white paper

From the Grid to the CPU: Minimizing Operational Expenses in Today's Data Centers

Rackable Systems, Inc.

The operational challenges of today's data centers have become a gloomy issue for IT managers and CIOs around the globe. As hundreds of thousands of servers and storage systems are deployed to keep pace with ever-growing business computing needs, data center managers are faced with increasing operational costs—namely power, cooling and maintenance. Today's data centers spend just 30% of their budget on capital expenditures, while a whopping 70% is consumed by operational expenses.

The natural shift away from antiquated RISC/proprietary Unix-based servers and the movement toward high-performance, open architecture, open source solutions has enabled countless data centers to proportionally reduce capital expenditures while increasing overall system performance. But while the price/performance ratio has improved, the cost to operate and maintain all those servers has spiraled out of control.

This white paper examines ways in which today's data centers can improve overall efficiency and reduce operational waste by leveraging unique server and data center-level technologies.

POWER CONSUMPTION

Today's X86 server designs have enabled dramatic improvements in density, with rack-mount systems packing more, ever-faster processors into a cabinet. Advancements such as the half-depth form factor and back-to-back mounting enable density levels as high as 80 1U rack-mount servers per cabinet. The newest multi-core processors from AMD and Intel are enabling never-before-seen CPU core densities, and the introduction of quad-core computing has increased compute levels to an unprecedented 640 processing cores in a cabinet just 7 feet tall and 3 1⁄2 feet deep.

Density levels that high with all those processing cores in such a small footprint is a necessity for data centers managing mission-critical applications, resulting in exponentially greater compute power in a small amount of space. But higher density comes with a price—a very high price in the form of power, cooling and maintenance. Thus, while delivering the flexibility and scalability that leading Internet and High Performance Computing (HPC) environments are clamoring for, dense server clusters raise power consumption levels and amplify thermal management challenges.

With data centers restricted by how much power can physically be delivered to the site, it is not uncommon for a data center to simply run out of power capacity long before running out of space. Similarly, data centers’ HVAC systems are limited in how many BTUs of cooling power can be delivered. In public co-location facilities, a fixed, predetermined level of wattage and heat dissipation is assigned per rack, making this an even greater challenge for companies hosting their servers in a public data center.

Given that data centers have among the highest density of energy-consuming equipment of any modern building and use 100 times the electricity of a typical office building on a per square foot basis, the issue of consumption continues to make headlines.

It is estim … 70 watts
of power … row
after row … costs,
and the

With the … vicious
cycle of

This document requires Adobe Reader version 6.0.1 or later. To view this document, please download Adobe Reader from: http://www.adobe.com/

INSIDE T

One of t … arefully
specify … stem
power d
exhausts

At a bas … tras
in the fo
overspec … e and
are cust … to the
end use

For exar
requiren … the
same 200-watt need. That would result in poor power efficiency given that typical power supplies function most efficiently when driven at a high percentage of their maximum rating. While a difference of tens of watts per server is not significant in small installations, in large deployments it drives up power consumption considerably, as well as the amount of heat that must be dissipated.

Once the proper ‘recipe’ for a specific server configuration has been identified to match performance and power requirements, the next step is to consider ways to reduce power consumption by choosing high efficiency, lower wattage components. This becomes especially important when selecting processors and power supplies.

CPUs and AC power supplies draw more power than any other components inside a system. Subsequently, they give off significant amounts of heat and are generally the first parts to fail inside a server. The processor drives a server's overall power draw requirement; the fastest processors consume the most power and generate the most heat.

Today's lower wattage CPUs from AMD and Intel continue to achieve high levels of performance with decreased power consumption. Higher CPU clock speeds will require higher wattage draws, which may be necessary in HPC and other compute-intensive environments. Regardless of what CPU is selected, the system should fit within the data center's so-called 'energy footprint' and systems should be designed to be cooled effectively.

Power su ... e.
Tradition ... ltages
required ... ency,
meaning ... d to
heat) wi ... ween
30% and

This document requires Adobe Reader version 6.0.1 or later. To view this document, please download Adobe Reader from: http://www.adobe.com/

And in t
mere 10 ... by
contribu ... em's
internal ... wer
efficienc

Recent ... of
90% or ... nds
of dollar ... , hard
dollar sa ... equent
increase ... stment
in the m ... e
server.

But wha ... ent,
reliable

DC POWER AS AN ALTERNATIVE

In 2003, the first distributed DC power technologies became available for large-scale data center server deployments. Long a standard in the telco industry, DC power options provide a more efficient power standard, eliminating the need for wasteful AC-DC power conversion inside a traditional AC power supply at the system level.

