An Improved Architecture for High-Efficiency, High-Density Data Centers

By Neil Rasmussen

White Paper #126

Executive Summary

Data center power and cooling infrastructure worldwide wastes more than 60,000,000 megawatt-hours per year of electricity that does no useful work powering IT equipment. This represents an enormous financial burden on industry, and is a significant public policy environmental issue. This paper describes the principles of a new, commercially available data center architecture that can be implemented today to dramatically improve the electrical efficiency of data centers.

Introduction

In a typical data center, less than half the electricity used actually makes it to the computer loads. More than half the electrical bill goes to the purchase of power consumed by the electrical power system, the cooling system, and lighting. The total electrical consumption therefore has two principal contributors – (1) the power consumed by the IT loads, and (2) the power consumed by the support equipment (**Figure 1**). The focus of this paper is the power consumed by the support equipment, which includes the *losses* of the power path equipment plus *all* power used by non-power-path support equipment.

Figure 1 – *Power consumption in the data center*
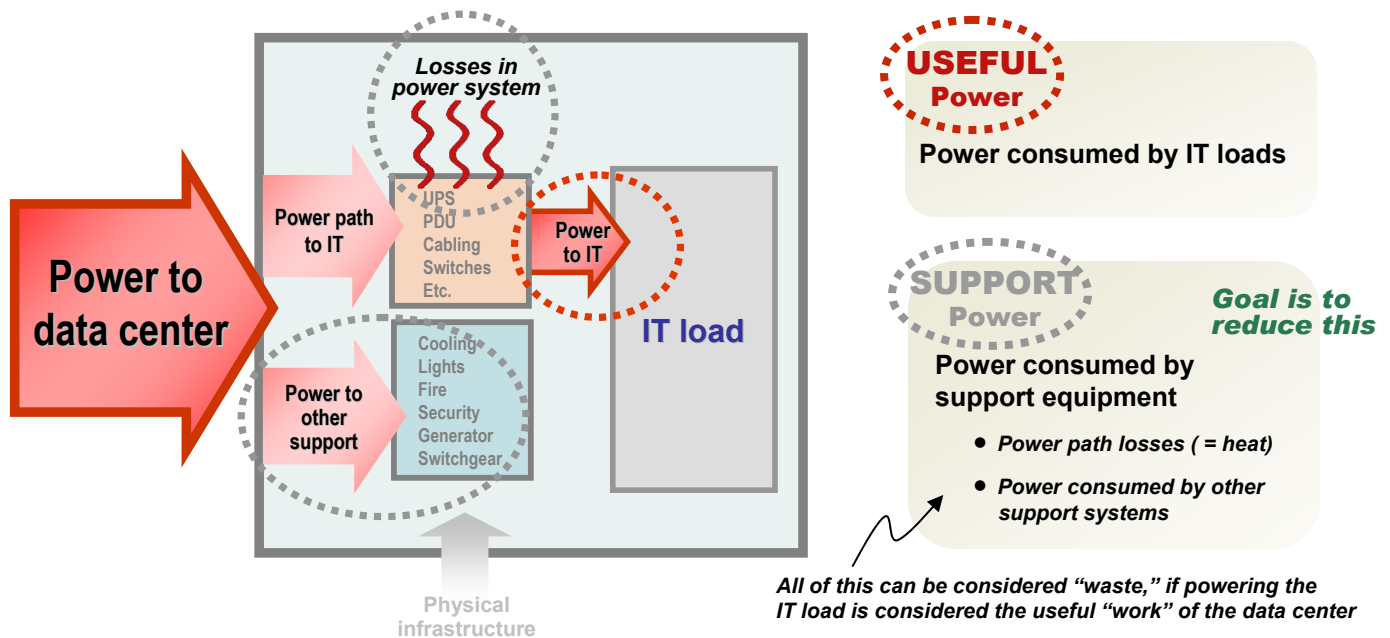
Vendors of computer equipment are providing new solutions such as virtualization that have the potential to reduce the total amount of IT equipment required to perform a specific function, which offers a means to reduce IT-load power consumption. Unfortunately, at the same time, the trend of IT systems operating at higher densities with time-varying power draw are driving *down* the electrical efficiency of the data center power and cooling systems (see box).

Various proposals to address the waste of power in data centers by improving the performance of power and cooling systems have been described in the literature. Some of these, such as direct water pipe connections to IT devices and DC power distribution, promise incremental improvements to

---

How do high density and varying IT Load reduce data center efficiency?

High density and dynamic loading actually offer an opportunity for **increased** efficiency, if supported by "smart" row-based power and cooling. However, without properly redesigned power and cooling – a common reality – the typical result can be

- Cooling waste due to increased room-wide cooling to cover hot spots

- Reduced operating loads and over-capacity of power and cooling – this pulls efficiency down because lighter load means lower efficiency of power and cooling systems

The loss of efficiency due to excess or misdirected power and cooling is addressed later in this paper.

---

system efficiency but are impractical today. This paper introduces an improved data center architecture – available and practical *today* – that reduces the energy consumption of the power and cooling systems by more than 50% in a typical installation.

The new architecture described in this paper is not simply a physical configuration of equipment or improved efficiency of individual devices – it is a whole-system makeover that combines the best elements of data center design:

- Engineering design of individual devices

- Power distribution

- Inter-component communication and coordination

- Cooling strategy

- System planning

- Management tools

When all of these elements are combined as an integrated system, the performance improvement can be dramatic.

Where Does All the Power Go?

The energy flow through a typical 2N data center is shown in **Figure 2**. Power enters the data center as electrical energy, and virtually all power (99.99%+) leaves the data center as heat. (The rest is converted to computing by the IT equipment.)

Figure 2 – Power flow in a typical 2N data center
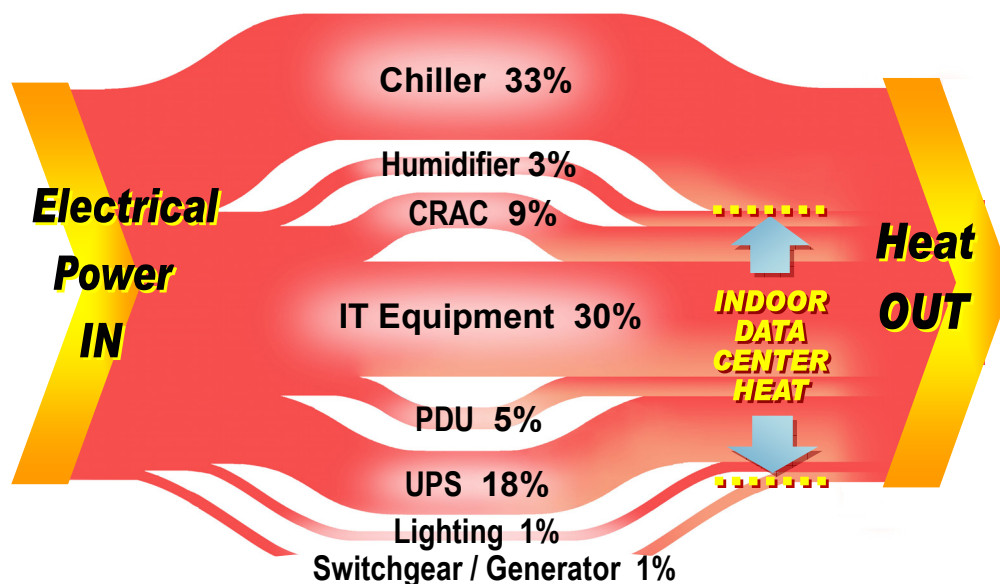
Note that in this example only about 30% of the electrical power entering the facility actually powers the IT load (called **USEFUL power** in **Figure 1**), and the rest is consumed – converted to heat – by power, cooling, and lighting equipment. (An insignificant amount of power goes to the fire protection and physical security systems, and is not shown in this breakdown). This data center is said to be 30% efficient, because 30% of total input power goes to the IT load. Therefore, 70% of the input power is not doing the "useful work" of the data center (powering the IT loads) and is therefore considered to be data center *inefficiency* (or "waste," in the terminology of an efficiency model). To understand how we can dramatically reduce this inefficiency – remember that ALL non-power-path support is considered inefficiency in this model – we need to understand the five key contributors, which are:

1   Inefficiencies of the power equipment

2   Inefficiencies of the cooling equipment

3   Power consumption of lighting

4   Over-sizing of the power and cooling systems

5   Inefficiencies due to configuration

Five contributors to electrical inefficiency

These all contribute to SUPPORT power in Figure 1

While most users understand that inefficiencies of the power, cooling, and lighting equipment are wasteful, **the other items on the list above actually dominate the inefficiencies** and are not well understood. Each of the above five contributors is analyzed in detail in APC White Paper #113, *Electrical Efficiency Modeling for Data Centers*, and their power consumption characteristics are summarized here:

1   Inefficiencies of the power equipment

Equipment such as UPS, transformers, transfer switches, and wiring all consume some power (manifested as heat) while performing their function. While such equipment may have name-plate efficiency ratings that sound impressive – 90% or higher – these efficiency values are misleading and cannot be used to calculate the power wasted in real installations. When equipment is doubled for redundancy, or when the equipment is operated well below its rated power, efficiency falls dramatically. Furthermore, **the heat generated by this "wasted" energy in power equipment must be cooled by the cooling system**, which causes the air conditioning system to use even *more* electrical power.

2   Inefficiencies of the cooling equipment

Equipment such as air handlers, chillers, cooling towers, condensers, pumps, and dry coolers consume some power while performing their cooling function (that is, some of their input power is dispersed as heat instead of contributing to the mechanical work of cooling). In fact, the inefficiency (waste heat) of cooling equipment typically greatly exceeds the inefficiency (waste heat) of power equipment. When cooling equipment is doubled for redundancy or when the equipment is operated well below its rated power, efficiency falls dramatically. Therefore, **an increase in the efficiency of the cooling equipment directly benefits overall system efficiency**.

3 Power consumption of lighting

Lighting consumes power and generates heat. The heat generated by lighting must be cooled by the cooling system, which causes the air conditioning system to consume correspondingly more electrical power, even if the outdoor temperature is cold. When lighting remains on when there are no personnel in the data center, or when unutilized areas of the data center are lit, useless electrical consumption results. Therefore, increases in the efficiency of the lighting, or controlling lighting to be present only when and where needed, materially benefits overall system efficiency.

4 Over-sizing

Over-sizing is one of the largest drivers of electrical waste, but is the most difficult for users to understand or assess. Over-sizing of power and cooling equipment occurs whenever the design value of the power and cooling system exceeds the IT load. This condition can occur from any combination of the following factors:

- The IT load was overestimated and the power and cooling systems were sized for too large a load

- The IT load is being deployed over time, but the power and cooling systems are sized for a future larger load

- The cooling system design is poor, requiring over-sizing of the cooling equipment in order to successfully cool the IT load
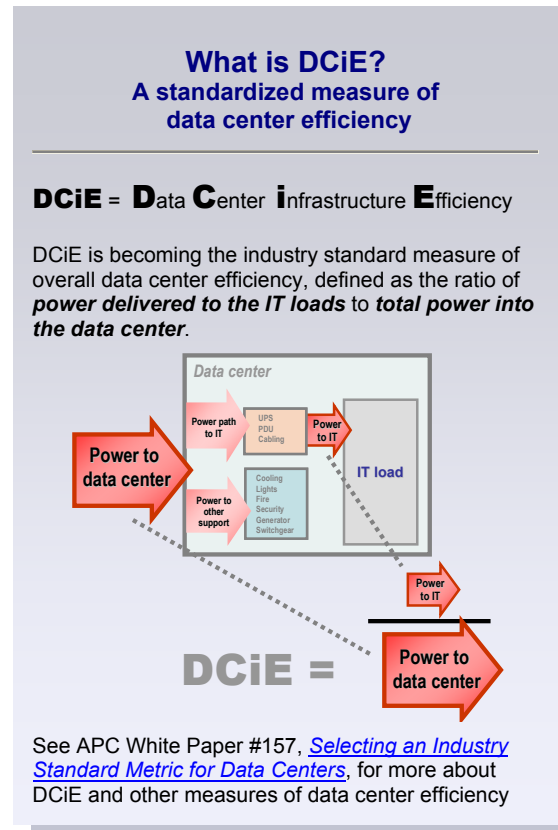
While it is clear that installing too much power and cooling equipment is wasteful from an investment standpoint, it is not obvious that such over-sizing can dramatically decrease the electrical efficiency of the overall system and cause excessive ongoing electrical consumption.

The fundamental reason that over-sizing of power and cooling equipment reduces the electrical efficiency of data centers is that the electrical efficiency of many powering and cooling devices declines dramatically at reduced load. While some electrical equipment such as wiring is *more* efficient at lower loads, most major equipment such as fans, pumps, transformers, and inverters exhibit decreased efficiency at lower loads (due to "fixed losses" that persist even when the IT load is zero). **This decline in efficiency is not readily determined from manufacturers' data sheets, which typically report efficiency at an optimum (usually high) load.**

For a detailed technical explanation of how the effects of over-sizing on electrical power consumption are quantified, see APC White Paper #113, *Electrical Efficiency Modeling for Data Centers*.

5 Inefficiencies due to configuration

The physical configuration of the IT equipment can have a dramatic effect on the energy consumption of the cooling system. A poor configuration forces the cooling system to move much more air than the IT equipment actually requires. A poor configuration also causes the cooling system to generate cooler air than the IT equipment actually requires. Furthermore, physical configuration may force various cooling units into a conflict where one is dehumidifying while another is humidifying, a typically undiagnosed condition that dramatically reduces efficiency. The current trend of increasing power density in new and existing data centers greatly amplifies these inefficiencies. These configuration problems are present in virtually all operating data centers today and cause needless energy waste. Therefore, an architecture that systematically optimizes the physical configuration can dramatically reduce energy consumption.

What is DCiE?
A standardized measure of data center efficiency

DCiE = **D**ata **C**enter **i**nfrastructure **E**fficiency

DCiE is becoming the industry standard measure of overall data center efficiency, defined as the ratio of *power delivered to the IT loads* to *total power into the data center*.

$$\text{DCiE} = \frac{\text{Power to IT}}{\text{Power to data center}}$$

See APC White Paper #157, *Selecting an Industry Standard Metric for Data Centers*, for more about DCiE and other measures of data center efficiency

An Optimized Data Center Architecture

The previous section describes the five principal contributors to the inefficiency of data centers. It is apparent from a review of these contributors that they are interrelated. Therefore, an effective approach to optimization must deal with the data center system as a whole – attempts to optimize the individual inefficiencies will be much less effective. A careful analysis of the contributors to electrical loss (inefficiency) leads to a finding that data center efficiencies can be substantially improved when an integrated system is developed based on the following principles:

- Power and cooling equipment that is not currently needed should not be energized

- Over-sizing should be reduced wherever possible, so equipment can operate within the optimum region of its efficiency curve

- Power, cooling, and lighting equipment should take advantage of the latest technologies to minimize power consumption

- Subsystems that must be used below their rated capacity (to support redundancy) should be optimized for that fractional-load efficiency, not for their full-load efficiency

DESIGN PRINCIPLES
of the new architecture

These guide the implementation of the system

- Capacity management tools should be used to minimize "stranded capacity" within the data center, allowing the maximum amount of IT equipment to be installed within the gross power and cooling envelope, pushing the system to the highest point on its efficiency curve[1]

- Optimized, integrated physical configuration should be inherent *within* the system, and not tied to the characteristics of the room where it resides — for example, row-based cooling should be integrated with the IT racks, independent of room-based cooling

- The system should be instrumented to identify and warn about conditions that generate sub-optimal electrical consumption, so that they can be quickly corrected

- The system should include installation and operation tools and rules that maximize operating efficiency and minimize or eliminate the possibility of sub-optimal configuration or installation

A commercially available integrated data center system using the above principles is shown in **Figure 3**.

