Power and Cooling Capacity Management for Data Centers

White Paper 150

Revision 3

by Neil Rasmussen

> Executive summary

High density IT equipment stresses the power density capability of modern data centers. Installation and unmanaged proliferation of this equipment can lead to unexpected problems with power and cooling infrastructure including overheating, overloads, and loss of redundancy. The ability to measure and predict power and cooling capability at the rack enclosure level is required to ensure predictable performance and optimize use of the physical infrastructure resource. This paper describes the principles for achieving power and cooling capacity management.

Contents

Click on a section to jump to it

APC by Schneider Electric White Papers are now part of the Schneider Electric white paper library produced by Schneider Electric's **Data Center Science Center**
DCSC@Schneider-Electric.com

Schneider Electric

Introduction

Link to resource
White Paper 155
Choosing and Specifying an Appropriate Power Density for a Data Center

Data center physical infrastructure capacity management is defined as the action or process for ensuring power, cooling, and space is provided efficiently at the right time and in the right amount to support IT loads and processes. This paper discusses power and cooling capacity management only. Issues related to space management are discussed in White Paper 155, *Choosing and Specifying an Appropriate Power Density for a Data Center*. The critical success factors for effective management of power and cooling capacities are:

- Providing accurate capacity forecasts
- Providing appropriate capacity to meet business needs

This forecasting and efficient provisioning of capacity is dependent on the ability to establish the power and cooling capability at the rack level. Having this capability is rare today. Data center operators typically do not have the information they need to effectively deploy new equipment at the rate required by the business, and are unable to answer simple questions such as:

- Where in my data center should I deploy the next server so I don't impact the availability of existing equipment?
- From a power and cooling availability standpoint, where is the best location to deploy the proposed IT equipment?
- Will I be able to install new equipment without negatively impacting my safety margins such as redundancy and backup runtime?
- Will I still have power or cooling redundancy under fault or maintenance conditions?
- Can I deploy new hardware technology, such as blade servers, using my existing power and cooling infrastructure?
- Do I need to spread out my blade servers to get reliable operation?
- When will I reach the limits of my current power and cooling infrastructure and require additional capacity?

The inability to answer these simple questions is common. For data centers which are grossly over-designed or under-utilized, the safety margins can allow successful operation with only a primitive understanding of overall system performance. The compromise in availability due to this lack of knowledge may result in a small, but tolerable amount of downtime. While not financially or energy efficient, in the short term, oversizing provides a safety margin until such a time as the available capacity equals capacity utilized. However, three factors are currently placing stress on data centers which are, in turn, exposing the inadequacies of current operating methods:

- Ultra, high density IT equipment
- Requirement to control total cost of ownership (TCO) and more fully utilize data center capacity
- Rapid pace of change due to virtualization and refresh cycle of IT equipment

Each of these factors leads to pressure to operate data centers in a more predictable manner.

High-density IT equipment

IT equipment drawing more than 8 kW per rack enclosure can be considered high density. Fully populated racks of servers can draw from 6 kW to 35 kW per rack. Yet the vast majority of data centers today are designed for a power density of less than 4 kW per rack. As mentioned earlier, more and more users are installing equipment that exceeds the design density of their data centers and the resultant stresses on the power and cooling systems can cause downtime from overloads, overheating, and loss of redundancy. Data center operators need better information regarding how and where to reliably deploy this equipment in both existing and new data centers.

Data center stress #1: High-density IT equipment
Overloads
Overheating
Loss of redundancy

Total cost of ownership

Most businesses cannot accept gross over-design or oversizing of data centers. The waste of capital and operating costs is significant. It is estimated that the typical data center today could hold up to 30% more IT equipment using the same facility power and cooling capacity if the capacity was properly managed. The typical data center today is not able to fully utilize its available power and cooling capacity, which reduces the system efficiency and drives up electrical power consumption by 20% or more when compared to a system where the capacity is properly managed. Capacity management tools can better utilize power and cooling resources and reduce electrical consumption.

Data center stress #2: TCO pressure
Unused capacity
Reduced efficiency
Unseen waste

Rapid pace of change

IT equipment in a typical data center is constantly changing. Equipment refresh cycles are typically below three years and equipment is constantly being added or removed on a daily basis. Furthermore, the power and cooling requirements of the IT devices themselves are not constant but vary minute-by-minute as a result of virtualization and power management features implemented by IT equipment vendors. The historic "try it and see if it works" method of deploying IT equipment is no longer viable, with overheating a common result. Capacity management tools must provide real time planning capabilities to address these challenges, and they must provide this capability in a cost effective, easy-to-install, easy-to-use, pre-engineered form. To better understand the effects of virtualization and cloud computing on the physical infrastructure and how to manage them, see White Paper 118, *Virtualization and Cloud Computing: Optimized Power, Cooling and Management Maximizes Benefits*.

