NETEZZA

The Power to Question Everything™

The Green Revolution in Data Warehousing:

The Power, Cooling and Footprint Advantages of the Netezza Analytic Appliance

Whitepaper

Executive Summary

Skyrocketing power costs and finite power availability are forcing many companies to re-examine their IT strategies and budgets. With their intense workload — dozens or hundreds of processors and large numbers of spinning disk drives straining over massive queries — data warehouse systems based on general-purpose designs have long been one of the major consumers of power. As data volumes continue to grow and complexity of the enterprise analytic environment increases, companies will continue striving for a competitive edge through better-performing data warehousing and analytic infrastructures. The Netezza Performance Server® (NPS®) analytic appliance redefines energy efficiency in data warehousing, while providing cost-effective query performance beyond the reach of other systems.

This paper examines the power crisis in the data center, and the environmental impact of general-purpose processor technology (web servers, email servers, OLTP servers, etc.). Today's tightly packed servers consume more power and give off more heat than their predecessors as the latest generations of x86 chips drive enormous amounts of data through microscopic circuitry. But is this type of computing, and the power it requires, optimal for data warehousing?

The paper then examines the NPS system from an energy efficiency perspective. Netezza's family of streaming analytic™ appliances uses a fundamentally different processing architecture, based on streaming technology originally developed for applications such as video games, industrial process controls and medical imaging that also happens to be perfect for accelerating query performance of massive databases. The Netezza architecture greatly reduces processor workloads, allowing the NPS system to outperform more expensive rack and blade servers at a fraction of their power consumption, heat and footprint. Even the new generation of more energy-efficient chips cannot compare to the 5 Watt processor in Netezza's intelligent storage nodes.

To show how technology differences translate into power ratings, the paper concludes with an analysis comparing the energy efficiency of the NPS system with data warehouse systems of equivalent storage capacity. Values are taken from the American Power Conversion Corp. (APC) on-line UPS selector that provides power and cooling ratings for all major blade and rack server configurations. The results are clear: Netezza's analytic appliance is setting the pace in energy efficiency as well as performance.

Power Crisis in the Enterprise Data Center

According to Moore's Law, the processing power of computer chips doubles about every 18 months. This pattern of innovation has been driving information technology forward for decades. Unfortunately, this progress has a downside, one that was less apparent when energy was cheap and components drew much less power than they do today. As processor performance has increased, so too has power consumption and its side effect, heat. Smaller servers packed with more-powerful components not only consume more electricity in their processing operations, they also generate more heat, requiring aggressive use of internal fans and air-conditioning that add a huge overhead to power consumption.

According to Gartner, "During the next three years, most CIOs will experience constraints in data center floor space and power that could limit an IT organization's ability to grow as the business grows.[1]"

As processing technology becomes smaller, faster and hotter, corporate data centers are consuming more electricity than ever before. Coupled with rising energy costs and a finite power grid, the result is an energy crisis affecting both individual companies as well as the power utilities that supply them. Recent industry studies have examined the surging demand for power from various perspectives, and the figures are startling:

- Space that's leased for $12 to $20 per square foot can cost $60 per square foot to cool[2]
- In 2005, power consumption for servers approached 3,800 Watts per square foot (equivalent to 38 standard light bulbs) for the most compact systems, up from 250 Watts per sqare foot in 1992 (a 1,400% increase) according to the American Society of Heating, Refrigerating & Air-Conditioning Engineers[3]
- Server [x86] power consumption has risen by as much as 300% in six years, increasing from 100-200 Watts per server in 2000 to 300-400 Watts per server in 2006[4]
- The [x86] server installed base has been growing at an approximately 11% to 15% compound annual growth rate during the past four years, and this rate will only accelerate[5]
- U.S. data centers are consuming 1.5% of total electricity consumption, a figure predicted to double in the next five years[6]
- By 2009 technology operations in the U.S. will spend twice as much for power and cooling as they did to buy the server hardware in their data centers[7]

Increasingly, power utilities are informing corporate customers that they can no longer provide the electricity needed to run their latest processing-intensive applications, a development that would have been unheard of only a few years ago. Companies then find themselves at a power crossroads: either scale back growth plans, or pay the astronomical cost of building an additional facility or relocating the data center to a region where power is less restricted.