Figure 3 – *High-efficiency integrated data center system*
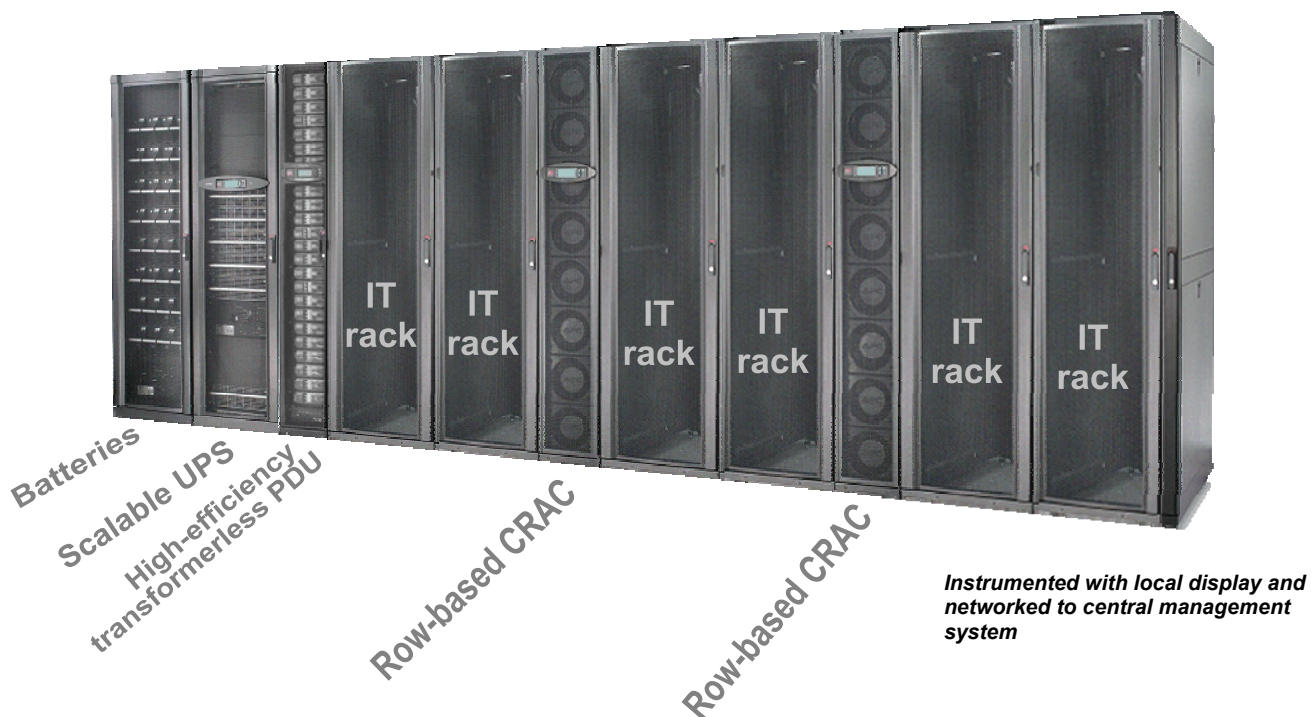
The data center system of **Figure 3** exhibits a 40% decrease in electrical consumption when compared with a traditional design, with the reductions in losses further classified as shown in **Figure 4**.

---

[1] For more about stranded capacity, see APC White Paper #150, *Power and Cooling Capacity Management for Data Centers*

Figure 4 – *Cost savings of improved architecture broken down by data center subsystem*
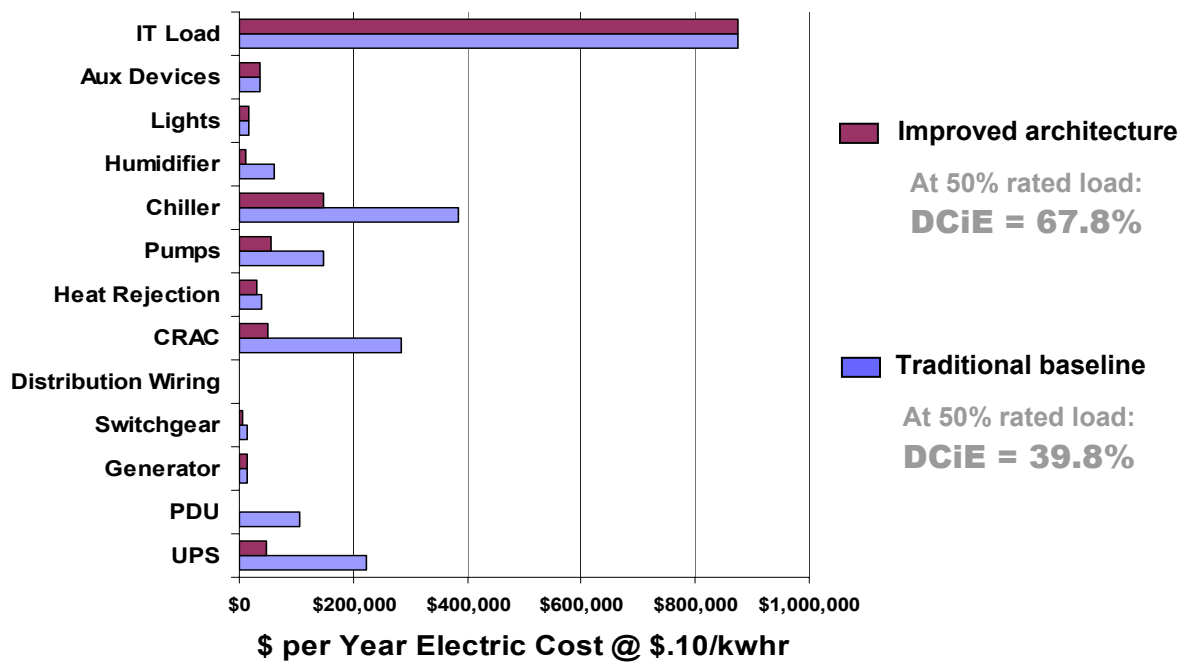
The efficiency gain of the improved system results in a dramatic reduction of electricity costs. For a 1MW IT load and a typical electrical cost of $0.10 per kW-hr, the savings in electrical costs would be approximately $9,000,000 over a ten-year period.

The above improvements are based on a data center with the following configuration:

- 2MW design capacity

- 1MW actual IT load

- All power and cooling infrastructure for 2MW installed and on-line

- Dual path power from service entrance to loads

- N+1 air handlers

- Chilled water system with cooling tower

- 7kW per rack average power density

- Hot-aisle/cold-aisle IT rack layout

- Efficiency curves for all devices from actual manufacturers' data

The power consumptions and savings are affected by these assumptions. For example, eliminating dual path power redundancy or N+1 air handlers would cause efficiencies to rise and savings to fall somewhat. The remainder of this paper examines these savings and the underlying assumptions in more detail.

Expressing the efficiency gain in terms of data center infrastructure efficiency (DCiE), the traditional data center design described above, operating at 50% of rated IT load, would have a DCiE efficiency value of 39.8% and the improved architecture would have a DCiE value of 67.8% under the same conditions.