Data center stress #3: Rapid change
3-yr. refresh cycle
Day-to-day equipment changes
Minute-by-minute load changes

Related resource
White Paper 118
Virtualization and Cloud Computing: Optimized Power, Cooling and Management Maximizes Benefits

Capacity supply and demand

To provide simple answers to the basic questions users have about capacity, a systematic approach to capacity management is required. **The foundation of capacity management is the ability to quantify the supply and the demand for both power and cooling.**

While having power and cooling supply and demand information at the room or facility level helps, it does not provide sufficiently detailed information to answer the questions about specific IT equipment deployments. On the other hand, providing power and cooling supply and demand information at the IT device level is unnecessarily detailed and difficult to achieve. An effective and practical level at which to measure and budget power and cooling capacity is **at the rack level**, and this paper utilizes that approach.

The model described in this paper quantifies power and cooling supply and demand at the rack level in four important ways:

- As-configured maximum **potential demand**
- Current **actual demand**
- As-configured **potential supply**
- Current **actual supply**

This information allows a complete description of the current status of a data center power and cooling at the rack level. These descriptions are explained below and illustrated in **Figure 1**.

As-configured power and cooling maximum *POTENTIAL DEMAND*

The power management systems in modern servers can cause the power to vary by 2 to 1 or more during typical operation. The maximum "as configured" power and cooling demand represents the peak values that can be caused by this variance in the rack. This information can be established at the time of system configuration via trending, it may be reported directly by the IT equipment, or it may be derived by other means.

The maximum power and cooling demand is always greater than or equal to the actual power and cooling demand and is critical information for capacity management.

Current power and cooling *ACTUAL DEMAND*

This is the value of power consumed and heat generated at each rack at any given point in time. Ideally, this is done by real-time measurement of electrical power consumption at the rack level. For virtually all devices, power consumed in watts equals the heat generated in watts. For other devices – including uninterruptible power systems (UPS), power distribution units (PDU), air conditioners, and VoIP routers – the heat output in watts is not equal to the power consumed, but can be mathematically derived. Rack power consumption can be measured by the power distribution system or it can be measured by the IT equipment itself, and the reported power consumed by the set of IT devices within a rack can be summed to obtain the rack power.

The as-configured power and cooling maximum *POTENTIAL SUPPLY*

The as-configured power and cooling maximum potential supply is defined as the amount of power and cooling that could potentially be delivered to the rack level by the installed infrastructure equipment. The potential power and cooling supply will always be greater than or equal to the actual power and cooling supply. If the maximum potential supply for any given load is greater than the actual supply being delivered to that load, this indicates that the system is in a degraded state. This can be caused by a number of factors, such as:

- Blocked air filters in the cooling system
- A decrease in outdoor heat rejection capability due to extreme environmental conditions
- The loss of a power module in a modular UPS

It is an important function of a capacity management system to recognize when the current actual supply is not the same as the design value, and to diagnose the source of the constraints of the system that are preventing realization of the design supply capacity.

The current power and cooling *ACTUAL SUPPLY*

The actual power and cooling supply at a rack is determined using information about the power and cooling distribution architecture of the data center power and cooling system, the actual current capacities of the bulk powering and cooling sources, and the effects on the available capacity of other loads.

The actual power supply at a given rack is determined by knowing the available branch circuit capacity to the rack, constrained by the availability of unutilized power of upstream sources such as PDUs and UPS. In some cases, the available capacity is further constrained by the design or configuration of the power system. For example, a modular system might not be fully populated or the design may call for dual power feeds.

