NETAPP VISION PAPER

10 TECHNIQUES NETAPP USES FOR IMPROVING DATA CENTER POWER EFFICIENCY

February 2009 | WP-7030-0209

IMPROVE POWER EFFICIENCY IN YOUR DATA CENTER USING TODAY'S TECHNOLOGIES

NetApp is committed to being an industry leader in energy conservation. We live this commitment by building highly efficient systems and consistently squeezing power inefficiencies out of our own data centers. In this paper, we'll reveal 10 techniques our team of experts used to design a data center that already exceeds EPA power efficiency recommendations for 2011. You can use these same techniques today to start reaping the benefits of a power-efficient data center.

TABLE OF CONTENTS

POWER EFFICIENCY TARGETS

In August of 2007, the U.S. Environmental Protection Agency (EPA) made a report to Congress highlighting the need for concerted efforts to reduce power consumption in data centers. The EPA estimates that if current efficiency trends continue, national energy consumption by data centers will nearly double to more than 100 billion kWh by 2011.

Today's data centers account for approximately 7 gigawatts (GW) of peak load on the power grid. That's equivalent to the output of about 15 base-load power plants. If current trends continue, this demand will rise to 12GW by 2011, which will require the building of an additional 10 power plants just to power data center growth.

In its 2007 report, the EPA constructed a set of scenarios that set objectives for improving the energy efficiency of data centers by 2011:

- **The state-of-the-art scenario** exploits the maximum efficiency potential for today's technology to reduce electricity use by up to 55% compared to current trends.
- **The best-practice scenario** uses today's technologies to reduce electricity use by up to 45%.
- **The improved-operational-management scenario** uses only low-cost measures to reduce electricity use by 20% relative to current trends.

The Green Grid suggests that the Power Usage Effectiveness (PUE) rating be used as a measure of power efficiency throughout the data center. The PUE is expressed as a decimal number. A total facility power demand to IT power demand ratio of 2:1 is equivalent to a PUE of 2.0, which is the current benchmark. The lower the PUE, the more efficient your IT infrastructure is in relation to your IT equipment.

The electricity savings of each scenario expressed as a PUE target appears in the table below.

Table 1) PUE targets for each EPA scenario in 2011.

Scenario	2011 Target PUE
Maximum achievable	1.20
State-of-the-art	1.45
Best-practice	1.55
Improved-operational-management	1.80

CHALLENGING OLD ASSUMPTIONS

Over the years, our rapid business expansion and corresponding growth in data have placed continuous pressure on our IT and data-center infrastructures. We've found that working together makes it possible to stay ahead of unbridled data growth.

Through the efforts of a multifunction NetApp team (and, in California, with the efforts of our utility, PG&E) we have been able to eliminate inefficiencies throughout our IT systems and facilities to accommodate our growth and extend the lives of our data centers. Our experience is a useful example of how real savings can happen when you involve a multidisciplinary team and take a holistic approach. Our experience also illustrates how big energy savings can result from challenging old assumptions.

ASSUMPTION #1: FACILITIES MANAGEMENT AND IT STRATEGY ARE DISCRETE FUNCTIONS

At NetApp, we discovered years ago that teaming up results in more efficient systems. Having each group aware of the needs and challenges of others makes it possible for facilities to quickly respond to changing growth demands from IT. At the same time, IT is aware of the challenges that face facilities—floor space, power, and cooling—and can balance these concerns against the growth needs of the business.

ASSUMPTION #2: WE'VE ALWAYS DONE IT THIS WAY

Raised floors in data centers are artifacts of mainframe days. Our cold aisle/hot aisle layout made a raised floor unnecessary, which enabled us to put our power and cooling delivery overhead. And our newer data centers look more like a series of enclosed tunnels than one large room. This design provides the ultimate physical barrier between hot and cold air.

ASSUMPTION #3: HIGH-TECH PROBLEMS REQUIRE NEW TECHNOLOGY TO SOLVE

Power efficiency doesn't always require expensive high-tech solutions. Consider the vinyl curtain. You can hardly get more low-tech. By creating a physical barrier between hot and cold aisles, vinyl curtains save us one million kWh of energy per year in our Sunnyvale data center alone.