Wasted Real Estate

There's another, rather ironic, dimension to the power crunch: inefficient use of data center space. As phenomenal processing power in smaller form factors outpaces the ability of server vendors to keep their equipment cool, workarounds are needed that leave extra space for air to circulate around miniature components. Rather than filling racks to capacity, the customer is required to skip rows or slots to allow heat to dissipate, giving up valuable real estate on the data center floor. Figure 1 illustrates the relationship between available floor space and rack density[8].

Figure 1 Effect of average rack density specification on the fraction of available space for IT racks

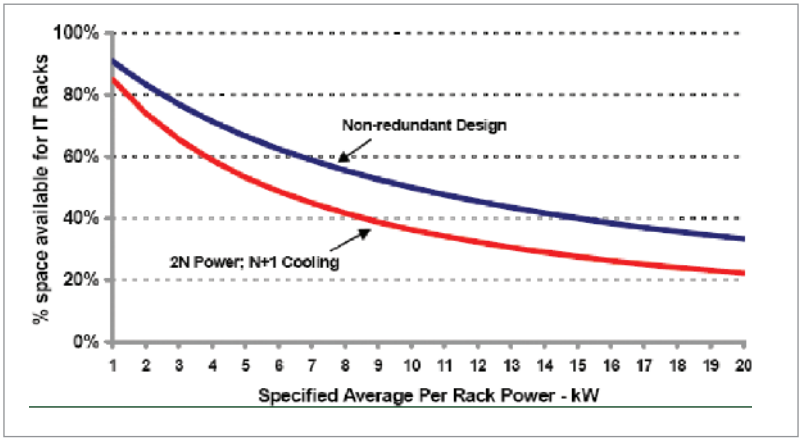

Hardware manufacturers use complex configuration tools that limit the number of blades that can be installed in a rack, and reserve much-needed space for cooling. Most hardware manufacturers also look for new and innovative technologies to cool their over-heated systems, but these new products cost more money, require additional power and, naturally, occupy more space. In a Gartner survey published in February 2007, 35% of respondents expect their data center racks to hold an average of seven to 10 KW each in the next two years, and 23% expect even higher.

Smaller…Faster…Hotter: Chip Technology on General-Purpose Servers

Data center servers used for general-purpose computing (as well as personal computers) are dominated by x86 architectures. This chip technology is at the heart of blade servers and rack servers, including those used in other data warehouse systems. These chips also contain features designed to improve performance for general-purpose computing, such as large cache memory to store results for upcoming operations.

But is this the right approach for data warehousing? A design that is ideal for processing ten records related to a credit card transaction may be woefully inefficient when churning through ten billion records to discover trends and patterns in customer and operational data. With billions of rows of data shuttling between storage and memory, and with a cache far too small to be effective on such a massive scale, the CPU is under a tremendous strain — with a corresponding increase in power requirements and heat.

Figure 2 shows a typical 3.4 GHz processor chip found on blades and rack servers. On top of the chip, a massive heat sink draws heat away from the chip like a car radiator. Fans pull air across the heat sink fins, which require their own electrical power source in addition to the power required to run the CPU. Even with this cooling apparatus, additional space is needed for air to circulate between the blades, driving up space requirements along with power consumption.

Figure 2 Typical x86 processor

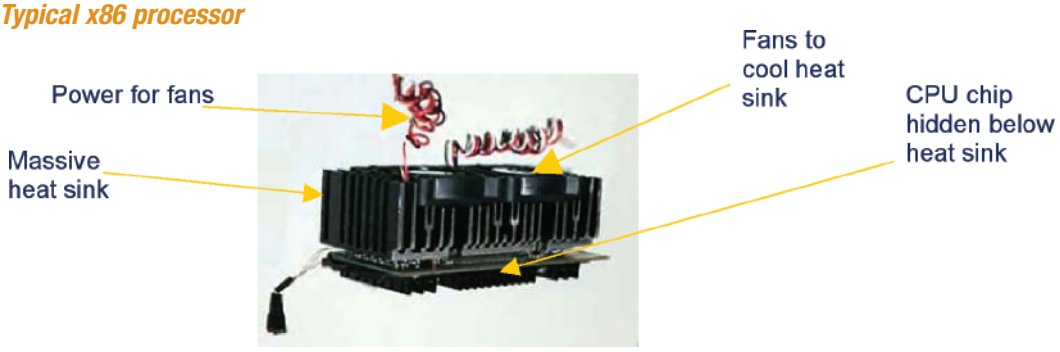

CPU chips based on x86 architectures are forcing more data through smaller components at faster speeds than ever before, magnifying previously minor power and cooling issues. Chip technology has evolved in several ways, including:

- Narrower gates (from 130 nm to 45 nm)
- Faster clock speeds (from 200 KHz to 3.4 GHz)
- More transistors (from 3,500 to 55,000,000)

…All on an area not much larger than a postage stamp. With so much processing activity condensed into a very small space, the chip is unable to dissipate the heat by simply relying on ambient air. As a result, chip manufacturers are forced to build in elaborate cooling measures like those shown above — even for the new "energy efficient" 45 nm chip technology. As manufacturers figure out how to deal with excessive heat at the chip level, it's important to remember that heat is still trapped inside the rack, and must be dissipated by system fans and finally by the data center's HVAC system.

Today, traditional data warehouse infrastructures incorporate the same architecture flaw, which is based on the movement of billions of records of data from storage to memory. These systems rely on high-power processors, volumes of memory, indexing, replications and other methods that consume storage space, money and time to analyze piles of irrelevant data in pursuit of insights. Dell, Hewlett-Packard, IBM, Teradata and other data warehouse vendors have all committed to these processing architectures in their data warehouse systems and they are all bound to the technology roadmap of chip manufacturers and its inherent penalty of power consumption, heat production and space requirements.

Now the industry is moving to multi-core processing — dual and quad-core and soon octal-core — creating multiple independent versions of the same CPU functionality in a single chip. It's another performance milestone, but also exacerbates the power/cooling issues mentioned above, since the chip is doing twice (or more work) in the same space for a corresponding increase in power consumption and heat.

Green and Greener

As power consumption becomes increasingly problematic, the industry is responding from various directions. A quick search of the term returns reams of information about ventilation systems, power management software and consulting services aimed at making the data center more energy-efficient. Chip manufacturers are announcing high-performance, low-power chips that deliver twice the performance but draw 30 percent less power. The new UltraSPARC® T1 from Sun Microsystems draws only about 70 Watts, a big improvement over earlier generations of multi-core devices. While these figures may seem impressive, they still fall short of addressing the problem.

The reason is that the new x86 chips rely on other on-board devices and chips to handle the flood of data between the CPU, memory and I/O, driving up power consumption substantially. The net result is that data warehouse systems relying on this technology still tend to be extremely power-hungry with individual processing nodes requiring nearly 100 Watts of power. The high power requirements of these systems are further extended when accounting for their high-performance RAM memories and separate, multi-terabyte disk storage arrays.

Over the next few years, the IT industry will be battling to bring power consumption under control. But for data warehousing, much greater energy efficiency, along with system performance, has already arrived. Netezza's architecture, rather than expensive components, makes the difference.

The Netezza Alternative: Low Power, High Performance

Netezza's patented Asymmetric Massively Parallel Processing™ (AMPP™) architecture takes a different approach to processing massive queries than architectures developed for general-purpose computing. This not only provides proven performance gains, it also dramatically lowers power consumption, heat and space requirements.

The Netezza architecture is unique in delivering streaming analytic processing, leveraging high-end embedded technology used primarily in areas such as video games, industrial process controls and medical imaging, and adapted by Netezza for data warehousing and advanced analytics. The approach uses commodity FPGAs (field programmable gate arrays) to do the bulk of the filtering, with an embedded PowerPC chip handling the remainder. The innovative architecture, rather than expensive, power-hungry components, is what provides the performance difference, along with substantial energy savings.

Unlike traditional systems that shuttle data between disk and memory for processing, a Netezza appliance streams data off the disk and through query logic loaded into the FPGA. By performing initial filtering as data streams off the disk, the FPGA reduces the workload on the PowerPC chip by approximately 90 percent. Unlike general-purpose chips, no cache is involved; none is needed. Since the PowerPC is not expected to handle gigantic workloads, it needs only 5 Watts of power, in contrast with conventional high-performance, "power-efficient" chips that require 70 Watts or more.

The FPGA and PowerPC chip, together with 400 GB of disk storage, reside on each of the massively parallel intelligent storage nodes called snippet processing units (SPUs). An entire SPU consumes less than 30 Watts of power, eliminating the need to skip rows, conserving space while holding power costs to a minimum.