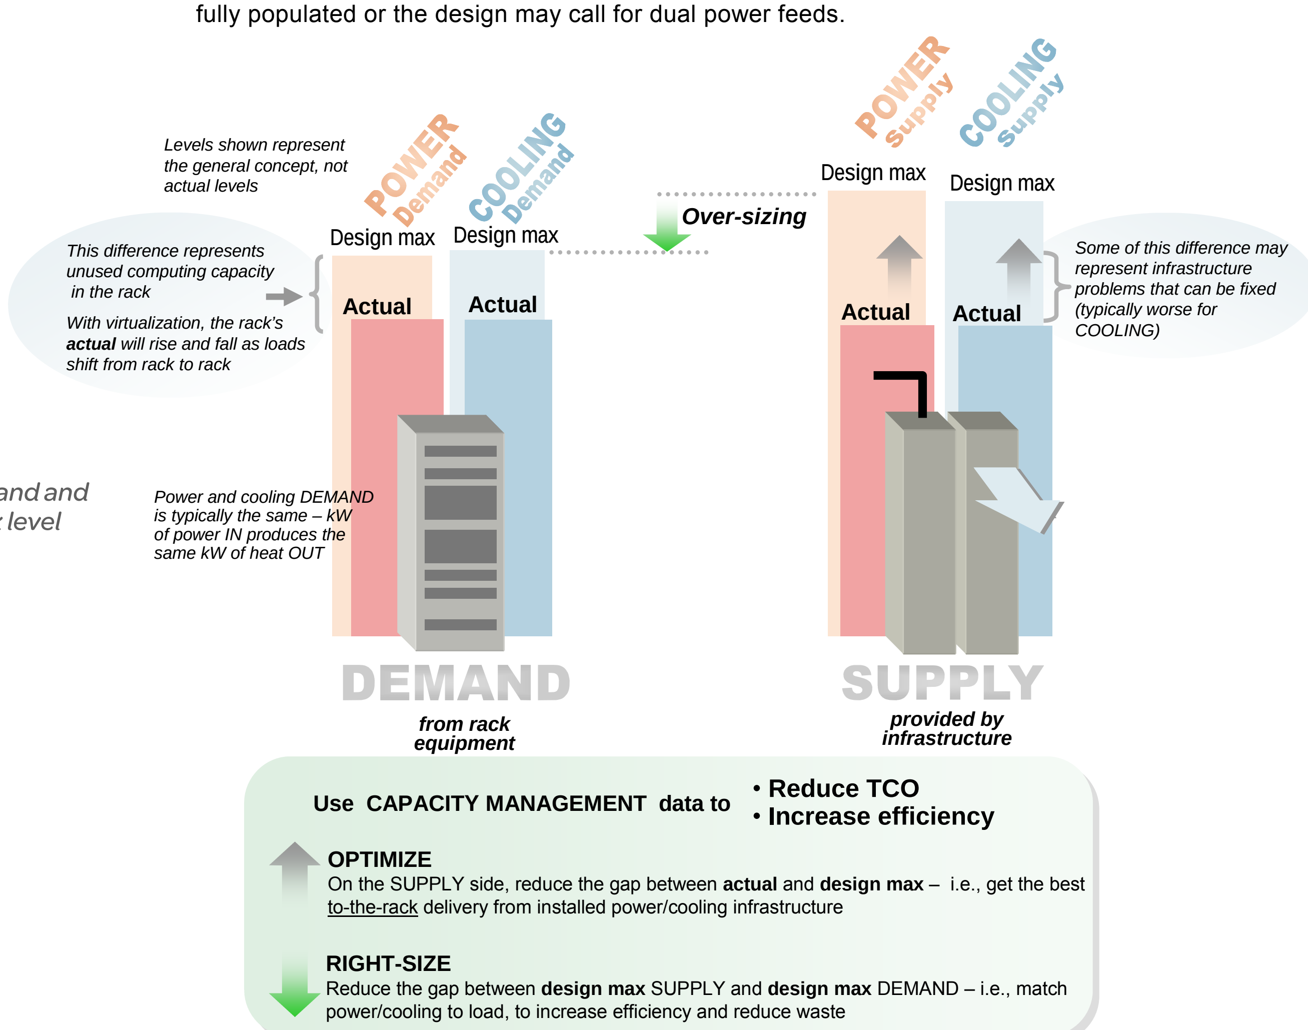

Figure 1

Quantifying demand and supply at the rack level

Determining the actual cooling supply at a rack is typically more complex than determining the power supply, and is highly dependent on the air distribution architecture. Unlike the power architecture, where the flow of power is constrained by wires, airflow is typically delivered to an approximate group of racks, where it spreads among the racks based on the draw of the fans in the IT equipment. This makes the computation of available air capacity more complex and sophisticated computer models are required. In cases where the supply or return air are directly ducted to racks, the cooling supply at a rack is better defined and therefore can be computed with improved accuracy.

System-level capacities

The demand on power and cooling is established at the rack as shown in **Figure 2**. The supply, as described in the previous section, must also be understood and quantified at the rack. However, the power and cooling supply system is not established rack-by-rack but is hierarchical, with supply devices such as UPSs, PDUs, and air conditioners supplying groups of racks. Bulk supply devices such as the power service entrance and cooling towers also represent sources of capacity supply that must be sufficient for the demand. Therefore, in addition to quantifying power and cooling supply capacity at the rack, it must also be quantified at the aggregate levels aligned with the supply devices.