ASSUMPTION #4: DATA CENTERS HAVE TO BE REALLY COLD

Over the years, we have gradually moved up the temperature in the data center from a starting point of 52°F (11.1°C). As we increase the temperature, we save in multiple ways: We need less cooling, we can get more hours of free cooling using outside air, and we can increase the water temperature in our chillers to maximize chiller efficiency.

What truly sets our data centers apart from other, less-power-efficient data centers is in the details. No two data centers are built alike, and it is unlikely that every technique we use will be applicable to your situation. However, we hope that this paper will inspire you to work collaboratively and develop ways to continuously squeeze inefficiency out of your data center.

NETAPP IT AT A GLANCE

# of users:	~7,999 at 113 worldwide locations
# of data centers:	6: Americas (3), Europe, India, Hong Kong
# of servers:	1,500
Storage:	2.5PB+
Applications:	Over 100 applications; six tier 1 applications
Densities:	Old data centers: From 4 kW/rack to 8kW/rack
	New data centers: 12 kW/rack

In the NetApp® data center located in Sunnyvale, California, a state with some of the highest electricity rates in the country. Power costs account for 65% of our annual data center budget, so maximizing power efficiency does have an impact on our bottom line.

10 TECHNIQUES FOR IMPROVING POWER EFFICIENCY

Through teamwork and continuous improvement, we have fine-tuned our data-center systems to optimize power efficiency. To date, we have already achieved the energy efficiency savings that the EPA is projecting for state-of-the-art, power-efficient data centers in 2011. We expect that our recently opened data centers will exceed the EPA's electricity reduction targets for the 2011 state-of-the-art scenario with PUEs of 1.20. These are the 10 techniques we use to maximize efficiency in our own data centers.

1 MEASURE TO CONTROL

An old maxim of operational efficiency is: "You can't control what you can't measure." We've discovered that efforts to reduce power inefficiencies need to begin with baseline measurements. Because utility bills don't tell you where your power is going, you need a way to measure the power consumption of various components of your data center. We measure power consumption for IT systems, UPS, chillers, and lighting.

Table 2) Power consumption components and measuring techniques.

Component	How to Measure
IT Systems	We use simulations of typical workloads to identify power requirements for the servers and storage in our labs. Power to the busway feeding the racks is monitored continuously so we can quickly identify problems and balance power loads. To identify hot spots, we install and continuously monitor temperature sensors at multiple locations around each rack.
UPS	We use a flywheel UPS system instead of a battery-based system. The flywheel UPS is highly energy efficient, in part because it eliminates one of the AC/DC/AC conversions typical in a data center. The UPS power in and out is monitored by our energy management system.
Chillers	Chiller power is measured in combination with power from the tower fans. As the chiller energy goes down, the power required to operate the tower fans goes up. Chiller flow, temperatures, and tonnage are monitored by our energy management system.
Lighting	We use occupancy sensors to automatically turn lights on and off as people enter and exit the room.

2 VIRTUALIZE AND CONSOLIDATE IT SYSTEMS

The EPA estimates that servers and storage account for 50% of all power usage. These areas are logical places to look for power savings. Server virtualization is a hot trend right now—and for good reason. Server virtualization can significantly extend the life of a data center by producing notable savings in space, power, and cooling. To get the full benefits of server virtualization you need a storage infrastructure that provides pooled networked storage. The same economics apply to storage virtualization: Fewer, larger storage systems provide more capacity and better utilization, resulting in less space, power, and cooling.

In the process of implementing storage and server virtualization, we have moved to more energy-efficient storage systems. We replaced 50 older storage systems with 10 of the latest storage systems running Data ONTAP® 7G. By upgrading to the latest technology, we were able to:

- Reduce our storage rack footprint from 25 to 6 racks.
- Lower our power requirements from 329kW to 69kW.
- Reduce our air conditioning capacity requirements by 94 tons.
- Cut our electricity costs to power those systems by $60,000 per year.