Comparison to Conventional Approaches

The reduction of electrical losses (increased efficiency) described in the previous section is dramatic. Earlier in this paper we identified five key contributors to the inefficiency of conventional designs. How does the proposed improved architecture achieve these remarkable efficiency gains? What are the new technologies, designs, and techniques that contribute to these gains? And what data supports these gains? To answer these questions, we will take a closer look at the key elements that together give rise to the improvements of the new architecture, which are:

- **Scalable power and cooling**, to avoid over-sizing

- **Row-based cooling**, to improve cooling efficiency

- **High-efficiency UPS**, to improve power efficiency

- **415/240 V AC power distribution**, to improve power efficiency

- **Variable-speed drives on pumps and chillers**, to improve efficiency at partial load and on cool days

- **Capacity management tools**, to improve utilization of power, cooling, and rack capacity

- **Room layout tools**, to optimize room layout for cooling efficiency

CONSTRUCTION ELEMENTS
of the new architecture

Technology to implement the design principles

While some of these elements can be implemented alone to gain efficiency, it is important to understand that the *integration* of these elements into an overall architecture is responsible for a significant part of the gain. For example, while row-based cooling technology is fundamentally a much higher-efficiency technology than conventional room cooling, it is also a key enabler of the cost-effective implementation of room layout tools, capacity management tools, and scalable cooling.

The architecture described in this paper is effective in any data center in any geographic region. An additional strategy that improves data center efficiency is to take advantage of the cool outdoor temperatures available in some locations to improve the efficiency of the cooling system using "economizer cooling modes" or "free cooling." These approaches typically add to the data center cost and have a payback time that depends on outdoor temperature and humidity conditions. "Free cooling" technologies complement the approaches described in this paper – taking advantage of cool outdoor air to decrease the amount of electric power expended on cooling, which increases the efficiency of the data center. This paper does not consider free cooling in any of the efficiency or power savings calculations.

The quantitative efficiency contribution of each of the elements above is discussed in the following sections.

Scalable power and cooling

▶ **Avoids over-sizing**

All data centers have an efficiency that varies with the IT load. At lower IT loads the efficiency always declines and is equal to zero when there is no IT load. The shape of this curve is remarkably consistent across data centers. An example is shown in **Figure 5**.

Figure 5 – Data center efficiency as a function of IT load comparing modular vs. non-modular designs
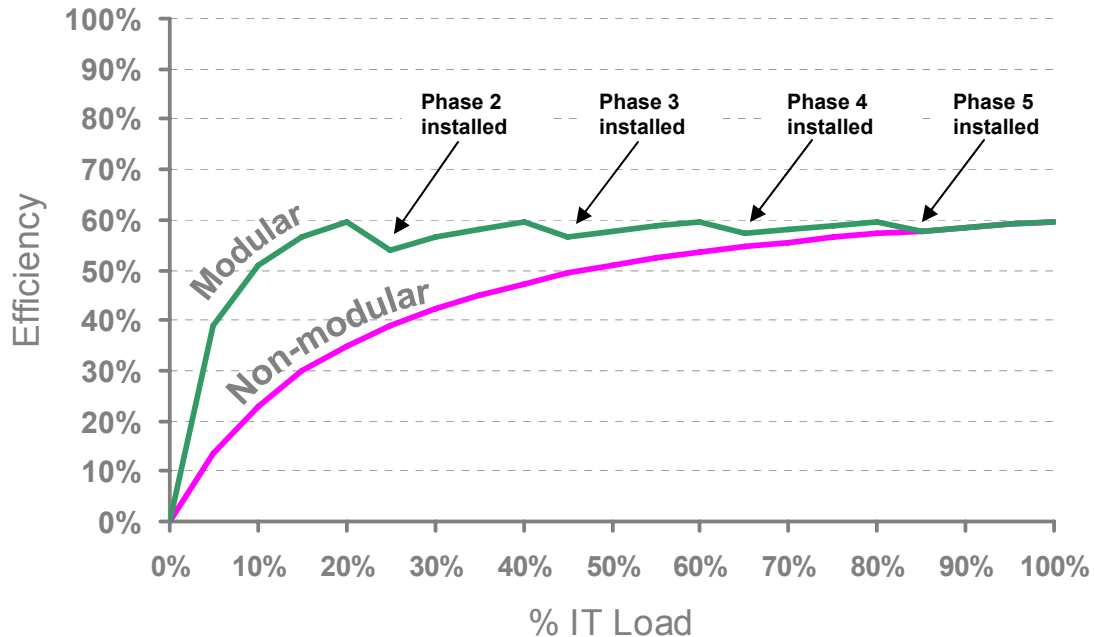
When the % IT load is well below the design value for the data center, the efficiency degrades and the data center is considered to be over-sized for that IT load. Many data centers operate in this condition – sometimes for years – typically because they are constructed for a hypothetical future IT load that has not yet been installed.

The reason why data center efficiency falls at light load is explained in detail in APC White Paper #113, *Electrical Efficiency Modeling for Data Centers*. It is similar to the reduction of fuel economy that an automobile experiences if it has a very large engine that is being used substantially below its power rating.

To correct the problem of reduced efficiency due to an over-sizing condition, the power and cooling equipment could be scaled over time to meet the IT load requirement. The upper curve in **Figure 5** shows what happens when power and cooling equipment is deployed in five phases instead of as a single system. At full load, the scalable power and cooling system has no efficiency advantage, but at lighter loads the efficiency is dramatically increased. At 20% load, the 1/5 deployed power and cooling system is running at full efficiency. This principle, as illustrated in **Figure 5**, is only partially attainable because some cooling infrastructure, such as coolant pumps, may be impractical to deploy in phases.

Many data centers are found to operate below the full IT load rating – especially smaller data centers or data centers early in their life cycle. The use of a scalable power and cooling solution can increase efficiency in these cases, as well as defer capital and operating costs until needed. In addition, some decisions such as the target power density for a future zone in a data center can be deferred until future IT deployment.

## Row-based cooling

▶ **Improves cooling efficiency**

Row-based cooling places air conditioning within the rows of IT equipment, rather than at the perimeter of the room. Shortening the air flow path reduces mixing of hot and cold air streams which improves the predictability of air distribution. Predictable air distribution to IT equipment allows for more precise control of variable airflow rates that automatically adjust to the needs of the nearby IT loads. Instead of wasting energy with constant speed fans, variable speed fans spin only as fast as required by IT loads. In addition, row-based cooling captures the hot air from the IT load while it is still hot, before it has a chance to mix with ambient air. Together, these effects dramatically improve the efficiency of the computer room air handler. The basic arrangement of row-based cooling is shown in **Figure 6**.

***Figure 6** – Cooling using row-based CRACs with shorter air flow paths*
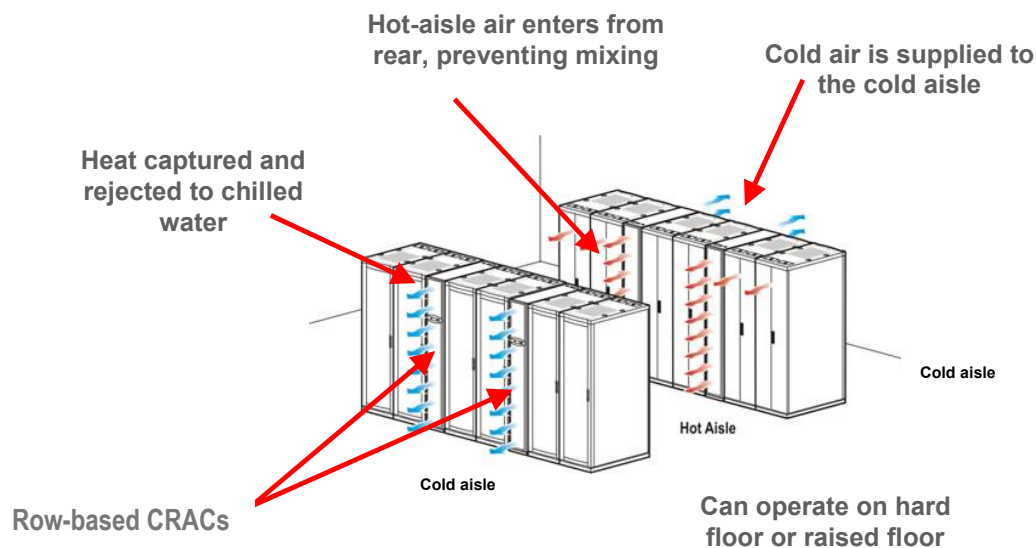
The increase in efficiency of row-based cooling architecture compared to traditional computer room air conditioners is shown in **Figure 7**.