Figure 2

Source of demand vs. source of supply

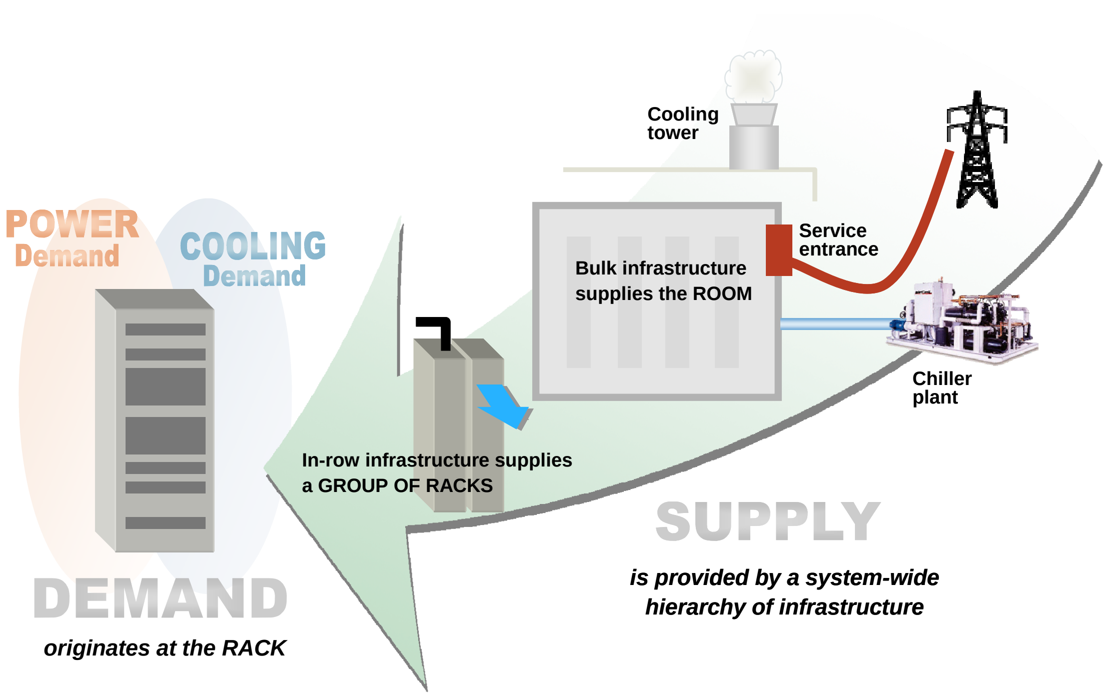

Supply must always be greater than or equal to demand to prevent the data center from experiencing a failure. This must be true at each rack, and it must also be true for each supply device supplying groups of racks. Therefore, at any given time, there is always excess capacity (as long as overall supply is greater than or equal to overall demand). Excess capacity comes in four different forms for purposes of capacity management, which are:

- Spare capacity
- Idle capacity
- Safety margin capacity
- Stranded capacity

Each of these types of excess capacity is explained in the following sections and illustrated in **Figure 3**.

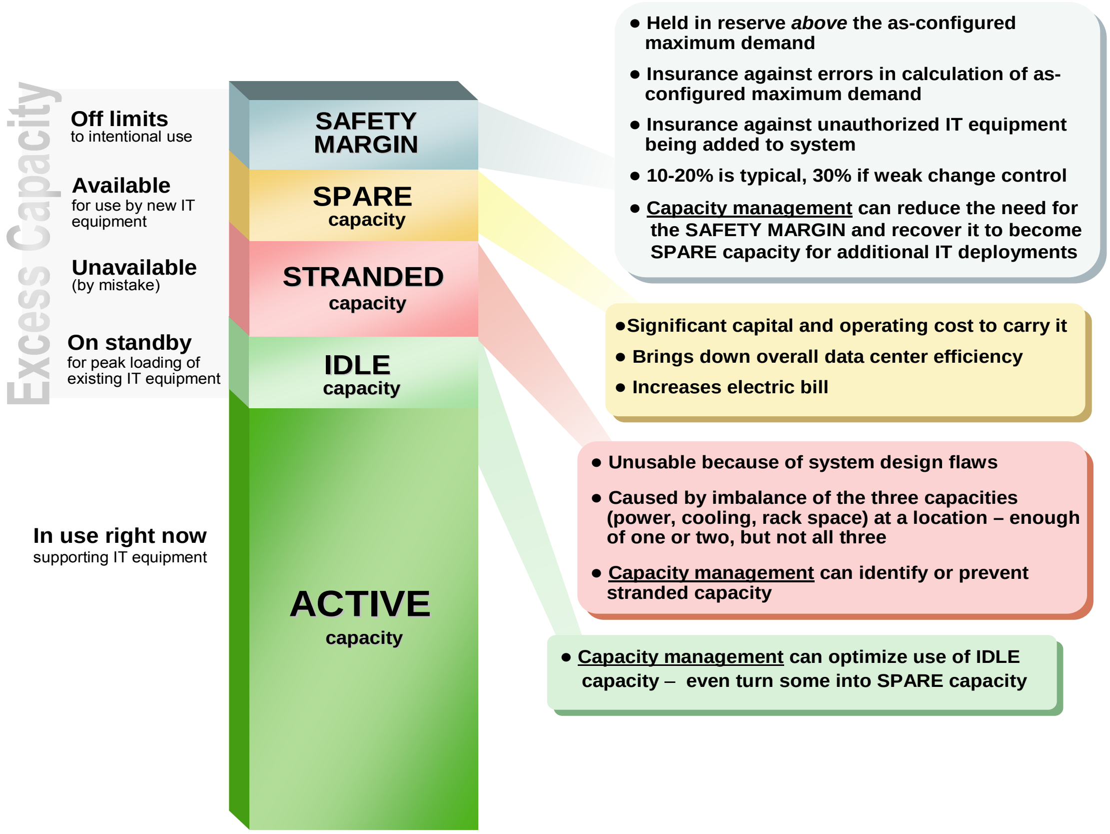

Figure 3

Types of excess capacity

Spare capacity

Spare capacity is the current actual excess capacity that can be utilized "right now" for new IT equipment. Carrying spare capacity has significant capital and operating costs related to the purchase and maintenance of the power and cooling equipment. Furthermore, spare capacity always brings down the operating efficiency of a data center and increases its electrical consumption.