By replacing a standard AC power supply with a 93% efficient DC power card inside each server, power losses at the server level are instantly reduced. Instead of converting AC power to DC power inside the server, the conversion happens at the cabinet-level via redundant AC-to-DC rectifiers. By converting the power within the external rectifiers—which have an MTBF of 250,000 hours a piece—20-40% less heat is dissipated within the servers themselves. The subsequent efficiency gain results in power savings while increasing overall system reliability.

In this rectified DC scenario, AC power is brought to the server cabinet through standard power distribution mechanisms. Rectifiers within the cabinet, which take up a mere 2U of rack space in a 44U cabinet populated with 80 systems mounted back-to-back, distribute redundant DC power to each system via DC bus bars housed inside the cabinet. Power savings of 10% and higher are achieved with rectified DC solutions, which can amount to millions of dollars in savings each year alone for a large-scale Web farm.

While the power savings achieved using DC solutions are certainly compelling, the real operational
advantag … r cards,
coupled … uced
mainten …
ensures … ified
DC solut … ucture
changes …

This document requires Adobe Reader version 6.0.1 or later. To view this document, please download Adobe Reader from: http://www.adobe.com/

REDUCI…

While ca … DC
power ca … dous
amount … ep to
the time …

In a typi … s from
the pow… he grid
through … 94%
efficient … supply
inside a … ith
upwards … sing that
today's …

New tec… l
power efficiency to levels as high as 79% for the data center as a whole. By leveraging large-scale rectifiers and exterior DC power converter units, AC-to-DC conversion is reduced to just two points on the electrical path, culminating in data center-wide DC power distribution directly to DC-based servers.

The increasing popularity of easily deployed rectified DC solutions and the even greater efficiencies achieved via data-center level DC power distribution have prompted IT leaders to consider DC as the new data center standard. New, 'greenfield' data centers are increasingly being built with DC input power—in lieu of AC—in order to achieve the highest possible efficiency and lowest possible power loss from the doorstep to the servers.

BETTER COOLING STRATEGIES

Reducing power consumption at both the server and data center levels clearly impacts operational expenditures. Reduced power consumption means lower power bills and less heat at the component level. Less heat in the chassis means greater system reliability and lifespan. And greater reliability means less human capital and money wasted on maintenance and repair.

But after all that attention on lower wattage components and better power distribution techniques, heat will naturally be a by-product of any machine, and remains the number one killer of systems in a data center. Effective heat evacuation will help reduce the burden on HVAC infrastructure while increasing server life expectancy.

This document requires Adobe Reader version 6.0.1 or later. To view this document, please download Adobe Reader from: http://www.adobe.com/

Tradition … -mount
systems … at back
into the … ture to
adequat … hese
tradition … reasing
the amb …

With coo … an
electric … ficant
portion … the
quest to …

One bes … inside
the cabi … s and
connect … n—or
chimney … the
cabinet … nates
both the … the
data cer … .

Some X8 … connect
directly … cooling
framewo … louvers,
heat can also be directed downward—toward the floor—and evacuated through in-floor cooling devices. Some data centers leveraging these unique heat evacuation techniques actually re-circulate the carefully contained hot air to warm office space during the winter months.

PUTTING IT ALL TOGETHER

The data center, the site of computing innovation for four decades, is continually challenged by ever-increasing operational costs driven by business demands for increased compute performance. With the rapid expansion of rack-mount, X86 systems, many data center managers are struggling to survive within their existing power envelope while keeping data center labor costs to a minimum.

Facilities built as little as two years ago are ill-equipped to hold today's high-density server configurations, largely due to inherent limitations in their power distribution and cooling infrastructure that very visibly impact the bottom line in any budget.

As reducing power consumption, lowering heat output and increasing system reliability become key factors for IT purchasing decisions, every data center needs to consider ways to optimize existing infrastructure and streamline operations:

Consider the power draw of the system. Do my servers contain any unnecessary, extraneous components (all of which, of course, draw power and expel heat)? Has my server provider selected low wattage components that fit within my energy footprint? Do I have enough power per square foot in my ... opulated data cen...

Choose ... most efficient ... the processo...

This document requires Adobe Reader version 6.0.1 or later. To view this document, please download Adobe Reader from: http://www.adobe.com/

Calculat ... rectified ... iency losses a...

Don't fo ... at cabinet- ... ollutes my data ...

Finally, ... ating systems ... inefficie ... the cost of h...

The ope ... er. As energy c ... s on eliminat ... ges of today's i ... spread adoption in the world's leading data centers, but careful consideration must be given to the growing post-purchase challenges as the sheer volume of deployments continues to rise. Data centers squarely focused on solutions that reduce power, cooling and maintenance costs will reap the benefits time and time again.