Server and storage consolidation and virtualization can be a mixed blessing. Many data centers discovered the hard way that high-density servers and storage fix one problem only to create another. Many high-density data centers require 25kW to power a single rack. We design the capacity in our existing facilities to support 4kW to 8kW per rack, and our new facilities will support 12kW per rack. With high-density systems requiring two to three times the power that data centers were designed to provide, cooling can be a critical issue. Even where power may be available, it may not be possible to cool the high-density racks with the existing cooling infrastructure. That is just another reason why power-efficient cooling is so important.

Note: A major server virtualization project is currently underway to help provide additional savings in utilization, floor space, power, and cooling. Stay tuned for those results in future publications.

3 MANAGE DATA

While planning storage virtualization, we performed an analysis of our existing business data. We discovered that 50% of the data we were storing could be eliminated.

The only way to prevent runaway data growth is to stop data proliferation before it can clog up the works. The average UNIX® or Windows® enterprise disk volume contains thousands or even millions of duplicate data objects. As these objects are modified, distributed, backed up, and archived, the duplicate data objects are stored repeatedly. The end result of this is inefficient use of storage resources and wasted energy to power them. We use many of our own technologies to reduce our raw storage requirements and eliminate duplicate data on our systems. Snapshot™, FlexClone®, deduplication, FlexVol®, and RAID-DP® are just a few of our technologies that have a huge impact on our storage footprint.

4 ELIMINATE OVERCOOLING OF SYSTEMS

In the data center, power and cooling go hand in hand. Traditionally, every watt brought into the data center to power equipment requires a second watt of power for cooling. Overcooling equipment is not an efficient use of energy. Most data centers cool equipment based on manufacturers' power-load recommendations. Because manufacturers typically base their power consumption estimates on running peak loads all the time—a condition that is rarely met—most of the time, data centers overcool their equipment.

Calculating accurate power loads can be tricky. Generally, busses in the data center are shared by servers and storage, making it hard to separate server and storage power requirements. To arrive at reasonable power-load estimates for our specific circumstances, we test equipment in our lab environment before deploying it in our data center. We have discovered that reasonable power-load estimates for our circumstances are 30% to 40% lower than manufacturer estimates. Once we deploy our systems, we monitor rack-by-rack power usage and balance the phases as needed. By constantly tuning our cooling systems based on our specific experience, we cut down the amount of energy that would be wasted by overcooling our systems.

We also save energy by using variable frequency drives on our air handlers. Instead of running fans at 100% speed, variable frequency drives vary the speed of the fans depending on what's actually needed to cool the IT equipment on a row-by-row basis. Sensors constantly monitor temperatures, and the air handler fan speeds are automatically adjusted to increase fan speeds in hotter areas and reduce them in cool areas. The energy saved by reducing fan speed is dramatic: a 50% reduction in fan speed yields and an 87% reduction in power consumption.

Note: This strategy can't be implemented in raised-floor environments in which floor tiles must be manually removed to increase cool air delivery to targeted locations.

5 WORK WITH PHYSICS IN DATA CENTER LAYOUT

We design our data centers to take advantage of a key principle of physics: hot air rises and cold air falls. Sending cooling up from the floor, the way older data centers with raised floors typically do, requires extra energy. So, instead of providing airflow from the floor, we drop a curtain of cold air down the front of the machines. The cold air is drawn into the machines and exits out the back as hot air, which naturally rises to the ceiling, where it is vented outside. It's common sense not to vent the hot exhaust of one machine into the air intake of another. To keep one machine's exhaust from heating another machine, we place our racks front to front and back to back. This arrangement, called "hot aisle/cold aisle," has become a best practice for data center design because of its energy efficiency.

By putting our cooling delivery above the racks, we eliminate the need for a raised floor, plus we gain other efficiency advantages. Modular power distribution busses positioned above the racks allow our data-center administrators to quickly provision power for new systems without outside help and without having to shut down the main power source. With this system, a data-center administrator can provision power to a new system in about five minutes.

Older data centers typically have raised floors to accommodate water pipes that carry cold water used to cool mainframes. Traditionally, raised floors were also used for cables and acted as a plenum to move cold air. As many modern-data-center personnel have discovered, raised floors can significantly limit the data center's ability to increase cooling capacity to support today's higher-density systems. To make more space for air to flow, you must either raise air pressure in the plenum or raise the floor. However, all too often, raising the floor of an existing data center is not a viable option.