Figure 7 – *Computer room air conditioner efficiency curves comparing row-based cooling to traditional room cooling*
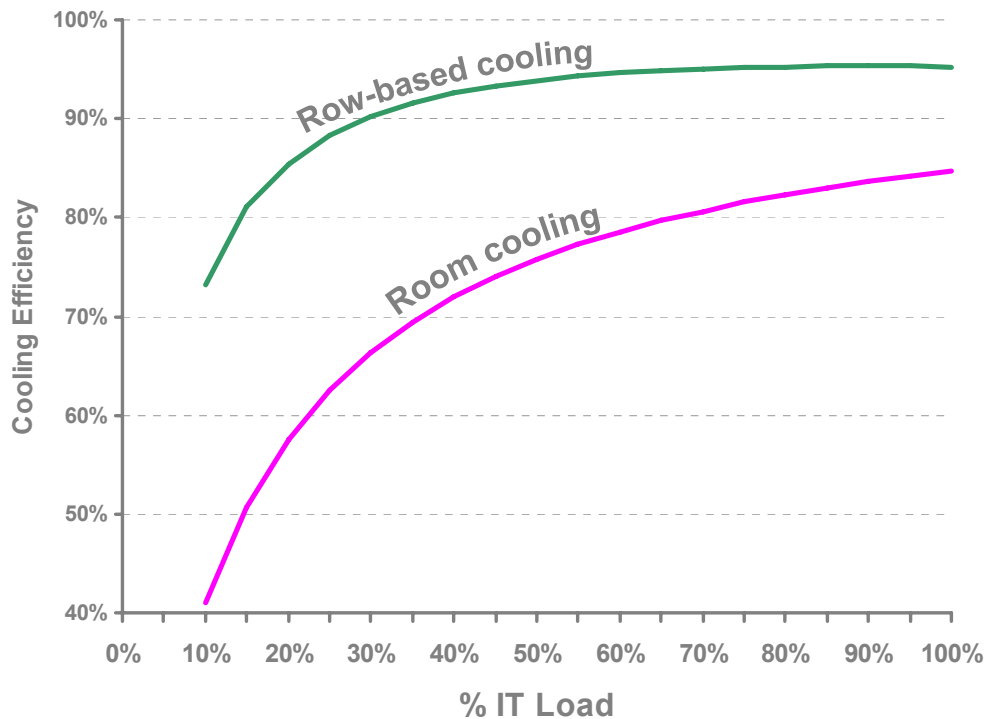
The curves represent the cooling efficiency expressed as the air conditioner output (heat processed) divided by the input (heat processed + electric consumption). This allows us to examine the performance of the computer room air conditioner using the typical 0-100% efficiency scale.[2] Ideally, the air conditioner would have an efficiency of 100%; the curve above shows that a typical CRAC has an efficiency of 80% at 70% IT load, which means that 20% of the input power is going to fan and humidification. By contrast, the row-based CRAC has an efficiency of 95% at 70% IT load, which means that only 5% of the input power is going to fan and humidification. This is a factor of four reduction in loss.

The above chart was developed assuming a chilled water design, tier 4 design, 3-foot raised floor for the room cooling solution, and an average power of 10kW per rack. This data is only for the CRAH unit and does not include the chiller, pumps, and cooling tower. These devices can be studied separately or combined with the CRAH to obtain the overall cooling efficiency of the data center. Note that the chiller and cooling tower do consume power and therefore do reduce the efficiency of the overall cooling system to values lower than shown in the figure.

---

[2] Computer room air handler efficiencies are typically expressed in other units, such as "coefficient of performance" or "IT load watt cooled per electrical watt." However, these are very difficult to relate to everyday experience and are not expressed in the common 0-100% form used to express efficiency for other types of equipment. All of the different methods are related mathematically and convey the same information.

High-efficiency UPS

➤ **Improves power efficiency**

Technologies are now available that substantially increase the efficiency obtainable by UPS systems. **Figure 8** compares efficiencies of a recently introduced high-efficiency UPS to UPS efficiency data published by Lawrence Berkley National Labs.[3]

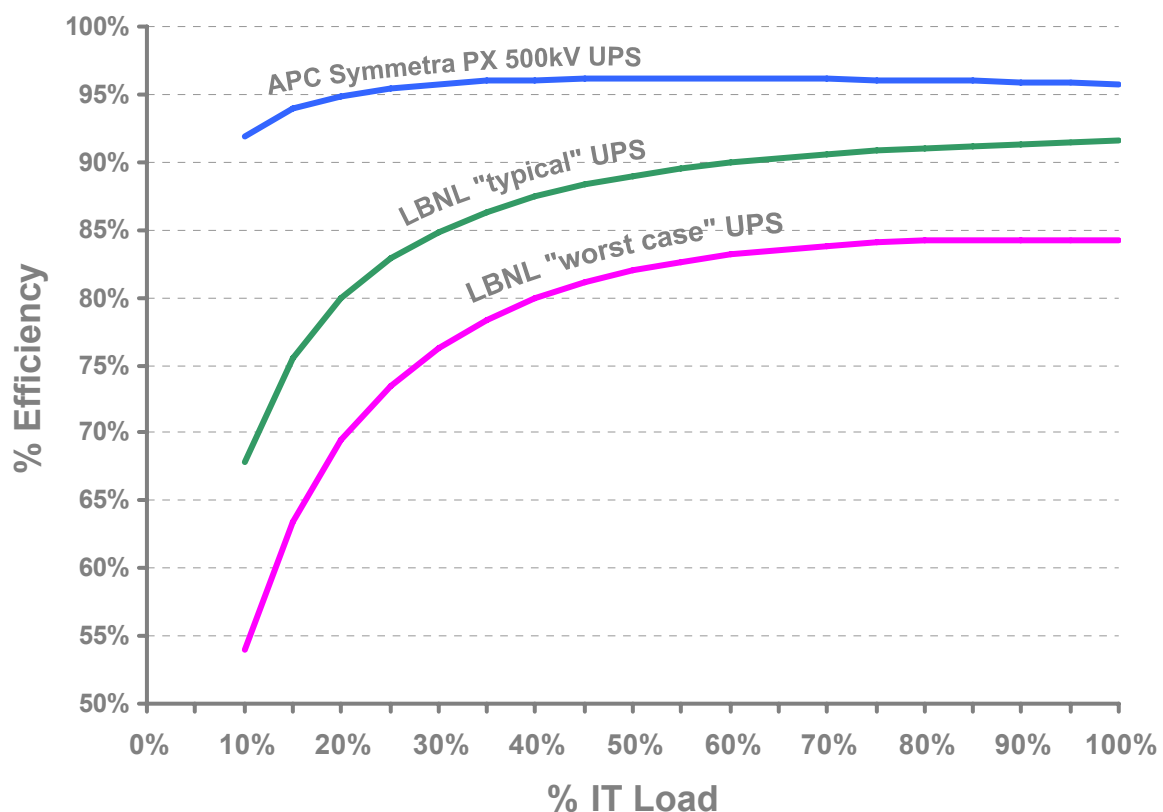*Figure 8 – UPS efficiency as a function of load comparing latest generation UPS to historic published data*

Figure 8 shows that the efficiency of the latest UPS systems is significantly higher for any IT load, and the efficiency gain is greatest at lighter loads. For example, at 30% load the newest UPS systems pick up over 10% in efficiency when compared to the average of currently installed UPS systems. In this case the actual wattage losses of the UPS can be shown to be reduced by 65%. It is important to note that UPS losses (heat) must also be cooled by the air conditioner, creating further power consumption.

---

[3] LBNL report on UPS efficiency: http://hightech.lbl.gov/documents/UPS/Final_UPS_Report.pdf, Figure 17, page 23 . Accessed March 21, 2008.