In an effective capacity management architecture for a growing and changing data center, certain types of spare capacity, such as spare utility connection capacity, are cost effective. However, power and cooling equipment should ideally be installed only when and where needed to meet growing demand.

Related resource
White Paper 143
Data Center Projects: Growth Model

An effective capacity management system must comprehend and quantify growth plans. For more information on quantifying growth plans see White Paper 143, *Data Center Projects: Growth Model*.

Idle capacity

Idle capacity is the current actual excess capacity that is held available to meet the as-configured maximum potential power or cooling demand. The existing IT equipment might need this capacity under peak load conditions, so this idle capacity cannot be used to supply new IT equipment deployments.

Idle capacity is a growing problem caused by power management functions and automated virtual machine movement within IT equipment. The idle capacity must be maintained for the times when power-managed IT equipment switches to high power modes.

Safety margin capacity

Safety margin capacity is planned excess capacity that is held available above and beyond the as-configured maximum potential power or cooling demand. Providing a safety margin allows system operation in the event of small errors in setting the maximum potential power and cooling demand, or in the event of some unauthorized IT equipment being added to the system. Safety margins in the range of 10% to 20% are typical, with up to 30% or more used in data centers with weak change control procedures. This represents capacity that cannot be used for IT deployments.

Stranded capacity

Stranded capacity is capacity that cannot be utilized by IT loads due to the design or configuration of the system. The presence of stranded capacity indicates a lack of one or more of the following capacities:

- Floor and rack space
- Power
- Power distribution
- Cooling
- Cooling distribution

A specific IT device requires sufficient capacity of all of the five above elements. Yet these elements are almost never available in an exact balance of capacity to match a specific IT load. Invariably, there are locations with rack space but without available cooling, or spaces with available power but with no available rack space. Capacity of one type that cannot be used because one of the other four capacities listed above has been used to its maximum capacity is called stranded capacity. Stranded capacity is undesirable and can seriously limit the performance of a data center. Unfortunately, most data centers have significant stranded capacity issues, including the following common examples:

- An air conditioner has sufficient capacity but inadequate air distribution to the IT load
- A PDU has sufficient capacity but no available breaker positions
- Floor space is available but there is no remaining power
- Air conditioners are in the wrong location
- Some PDUs are overloaded while others are lightly loaded
- Some areas are overheated while others are cold

Depending on the situation and the architecture of the power and cooling system, it might be impossible to utilize stranded capacity or it might be that only minor investments are needed to free stranded capacity so that it can be effectively used. By definition, using stranded capacity comes at a cost. It is often necessary to take down part of the installation or install new power and cooling components.

Stranded capacity is a very frustrating capacity management problem for data center operators because it is very hard to explain to users or management that a data center, for example, with 1 MW of installed power and cooling capacity can't cool the new blade servers when it is only operating at 200 kW of total load.

An effective capacity management system not only identifies and highlights stranded capacity, but also helps customers avoid creating it in the first place.

Managing capacity

The previous sections have established the framework for quantifying power and cooling supply and demand.

A power and cooling capacity management system based on measurement by technicians combined with paper calculations could be envisioned, and in fact this method is used in crude form in some data centers. However, with the advent of server virtualization and IT equipment that changes its own power and cooling demand dynamically, **the use of networked power and cooling instrumentation combined with power and cooling capacity management software is the only practical and feasible solution**. From a user's perspective, such a system would provide the following functionality:

- Presentation of capacity data
- Setting the capacity plan
- Alerting on violations of the capacity plan
- Modeling proposed changes

Presentation of capacity data

The current supply and demand conditions of the data center, including spare and stranded capacity and other capacity attributes described earlier in this paper should be presented at these levels:

- *Room level*: The bulk level supply and demand as well as the various capacities for the entire room. Typically focuses on facility level UPS, generator, chiller, cooling tower, and service entrance equipment.
- *Row level*: Power and cooling supply and demand associated with a row or other logical zone within the data center. Often associated with cooling or power distribution equipment that is row-oriented, such as PDUs, or row-oriented cooling systems. Particularly valuable for planning purposes when rack-level details about configuration of specific racks are not yet known.
- *Rack level*: Power and cooling supply and demand associated with a specific rack or cabinet. Information at this level is required to diagnose problems or to assess the impact of specific IT device deployments. May be associated with rack level distribution circuits or rack-oriented cooling systems.
- Organization level: As the focus on efficient computing increases, and data center operatons are being tasked to significantly reduce their energy budget, executives require insight into both consumption and capacity usage of the company data centers.

An effective capacity management system will provide a display of the above types of information in a hierarchical drill-down model, including a graphical representation of the layout of the data center. **Figure 4** illustrates the room-level view and **Figure 5** illustrates the organizational-level view.