6 CONTINUOUSLY IMPROVE HEAT CONTAINMENT

Wherever you have high-density racks in a hot aisle/cold aisle arrangement, you need additional airflow measures to prevent hot exhaust from getting into the cold aisle. In Sunnyvale, we accomplish this in two ways. First, we draw up the hot air and exhaust it at the top of the racks. Then, we further isolate the hot air from the cold aisles using low-cost, low-tech techniques. Blanking panels on unfilled racks block cold air from being wastefully drawn through empty racks. To further isolate the heat, we use vinyl curtains—similar to the ones you see in meat lockers—at the ends of the hot aisles and around the cooling outtake system above the racks.

Whether to curtain the hot aisles or the cold aisles is a hotly debated issue. We choose to use vinyl-strip curtains to contain the air in hot aisles and the same vinyl to create a physical barrier around ducts and equipment. The vinyl melts at 120°F (48.8°C), allowing the curtains to melt and drop off before sprinklers are triggered so they do not get in the way of fire suppression. We estimate that the curtains in Sunnyvale alone will save us one million kWh per year.

In our Research Triangle Park (RTP) labs in North Carolina we employ a variety of other heat-containment methods that yield PUEs from 1.41 to 1.20. It is worth noting that we are unable to bring in outside air for economizer cooling in RTP Building One, so all PUEs for this building are achieved without the benefit of economizer cooling.

Note: We are able to achieve such good PUEs at our RTP site because labs do not require the same redundancy and availability as a typical data center.

Table 3) Summary of PUEs in RTP labs.

Cooling Type	Equipment kW/rack	PUE
Existing Labs Without Outside Air		
VAV Damper	4	1.41
Overhead High Density	8	1.41
Pressurized Cold Room	8	1.37
New Lab with Outside Air		
Pressurized Cold Room with Outside Air	12	1.20

Overhead VAV Damper

Our second- and third-floor labs in Building One house 650 racks with a 4kW-per-rack power capacity. We use the hot aisle temperature to control our overhead control damper. Our variable air volume (VAV) device uses a variable frequency drive (VFD), resulting in lower fan speeds that have a high power-savings impact. With these adjustments, we are able to achieve a PUE of 1.41 in our Building One second- and third-floor labs.

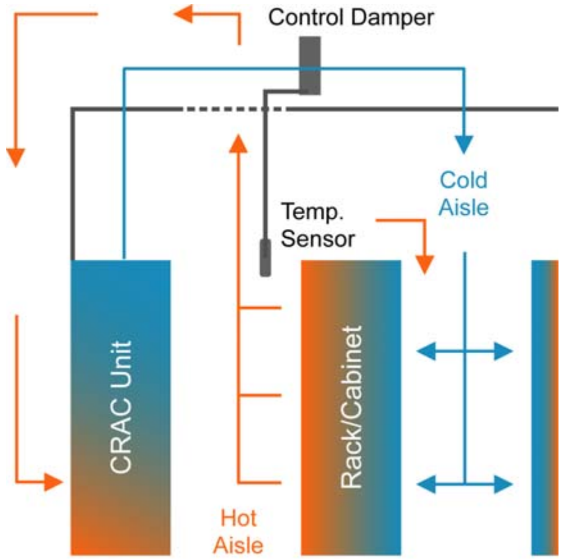

Figure 1) VAV system.

Overhead High-Density Cooling

In our Kilo-Client lab, power capacity is 8kW per rack. Our overhead high-density cooling system uses refrigerant for coolant and low-horsepower fan units. This system, which requires a separate dehumidifier, also yields a PUE of 1.41.

Pressurized Cold Room

Our first-floor lab in Building One houses 360 racks with an 8kW-per-rack power capacity. Instead of cold aisles, pressure-controlled cold rooms are enclosed by doors at both ends. Although this solution does not allow us to use outside air for free cooling, it dramatically reduces hot and cold air mixing and results in lower fan speeds (which significantly reduces power consumption), maximizes cooling capacity, and virtually eliminates hot spots. With this setup, we can get a PUE of 1.37, and, as a nice side benefit, the lower fan speeds make the environment about 10 times quieter.