Some newer UPS systems offer an "economy" mode of operation that allows the UPS manufacturer to claim higher efficiency. However, this mode does not provide complete isolation from utility-mains power quality problems and is not recommended for data center use. The high-efficiency UPS and the efficiency data used in the architecture described in this paper and shown in **Figure 8** are for a double conversion on-line UPS with complete isolation from input power irregularities.

415/240V AC power distribution

▷ **Improves power efficiency**

High-efficiency AC power distribution – using the European standard of 415/240V instead of the current North American standard of 208/120 V – offers an opportunity for significant efficiency improvement in North America. Sending power to IT loads at 415/240 instead of 208/120 eliminates power distribution unit (PDU) transformers and their associated losses. In addition to this efficiency gain, the elimination of PDUs has the added benefits of reducing copper costs, reducing floor loading, and freeing up additional space for the IT equipment footprint. Data centers that use transformer-based PDUs typically suffer an efficiency reduction of from 2% to 15%, with the larger percent losses occurring in data centers operating with redundant power paths and lighter IT loads.

The use of 415/240 V AC power distribution creates efficiency benefits only in North America – most of the rest of the world already operates using this or the nearly equivalent 400/230 V AC standard. The baseline efficiencies described in this paper are based on North American designs, and therefore the efficiency gains include the effect of removing PDU transformers.

For more about the use of high-efficiency 415/240 V AC distribution in North American data centers, see APC White Paper #128, _Increasing Data Center Efficiency by Using Improved High-Density Power Distribution_.

Variable-speed drives on pumps and chillers

▷ **Improves cooling efficiency**

Pumps and chillers in the data center cooling plant traditionally operate with fixed speed motors. The motors in such arrangements must be configured for maximum expected load and worst case (hot) outdoor environmental conditions. However, data centers typically run at only part of their design capacity, and they spend most of their operating life with outdoor conditions cooler than worst-case. Therefore, chillers and pumps with fixed-speed motors spend much of their operating time with their motors working harder than necessary.

Pumps and chillers equipped with variable-speed drives (VFDs) and appropriate controls can reduce their speed and energy consumption to match the current IT load and the current outdoor conditions. The energy improvement varies depending on conditions, but can be as large as 10% or more, especially for data

centers that are not operating at full rated IT load, or for data centers with chiller or pump redundancy. Variable-speed drives on pumps and chillers can be considered a form of "automatic rightsizing."

Some of the efficiency gains of variable-speed drives can be obtained by stage-control or multiple fixed-speed pumps and chillers. However, these systems can require substantial engineering and typically deliver less than half of the gains of VFDs.

Variable-speed drives on pumps and chillers are an additional cost, compared with fixed-speed devices. For some seasonal or intermittent applications the energy savings from this extra investment may have a poor return on Investment. However, for data centers that run 7 x 24 all times of the year the payback times can be as short as a few months, depending on the specific data center.

Capacity management tools

▸ **Improves utilization of power, cooling, and rack capacity**

Most data centers do not fully utilize power, cooling, and rack capacity. The primary symptom of this condition is the low average operating power density of data centers; while the power density of modern IT equipment is in the range of 5-20kW per cabinet, the typical data center is operating at an average of 3kW per cabinet or less. This difference means that data centers are physically much larger than needed, with longer airflow patterns, more air mixing, longer power distribution wiring runs, and more lighting than is actually required.

The primary reason that data centers typically operate at low power density is the inability to manage power, cooling, and rack space capacity in an efficient and predictable manner. The result of physically spreading out the IT loads is that the efficiency of the power and cooling systems is reduced. An effective system of tools and rules allows data centers to operate at higher power densities with the following efficiency benefits:

- Shorter airflow paths, resulting in the need for less fan horsepower

- Less air mixing, resulting in higher heat rejection temperatures

- Higher heat rejection temperatures, resulting in improved chiller efficiency

- Higher heat rejection temperatures, resulting in increased air conditioner capacity

- Shorter wiring lengths, resulting in lower wiring and PDU losses

- More IT load can be powered by the same power and cooling infrastructure

In addition to the problem of operating at low power density, most data centers operate with power and cooling "safety margins" ranging from 15% to 50%. The safety margin is the minimum allowable percent difference between the IT load and the ratings of the power and cooling devices. Safety margins are a form of intentional over-sizing, where the over-sizing is used to protect the system from overloads or overheating due to an imprecise understanding of the performance of the system. Safety margins are, in effect, a way to account for ignorance about the system.

 WP126 Rev 0

Safety margins impose two major penalties on the performance of a data center system. First, they significantly increase the capital costs of the data center, because they force the purchase and installation of equipment (capacity) that you are not allowed to use. Second, they drive down the efficiency of the data center by forcing it to operate away from the maximum efficiency point on the load curve.

An effective capacity management system consists of the tools and rules that allow a data center to operate at a higher density and with smaller safety margins (without compromising safety). The benefits of a practical system are on the order of 5% in overall infrastructure electrical efficiency, in addition to the capital savings due to higher power density which can be on the order of 5-10%. A system that allows more IT equipment to operate within a given power and cooling infrastructure "envelope" increases both the economic efficiency as well as the electrical efficiency. It can be shown mathematically that the **incremental** electrical efficiency associated with squeezing another watt of IT load into a given power and cooling envelope is greater than the overall efficiency of the data center, which means that it is generally more efficient to squeeze another watt of IT load into an existing data center than to place it in a new data center.

For a discussion on the principles and operation of an effective power and cooling capacity management system, see APC White Paper #150, *Power and Cooling Capacity Management for Data Centers*. An example of a commercially available capacity management system is shown in **Figure 9**:

Figure 9 – *Example of an operating screen from a commercially available power and cooling capacity management system*
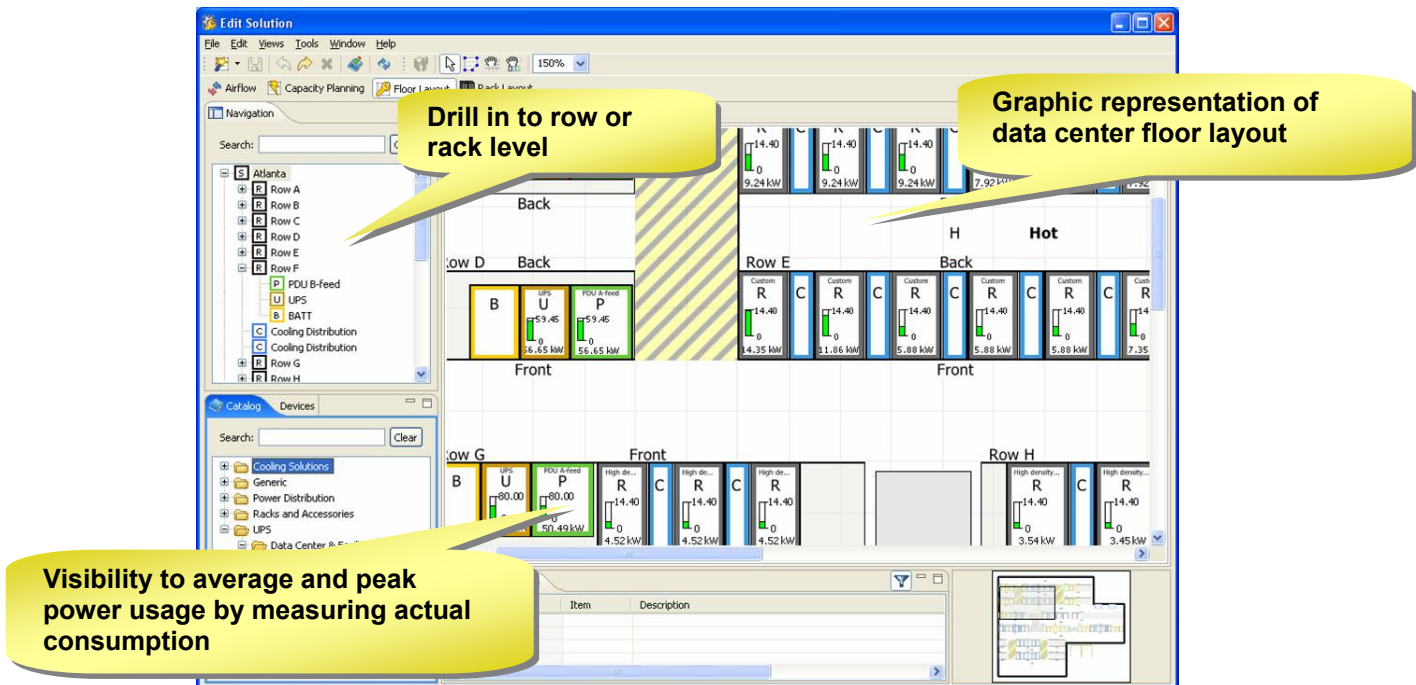
Room layout tools

### Optimizes room layout for cooling efficiency

Many electrical inefficiencies in data centers result from how power and cooling equipment is assembled into a complete system. Even if very high-efficiency power and cooling equipment is used, low overall efficiency is often the result. One of the biggest contributors to this problem is the physical layout of cooling equipment and the IT equipment.

