White Paper: "Sustainable Flying: Biofuels as an Economic and Environmental Salve for the Airline Industry" http://www.jdasolutions.aero/downloads/EQ2_Report_Aviation_biofuel.pdf
- Alliance for Telecommunications Industry Solutions (ATIS) Green Initiative Website http://www.atis.org/Green/index.asp
- San Jose Mercury News, May 17, 2011 "Intel CEO Paul Otellini touts focus on power-saving chips" http://www.mercurynews.com/business/ci_18081305?nclick_check=1
- Wikipedia article on US Green Building Council (USGBC) http://en.wikipedia.org/wiki/U.S._Green_Building_Council
- Koomey, Worldwide Electricity Used in Data Centers, 2008
- "How Dirty is your Data?" Greenpeace, 2010
- "Make IT Green: Cloud Computing and its Contribution to Climate Change," Greenpeace, 2010
- US Energy Information Administration (EIA) Website http://www.eia.gov/cneaf/electricity/epm/table5_6_a.html
- Extreme Networks Inc. BlackDiamond 8810 Core Switch "Green Evaluation of Energy Consumption vs. Cisco Catalyst 6509 and Foundry BigIron RS-16", Tolly Group, January, 2008
- Uptime Institute Symposium 2011 Website http://symposium.uptimeinstitute.com/

Make Your Network Mobile

Corporate and North America
Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, CA 95051 USA
Phone +1 408 579 2800

Europe, Middle East, Africa and South America
Phone +31 30 800 5100

Asia Pacific
Phone +65 6836 5437

Japan
Phone +81 3 5842 4011

www.extremenetworks.com