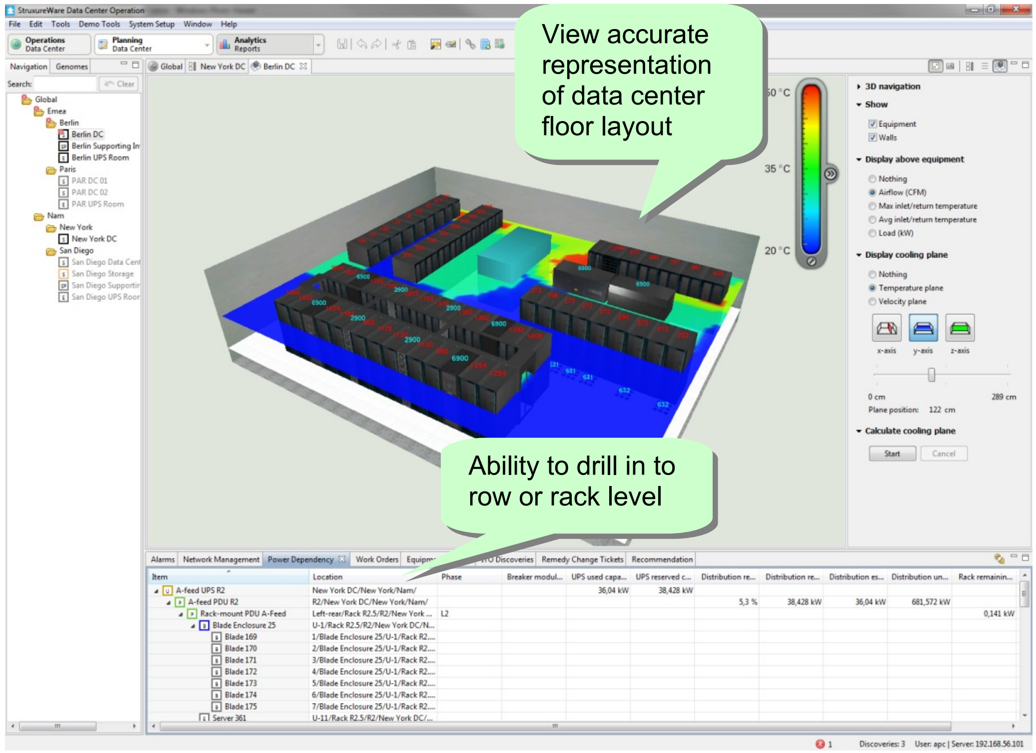

Figure 4

Example of data center layout views using Schneider Electric StruxureWare Data Center Operation application

Figure 5

Example view using Schneider Electric StruxureWare Data Center Operation: Vizor app. showing organizational-level capacity data for mobile and iPad usage.

Setting the capacity plan

A capacity plan must be established during the design of the data center. Once the power and cooling devices are installed in the data center, they constrain and in many ways "become" the supply side of the capacity plan. With today's scalable data center power and cooling solutions, it is possible to have a capacity plan that can adapt to IT growth plans in order to optimize capacity expense and electrical efficiency.

It's important to not just ensure sufficient capacity, but also ensure the **appropriate amount of** capacity. Too often the focus is on assuring sufficient capacity without regard for right-sizing to the actual IT needs. The common result is oversizing with the associated waste of capital expenditures, energy, service contracts, and water consumption.

Data center design tools help establish capacity plans and therefore should integrate into the data center infrastructure management (DCIM) system. An example of such a suite of software tools is Schneider Electric's InfraStruXure Designer data center design tool and their StruxureWare for Data Centers DCIM suite.

Alerting on violations of the capacity plan

Capacity related alerts should be triggered when actual conditions are outside the boundaries of the capacity management plan. These warnings can take the form of local, visual, or audible alerts, or can escalate via the management system as pages, e-mails, etc.

Capacity alarms are generated in response to user-defined events such as:

- Increased power consumption of installed equipment in a rack beyond the peak specified in the capacity management plan for a rack, a row, or the room
- Reduction in available cooling or power capacity at the row, rack, or room level due to loss or degradation of a power or cooling sub-system
- Cooling or power systems entering a state where they are not able to provide the redundancy specified in the capacity management plan

For many of these events, no actual hardware fault has occurred and hence no events would be triggered by traditional monitoring systems. In fact, most alerts provided by a capacity management system are predictive or proactive in nature. Note that in an actual data center, the capacity management system complements other monitoring tools such as real time fault, security, water leak, and temperature monitoring. An example of a monitoring system that provides both real time alerts as well as capacity management alerts is Schneider Electric's *StruxureWare Central* (**Figure 6**).