A room layout optimized for efficiency is one in which:

- Airflow path lengths are minimized to reduce fan power

- Airflow resistance is minimized to reduce fan power

- IT equipment exhaust air is returned directly at high temperature to the air conditioner to maximize heat transfer

- Air conditioners are located so that airflow capacities are balanced to the nearby load airflow requirements

Some of these objectives are encouraged or enforced by the basic design of the cooling equipment, such as row-based cooling. However, floor layout of both the IT equipment and the air conditioners have a major effect on the optimization. The optimum floor layout varies between data centers and depends on the shape and size of the room, target IT power densities within the room, and other site-dependent factors. To create an optimized layout requires adherence to certain rules and requires complex calculations. Fortunately, these rules and calculations can be automated by computer aided design tools. An example of a data center layout tool that optimizes air conditioner layout is shown in **Figure 10**.

Figure 10 – *Room layout tool to optimize positioning of air conditioners*
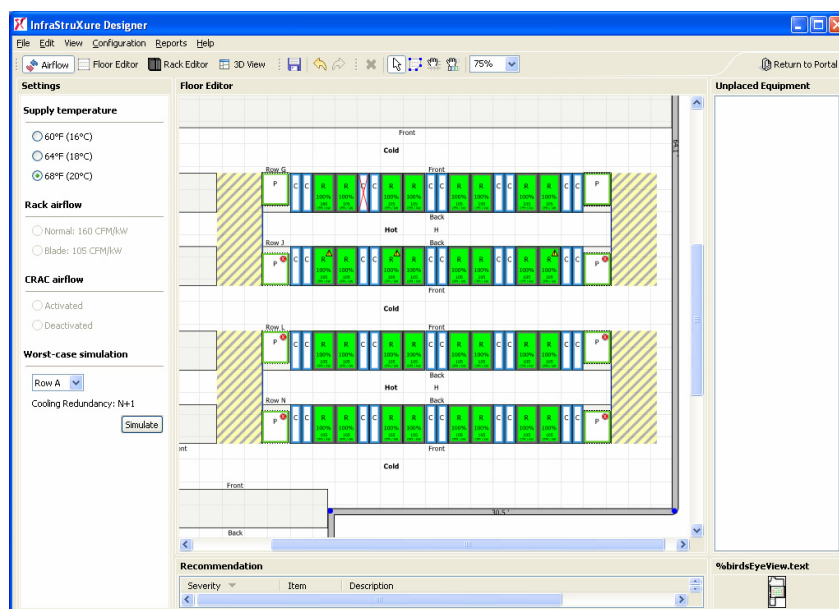
Overall efficiency gain from new architecture

When the elements of the improved architecture are combined, the total reduction in electrical consumption is as much as 40% when compared with traditional designs, as described earlier in this paper. The breakdown of savings by data center subsystem is given earlier in **Figure 4**. The overall *data center infrastructure efficiency* (DCiE – Green Grid's established measure) is a curve that varies as a function of the IT load as shown in **Figure 11**.

Figure 11a shows DCiE as a function of load for a high-availability data center with dual power path design and N+1 air handlers. **Figure 11b** shows the same data, but for a typical data center without power or cooling redundancy. In comparing these graphs we learn that:

- For traditional data centers, power and cooling redundancy reduces the overall DCiE by about 5%

- Power and cooling redundancy have a negligible effect on the efficiency of the improved architecture

- Phased deployment of modular power and cooling has the most significant efficiency benefit for data centers with power and cooling redundancies, particularly at light load

The efficiency gains described in this paper are affected by many other factors that vary between actual data centers. Some of these factors include:

- Dropped ceiling for air return in a traditional data center

- Uncoordinated room-perimeter air conditioners fighting each other

- Lack of hot-aisle/cold-aisle rack layout

- Energy-efficient lighting

- Powering air handlers from the UPS

- Imbalanced sizing of power and cooling systems

- Full dual-path air handlers

- Full dual-path chillers

- Packaged chillers or DX glycol systems

- Shallow raised floor (0.5m or less)

- Large auxiliary loads (personnel space, network operations center)

- Hot and/or humid climate

- Very long coolant pipe runs

None of these factors or conditions were assumed in the development of the data in this paper. Nevertheless, all of these factors can be quantified, modeled, and analyzed. The models, techniques, and analysis used in this paper can be applied to a specific existing or planned data center. For example, this is done as a normal part of the Data Center Efficiency Assessment Service offered by APC.

Figure 11 – *Data center efficiency curve showing effect of improved architecture*

11a – Dual power path, N+1 air handlers
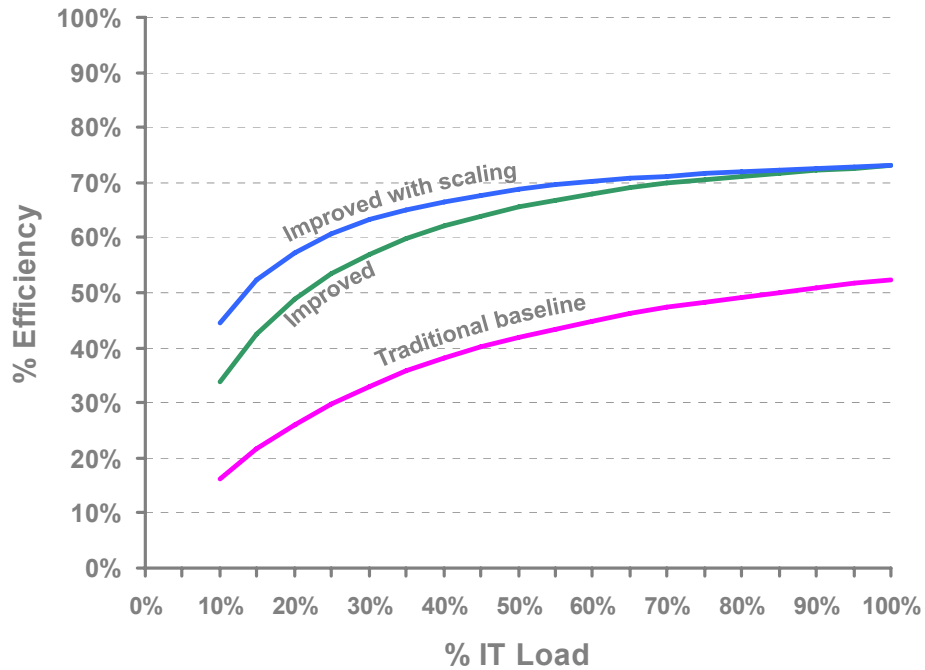
11b – Single power path, N air handlers
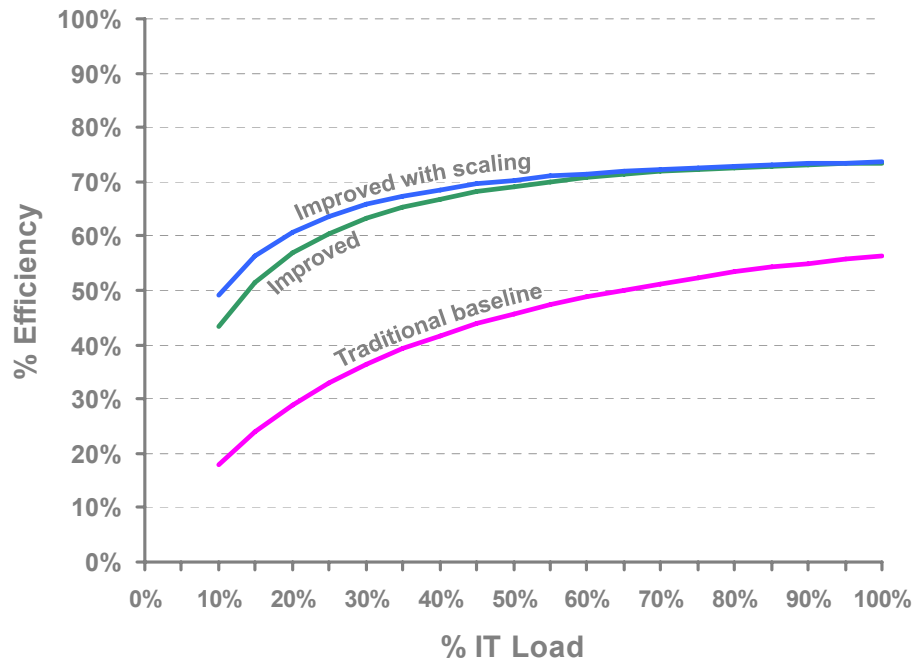
Comparison to Other Proposed Approaches