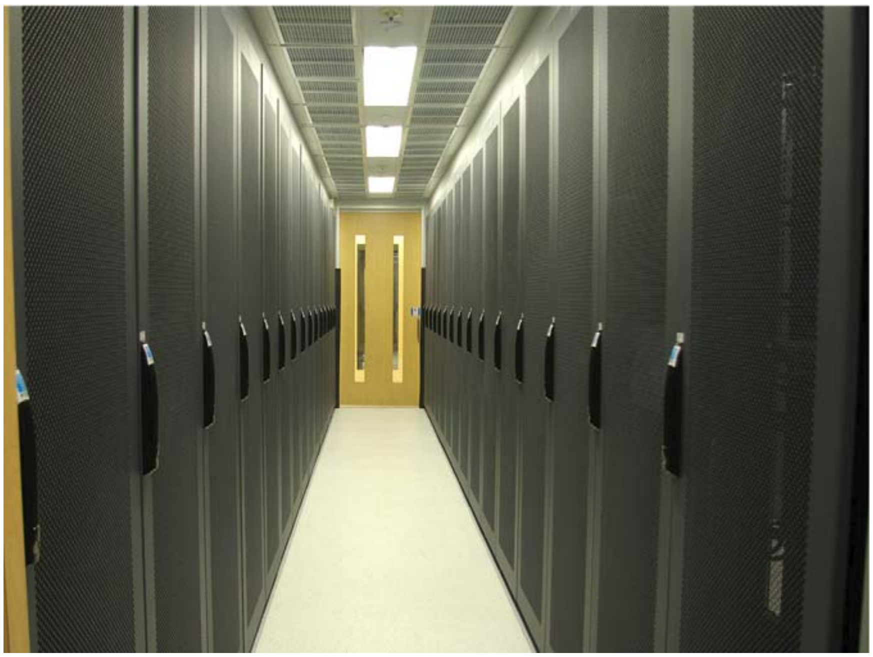

Figure 2) Cold room tunnel.

7 MAXIMIZE FREE COOLING

Our Sunnyvale data center was designed to use outside air for cooling. Dampers built into the side of the building automatically modulate the outside air coming in. When the temperature outside is lower than the established set point, the dampers automatically open and outside air flows through the cooling system. As the outside air heats above the temperature set point, the dampers close and our chillers kick in to cool the equipment. Today, free cooling saves us $1.5 million per year in energy costs. In the future, all our new facilities will include air economizers to use outside air for free cooling and provide even greater savings.

Our environmental engineers are constantly working to raise the temperature set point for which we can use air from the outside to cool equipment inside. We started with a set point of 52°F (11.1°C) and have gradually moved it up. Now, with a set point of 70°F (21.1°C), Sunnyvale weather provides free or mixed cooling 68% of the year. We are in the process of raising the set point to 75°F (23.8°C), which could raise our free cooling hours to 85% of the year. We expect that by the time we finish, we will be able to use free or partial free cooling 97% of the time.

When we raise data-center temperatures, not only can we take advantage of a wider range of outside temperatures for free cooling, the water we use to cool the air doesn't have to be as cold. We can raise the chiller water temperature by 10 degrees, giving us a 17% increase in chiller efficiency and a 14% increase in chiller capacity. Increased chiller capacity yields an infrastructure savings in addition to our ongoing power savings.

Raising the set point has been a balancing act with the systems that have variable internal speed fans. When the set point is too high, equipment fans kick in and consume the energy saved in raising the set point. Discovering the optimal set point—one that triggers neither the computer room air conditioning units nor the fans in the systems—requires fine tuning.

If you choose to employ this tactic, you need to remember that raising the temperature set point reduces your margin of error and the amount of time to critical temperature overload should there be a failure. Be sure to balance cost savings against the risk of such an overload.

8 MINIMIZE ELECTRICAL CONVERSION LOSSES

Traditional data centers use lead acid batteries to store power for use in the event of a grid failure. Battery UPS systems convert from AC to DC and back to AC.