Figure 3 shows the overall power requirements for three of the primary new "low-power" CPU device vendors (IBM, AMD and Intel), including the surrounding devices for I/O discussed previously. As shown in the chart, these lower-power CPU devices consume far more power than the 5 Watts required by the embedded PowerPC processor in the NPS system's SPU. Furthermore, each of the other vendors' CPU alternatives alone require much more power than a fully-populated Netezza SPU intelligent storage node.

Figure 3 I/O power cost on new x86 chips

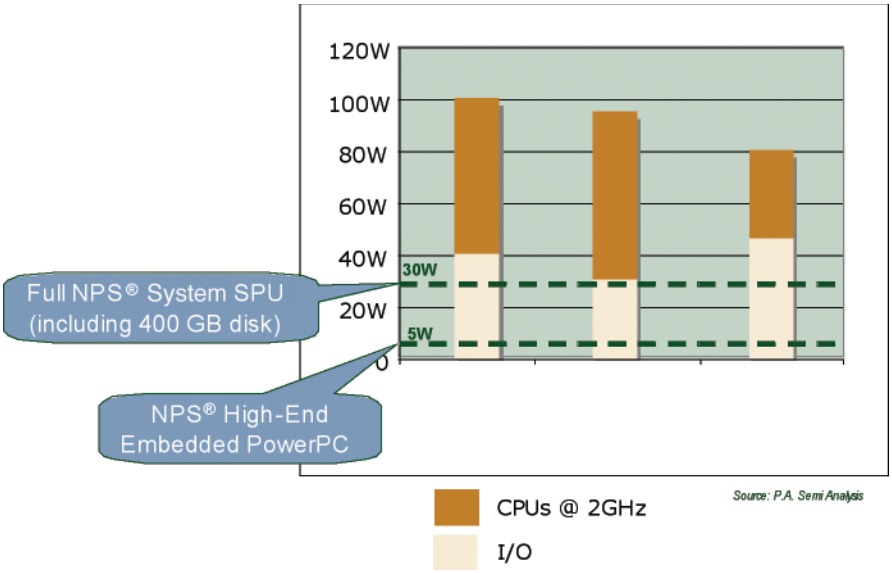

A Linux host server in the NPS appliance aggregates SPU results and manages the query workload. The host is the only element in an NPS system where an x86-based processor is used today. Power and cooling requirements for a full one-rack system, including the host and 112 SPUs (with 12.5 TB of user data capacity) are only 3,500 Watts together with only 12,000 BTU/hour of cooling. Compare this with two widely used similarly sized configurations:

- Teradata 5500H with 2X dual Intel Xeon processors and NCR 6843 storage array requires five racks of equipment, 18,500 Watts to power and produces 63,000 BTU/hour that needs to be ventilated and cooled.
- IBM BCU AIX on P5-575 servers requires four racks of equipment, 29,700 Watts to power and produces 101,400 BTU/hour that needs to be ventilated and cooled.

Power Math: You Make the Call

So when it comes to power savings, what's the bottom line? The following tables compare power and cooling on two NPS system models with well-known systems of equivalent processing power and storage capacity. Figures are based on APC's online calculator that provides power and cooling ratings for major blade and rack server configurations (www.apc.com/tools/ups_selector/index.cfm). Netezza encourages companies to use the APC calculator to make their own independent assessment of power requirements to validate all vendor claims.

- To use the online calculator, simply click "Configure by Devices," click "Server" and select the desired model from the drop-down list. Now click "Add to Configuration." The calculator displays the power and heat ratings for the server selected.
- To view the storage ratings, select "Add Another Device," click "Storage Array," and select the desired device.
- The power and BTU values in the tables below are based on multiplying the server and storage ratings by the total number of devices.

Performance in a Small Footprint

Orange UK, the UK's most popular mobile phone service, chose the NPS system to analyze billions of Call Detail Records (CDRs) in a fraction of the time required by other systems. The need for in-depth analysis about the success of targeted products and services is increasingly important within the telecom industry, and companies like Orange can leverage the NPS system to ensure that CDR data can be turned into actionable information.

Even as the amount of data stored in Orange's data center continues to grow, it has managed to reduce its equipment footprint. The new infrastructure in the data center has seen the number of cabinets spaces drop from 26 to nine. With complexity and floor space now seen as two of the biggest costs facing IT departments, Netezza has helped Orange to prepare itself for future growth.