There are a number of hypothetical approaches to improving power and cooling efficiency that are not incorporated in the proposed architecture described here. In particular, **DC power distribution** and **direct piping of coolant to servers** have been suggested as architectural changes that might improve the efficiency of data centers in the future. It is instructive to assess how much additional gain is possible using these approaches when compared with the high-efficiency architecture described in this paper.

DC power distribution

It has been suggested that replacing conventional AC power distribution with DC could reduce data center electrical losses significantly. The benefits are said to accrue from the elimination of PDU transformers, replacing a UPS with an AC-DC converter, and the creation of new IT equipment that accepts high-voltage DC input in addition to or instead of AC. These efficiency benefits are quantified in **Table 1**.

Table 1 – *Data center efficiency improvement from DC distribution, compared to conventional design and to this paper's high-efficiency architecture*

DC distribution element	Savings compared to conventional design	Savings compared to high-efficiency architecture described in this paper
Elimination of Transformers	5%-10%	none
Replacing UPS with AC-DC converter	5%-15%	none
New IT equipment that accepts high voltage DC input	4%	2%
TOTAL IMPROVEMENT	**13%-28%**	**2%**

Table 1 shows that DC power distribution offers significant improvement over conventional AC data center designs but very little advantage over the high-efficiency architecture described in this paper. The key reasons why the architecture of this paper achieves nearly the same performance as the hypothetical high-voltage DC architecture are:

- Both systems eliminate the electrical waste of PDU transformers

- New high-efficiency AC UPS systems achieve the same efficiency as AC-to-DC high-voltage UPS systems

- Both systems operate IT power supplies at a higher input voltage, which improves their efficiency.

High-voltage DC distribution offers a slight theoretical advantage over the architecture described in this paper, but it is experimental, not yet commercialized, and lacks standardized regulations. By contrast, the approach described in this paper has worldwide regulatory approvals and is available today. For a more complete discussion on the various types of DC power distribution and a detailed quantitative analysis of the efficiency of AC vs. DC distribution, see APC White Paper #63, *AC vs. DC Power Distribution for Data Centers*, and APC White Paper #127, *A Quantitative Comparison of High-Efficiency AC vs. DC Power Distribution for Data Centers*.

Direct piping of coolant to servers

Many of the inefficiencies of data centers today are due to problems with airflow in the data center. The distribution of electrical power is predictable because it is transmitted directly to IT equipment from the supply via specific wires. Contrast this to the distribution of air cooling, which follows an invisible and often incomprehensible path from the air conditioners to the IT loads. To make the cooling system more effective it is logical to assume that directly connecting the cooling fluids to the IT loads, analogous to the flow of power, would make the system much more predictable and possibly more efficient.

Modeling of the efficiency of direct coolant piping to servers does show a significant potential gain in efficiency, when compared to conventional designs. However, **Table 2** shows that when compared with the high-efficiency row-based cooling architecture described in this paper, direct coolant piping shows only very small potential gains. The fact that row-based cooling provides most of the efficiency gain of direct coolant piping is not surprising, because row-based cooling is effectively bringing the coolant supply much closer to the IT loads.

Table 2 *– Data center efficiency improvement obtained by using direct piping of coolant to servers, compared conventional design and to this paper's high-efficiency architecture*

Direct piping element	Savings compared to conventional design	Savings compared to high-efficiency architecture described in this paper
Higher temperature coolant return	**5%**	**none**
Fan losses	**10%**	**5%**
Pumping losses	**-5%**	**-2.5%**
TOTAL IMPROVEMENT	**10%**	**2.5%**

Unfortunately, IT equipment designed to be cooled by direct coolant supply is not available and not expected to be available in the near future. In addition, there are significant cost and reliability issues to be solved. Fortunately, the high-efficiency architecture described in this paper achieves most of the benefit of direct piping of coolant to servers, but achieves it with existing air-cooled IT equipment and with equipment that can be deployed today.

Practical Performance Limits

The earlier review of the five contributors to electrical "waste" in data centers – unnecessary inefficiency and suboptimal configuration of devices – suggests areas for improvement. This naturally raises the question: What are the practical limits to the reduction of inefficiencies, such as the existence of fundamental laws of physics or practical engineering principles that constrain the potential for energy savings?

Surprisingly, there is no theoretical limit as to how low the non-IT losses could be for a data center. For this reason, *all* power consumed by power equipment, cooling equipment, and lighting must be considered to be waste (this is the **SUPPORT** power in **Figure 1**). For example, a data center using natural convection of outdoor air, in combination with superconducting electrical systems, is theoretically possible with no losses and 100% delivery of input power to IT loads.

However, there are practical limits to power and cooling efficiency today, given available technology and reasonable budgetary limits.

The most significant practical barriers to further efficiency gains of power and cooling systems, beyond those described for the improved architecture in this paper, are related to the cooling systems. The pumping and transport of heat via coolants and air conditioning systems is a stable and mature technology. While we can expect further optimization and integration of these systems in the coming years, efficiency gains for traditional air conditioning systems of only about 5% are expected beyond those described in this paper.

Free cooling, and air conditioning systems designed to take advantage of it, have the potential to gain another 5-10% in efficiency, depending on geographic location. When combined with the expected incremental gains in the performance of air conditioning technology, this would allow the DCiE efficiencies to reach the range of 88%, compared to 73% for the system architecture described in this paper.

Conclusion

Conventional legacy data centers operate well below the efficiency that is possible using proven designs incorporating readily available power and cooling equipment. This paper provides an example of an improved architecture that incorporates high-efficiency power and cooling equipment, combined with configuration and operation strategies that optimize efficiency.

One key finding of this paper is that purchasing high-efficiency devices is not sufficient to ensure a high-efficiency data center. An architecture and strategy that uses such high-efficiency equipment in an efficient manner, and reduces over-sizing, is just as important as the efficient hardware itself. When high-efficiency equipment is combined with an effective architecture, savings of 40% of the total electrical power of the data center are possible when compared with conventional designs.

About the Author:

Neil Rasmussen is the Senior VP of Innovation for APC-MGE Critical Power and Cooling Services. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is the principal architect of the APC-MGE InfraStruXure system.

Prior to founding APC in 1981, Neil received his Bachelors and Masters degrees from MIT in electrical engineering where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981, he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.

Related APC White Papers