Our Sunnyvale data center currently features kinetic flywheel UPSs that store energy as motion. Energy comes in from our switching infrastructure and spins an electric motor in each of the UPS units. Flywheels store energy to produce 15 to 20 seconds of energy—just enough to carry out any of our switching operations. Older battery UPSs are 85% efficient, while today's best UPSs average roughly 94%. Flywheel UPSs, with an efficiency rating of 97.7%, lose less than half the energy that batteries do.

9 USE HEAT THAT WOULD OTHERWISE BE WASTED

Power demand and energy prices in California soar with the temperature. As the temperature and price of electricity peak, our natural gas–powered cogeneration system goes online to economically power our one-megawatt data center. Generating electricity close to its use, as we do, is known as distributed generation (DG). DG can lower power costs and reduce the amount of electricity lost in transmission.

To add to the cost benefits of DG, we also employ cogeneration—a thermodynamically efficient use of fuel. Electricity production often generates large quantities of wasted heat; cogeneration puts this energy to work. Cogeneration in our Sunnyvale data center starts with the electricity produced by natural gas–powered generators. The heat that is created in the process is used to power an adsorption chiller that chills the water used in the cooling system. The cogeneration system has an overall efficiency rating of 75% to 85%. Using our cogeneration system saves $300,000 annually.

To always deliver power to the data enter at the lowest possible cost, we steadily monitor the price of electricity and the cost of natural gas to determine if and when to run our natural gas generators. As you can see from Figure 3, our cogeneration system significantly reduces the amount of power we get from the grid.

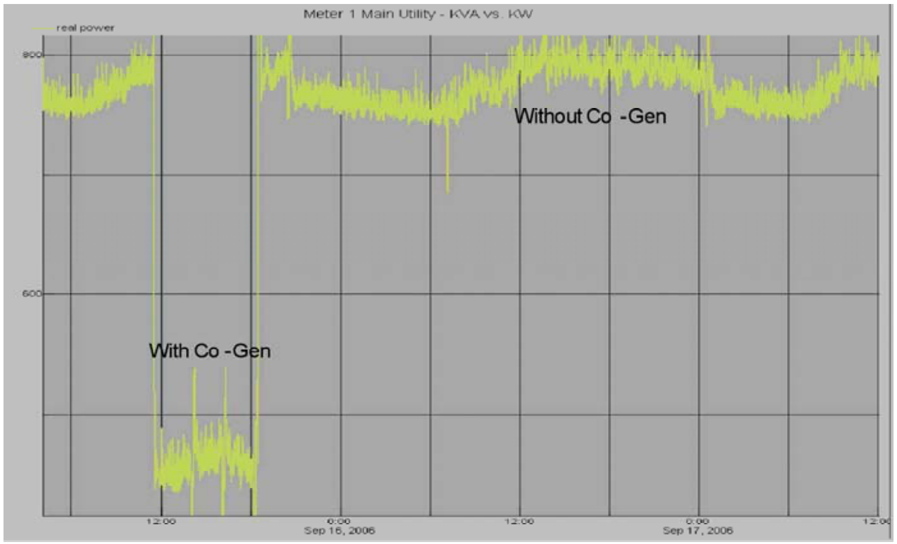

Figure 3) Power usage with and without cogeneration.

10 CONSTANTLY MONITOR AND TUNE

Legacy measurement of watts per square foot in a data center may indicate that your room is running fine on average, while, in reality, you have hot spots throughout the room that are damaging equipment and going unnoticed. Most data centers measure load at the perimeter, which, predictably, makes things unpredictable. It is important to take metering one step further. To truly enable energy efficiency, you should measure your data center at the rack level (watts per rack) instead of using the traditional approach of cranking up the fans when a particular area in the data center starts to run hot. We constantly test and tune our environment. We use vinyl curtains to eliminate heat syncs, and we have multiple temperature sensors at the mid-level that register an average differential of 10 to 12 degrees.

TAKE IT TO THE LIMIT

Our recently constructed, dedicated lab building in North Carolina and our new lab in Sunnyvale went live with projected PUEs of 1.20. We've reached this "maximum achievable" PUE three years ahead of the EPA's target by applying all we have learned about data-center power efficiency. Here's what we're doing to minimize inefficiency and reach our goal.