As a Senior Data Center engineer at Orange explains, "Space and power are always at a premium within our data centers. The NPS solution saved large amounts of both while providing increased performance. The speed and ease of deployment was an added bonus."

Data Size = 12.5 TB	Netezza	MPP Solution	SMP Solution	Packaged Appliance Solution	Cluster Solution
# of Data Center Tiles	1.5 Tiles	6 Tiles	8.5 Tiles	15 Tiles	10 Tiles
Power - Watts	3,500	18,500	15,100	26,800	15,900
Cooling - BTU/hr	11,900	63,000	51,700	91,300	54,300
Cooling Watts [2]	3,600	18,900	15,500	27,400	16,300
TOTAL WATTS	7,100	37,400	30,600	54,200	32,200
NETEZZA POWER EFFICIENCY (%)		81%	77%	87%	78%

Data Size = 25 TB	Netezza	MPP Solution	SMP Solution	Packaged Appliance Solution	Cluster Solution
# of Data Center Tiles	3 Tiles	12 Tiles	11.5 Tiles	20 Tiles	14.5 Tiles
Power - Watts	7,500	38,000	28,600	32,000	30,600
Cooling - BTU/hr	25,600	129,900	97,600	109,200	104,500
Cooling Watts [2]	7,700	39,000	29,300	32,800	31,300
TOTAL WATTS	15,200	77,000	57,900	64,700	61,900
NETEZZA POWER EFFICIENCY (%)		80%	74%	77%	75%

Conclusion: The Energy-Efficient Data Warehouse

Netezza has earned worldwide recognition for its breakthrough architecture that processes terabytes of data in record time. Along with unmatched performance, Netezza's analytic appliance provides an extremely efficient processing approach that requires only a fraction of the power used by traditional architectures. Rather than struggle with growing power and cooling constraints, more and more organizations are discovering the impact that the high-speed, low-power NPS system can have on their business. Netezza customers reap the benefits: huge power savings, minimal space requirements and lower cost of ownership — along with performance that provides a competitive advantage.

[1] *A Message from Data Center Managers to CIOs: Floor Space, Power and Cooling Will Limit Our Growth* – Gartner, August, 2006

[2] *Power Surge* – Darrell Dunn, Information Week, February 27, 2006

[3] *Power-Hungry Computers Put Data Centers in Bind* – Don Clark, Wall Street Journal – November 14, 2005

[4] *The Impact of Power and Cooling on Data Center Infrastructure* – An IDC presentation by John Humphreys and Jed Scaramella, May, 2006

[5] *Data Center Power and Cooling Scenario Through 2015* – Rakesh Kumar, Gartner, March 14, 2007

[6] *From the EPA's report to Congress on August 2, 2007* – Presented by Andrew Fanara

[7] *IDC's IT Forum & Expo, June 14, 2006. Data Center News* – Bridget Botelho, SearchDataCenter.com, February 14, 2007

[8] *Guidelines for Specification of Data Center Power Density* – Neil Rasmussen, White Paper #120, APC

[9] *2006 Data Center Polling Results: Power and Cooling* – Michael A. Bell, Gartner, February, 2007

About Netezza

Netezza (NYSE Arca: NZ) is the global leader in analytic appliances that dramatically simplify high-performance analytics for business users across the extended enterprise, delivering significant competitive and operational advantage in today's information-intensive marketplaces. The Netezza Performance Server® (NPS®) family of streaming analytic™ appliances brings appliance simplicity to a broad range of complex data warehouse and analytic challenges. Customers who are realizing the benefits of Netezza appliances include Ahold, Amazon.com, CNET Networks, Debenhams, Department of Veterans Affairs, Epsilon, Neiman Marcus, Orange UK, Premier, Inc., Ross Stores, Ryder System, Inc., The Carphone Warehouse, the US Army and Virgin Media. Based in Framingham, Mass., Netezza has offices in Washington, DC, the United Kingdom and Asia Pacific. **For more information about Netezza, please visit www.netezza.com.**

Netezza Corporation : 200 Crossing Boulevard : Framingham, MA : 01702-4480
+1 508 665 6800 tel : +1 508 665 6811 fax : www.netezza.com

NETEZZA
The Power to Question Everything™