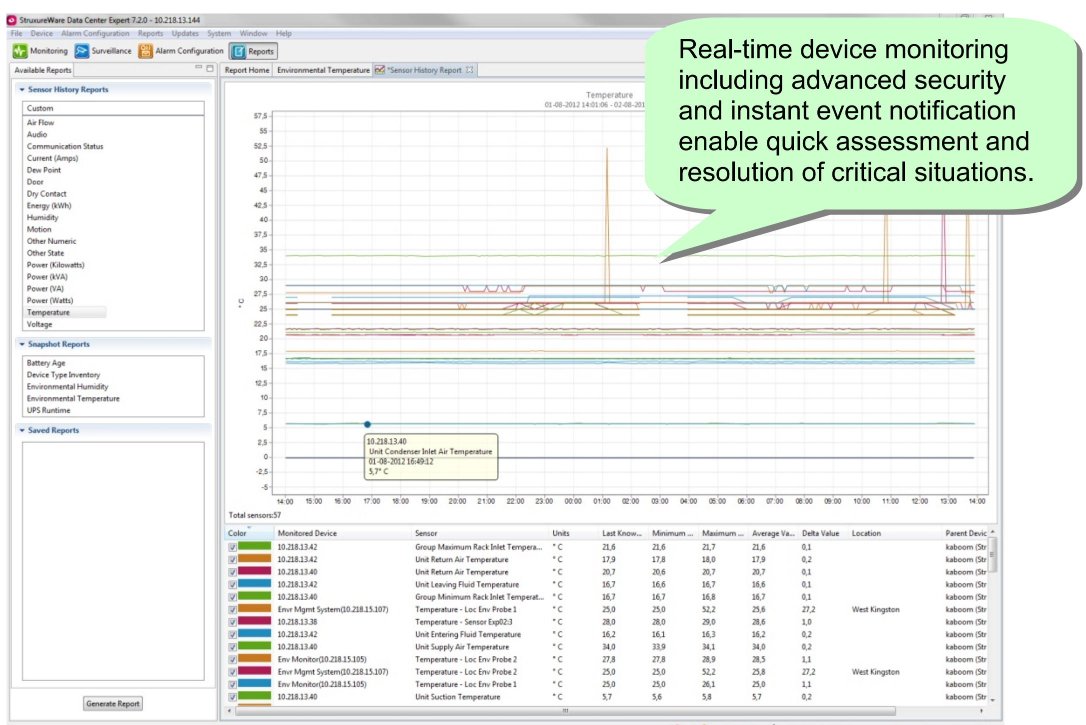

Figure 6

StruxureWare Data Center Expert as an example of a centralized monitoring system

Modeling proposed changes

In addition to the determination of current conditions, an effective DCIM system must provide the ability to analyze the capacity in historical and hypothetical situations. These scenarios may include:

- Simulating fault conditions, such as a loss of one or more power or cooling devices
- Analyzing planned growth versus actual capacity usage
- Proposals of equipment adds, removes, and relocations
- Trending based on historic data

The DCIM's capacity management functions should allow these scenarios to be evaluated against the current capacity management plan. An effective model would guide the user to select the best scenario from options, for example to maximize electrical efficiency or minimize floor space consumption. **Figure 7** shows an example of how DCIM can help with equipment changes by ensuring there is sufficient power, cooling and space resources for adding new IT equipment in a given location.

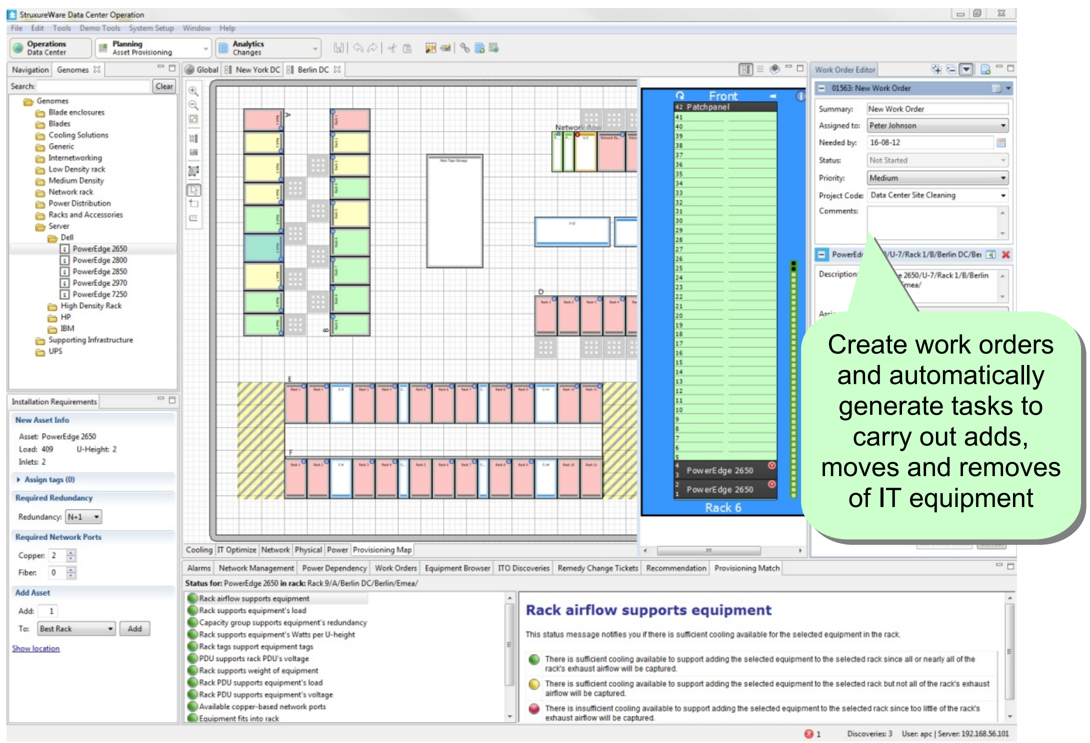

Figure 7

Schneider Electric's StruxureWare Data Center Operation application provides the ability to quickly and easily make optimal IT equipment changes