CREATING COLD ROOMS

We've taken to the limit what we learned about the effectiveness of creating barriers between hot and cold air. Instead of having curtains separate hot and cold aisles, our hot and cold aisles are actual rooms that resemble hallways or tunnels. Each cold room houses 30 to 60 equipment racks.

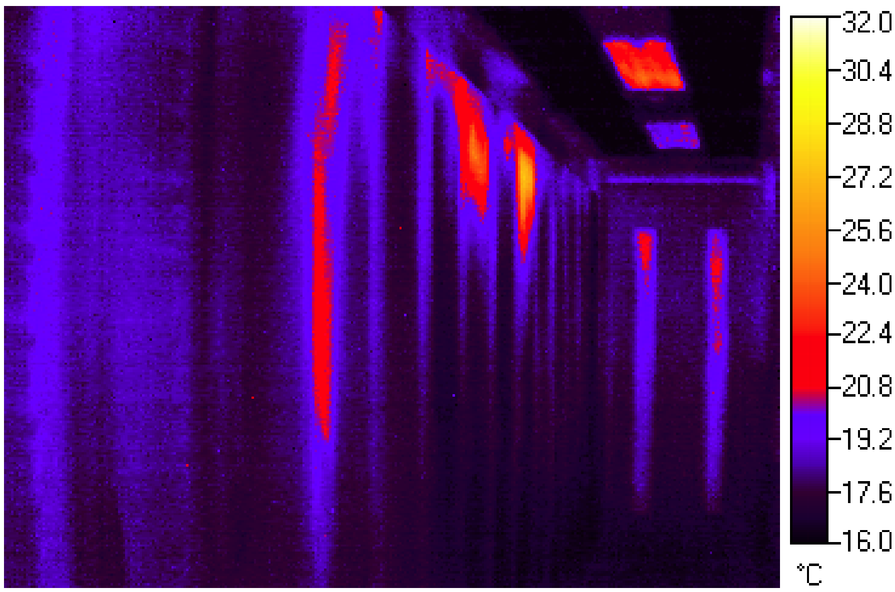

Figure 4) Cold room.

A cold room virtually eliminates the short-cycling of air around the cabinets. The net result: Less fan horsepower is required for full utilization of the installed equipment capacity.

PRESSURIZING THE COLD ROOMS

We measure the pressure difference between the cold room and the hot aisle at the back of equipment racks and control the amount of air we deliver into the cold room to match the volume of air flowing through the racked equipment. This lowers the fan speeds in the equipment.

ELEVATING THE SUPPLY AIR TEMPERATURE

We increase the supply air temperature delivered to the cold room to 74°F (23.3°C). At this temperature, return air from the equipment approaches 95°F (35°C).

USING OUTSIDE AIR FOR FREE COOLING

In our new labs, we monitor both humidity and temperature and calculate our use of outside air based on which is more efficient. We expect to average free cooling 65% to 75% of the year and partial free cooling a full 97% of the time annually. Doing so will allow our mechanical refrigeration plant to work less intensely when it is needed. As an added benefit, should our mechanical refrigeration plant fail, we can use outside air for emergency cooling.

MOVING FANS CLOSER TO COMPUTER LOADS

By moving the fan systems closer to the actual equipment load, we are able to reduce energy losses caused as air flows through the ducts.

MONITORING OCCUPANCY

We estimate that our labs are unoccupied about 90% of the time. We monitor occupancy so that as soon as someone walks into an area, the system detects his or her presence. Lights come on and the temperature is lowered to make working conditions more pleasant. When the person leaves the area, the temperature goes back up and the lights go out.

CONCLUSION

How you implement a power-efficient data center depends on your data-center equipment and conditions. For example, when elevating a temperature set point, you need to consider the temperature tolerances for the cabling that supports your IT equipment. Implementing power efficiency in the data center doesn't require expensive, high-tech solutions. Something as simple and inexpensive as vinyl curtains can help you save energy and start reaping the benefits of a power-efficient data center today.

ABOUT NETAPP

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover our passion for helping companies around the world go further, faster at www.netapp.com.