Monitoring power and cooling for IT devices

For most purposes, a rack may be viewed as a black box consuming power without awareness of the quantity, type, or location of the IT devices in the rack. The capacity management system as described is only weakly dependent on device-specific information, as long as the power consumed at the rack level can be directly specified or measured. When the power consumption is known, the cooling requirements can be estimated with high accuracy. When these rack level values are inside the boundaries of the capacity plan, no additional information is needed to ensure that the power and cooling systems can meet the demand. Therefore, with the right tools, capacity management offers many advantages to IT operations without the requirement of detailed inventory management of the IT devices. However, there are advantages to having information about the IT devices in the rack, or directly communicating with the IT devices. These advantages include:

- Awareness of power consumption characteristics of individual IT components
- Awareness of minimum and maximum time-varying power characteristics of the devices
- Awareness of unusual airflow requirements, or other unusual fan operating modes
- U-space utilization at the rack level (space capacity)
- Ability to give accurate assessments of the effects associated with adds, removes, or changes

To effectively use knowledge gained from detailed inventory management, the data must be understood by a capacity management system.

In general, most small to medium data centers do not have the process maturity and staffing needed to maintain rack-related IT equipment installation inventories and change history. Therefore, a capacity management system cannot depend on the presence of this information, but should be able to take advantage of it when available. As organizations mature, they can migrate from simplified capacity management to a more detailed solution that incorporates change and inventory management. The interaction between change management and capacity management is bi-directional as change management is highly dependent on capacity management information to predict the impact of proposed changes.

Figure 8

Summary of a robust capacity management system

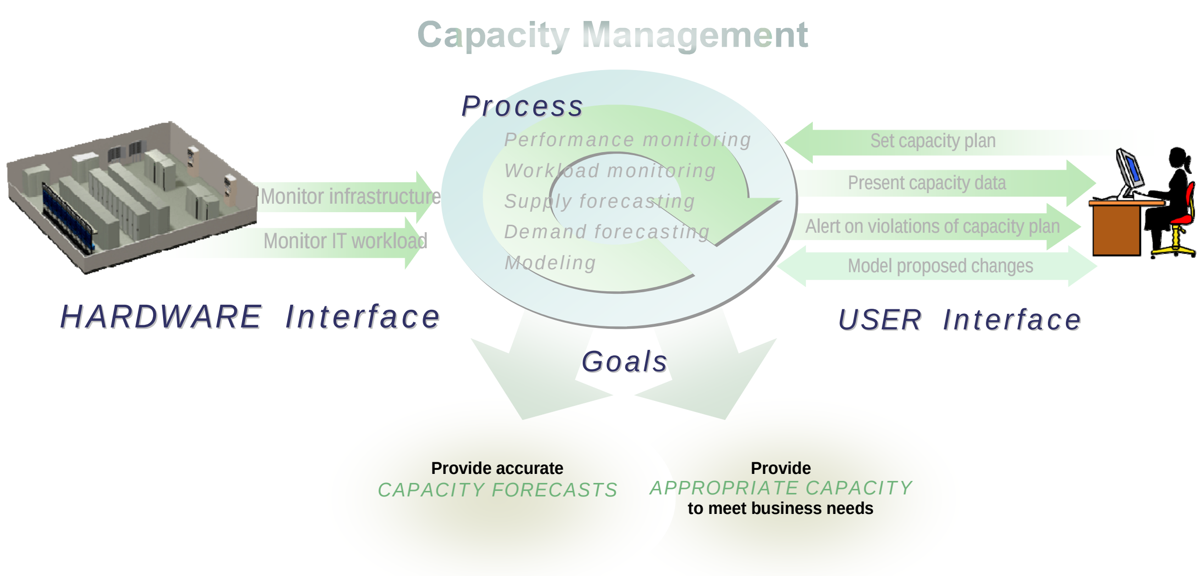

Conclusion

Capacity management is an essential part of the efficient planning and operation of data centers. The need for capacity management grows with the density, size, and complexity of the data center. A methodology for capacity management has been described. It has been shown that capacity management is not dependant on detailed information about the IT devices at the rack level and requires less effort to implement and maintain, compared to traditional detailed inventory management systems, while still providing most of the key benefits. If capacity management is implemented as described in this white paper, it can provide critical information about the state of the data center which is not provided by traditional monitoring systems. When combined with network power and cooling instrumentation, today's newer DCIM software systems provide the necessary tools for efficient and effective capacity management "out of the box".

About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 25 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC 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.

Resources

Click on icon to link to resource

Choosing and Specifying an Appropriate Power Density for a Data Center
White Paper 155

Data Center Projects: Growth Model
White Paper 143

Virtualization and Cloud Computing: Optimized Power, Cooling and Management Maximizes Benefits
White Paper 118

Browse all white papers
whitepapers.apc.com

Browse all TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative at
www.apc.com/support/contact/index.cfm