Power and Cooling Guidelines for Deploying IT in Colocation Data Centers

White Paper 173

Revision 0

by Paul Lin and Victor Avelar

Executive summary

Some prospective colocation data center tenants view power and cooling best practices as constraining. However, an effective acceptable use policy can reduce downtime due to thermal shutdown and human error, reduce stranded capacity, and extend the life of the initial leased space, avoiding the cost of oversized reserved space. This paper explains some of the causes of stranded power, cooling, and space capacity in colocation data centers and explains how high-density rack power distribution, air containment, and other practices improve availability and efficiency. Examples of acceptable use policies that address these issues are provided.

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

Colocation data centers can provide tenants a reliable and secure place to put their IT equipment. Additionally, tenants have the ability to take advantage of the shared power infrastructure, cooling systems, physical security, and redundant architecture. But, tenants generally try to deploy their equipment the way they see fit for various reasons including:

- Comfort level with the way they have always done things
- The need to adhere to design principles in previous data centers (even if they are antiquated)
- Belief that the cost of best practices do not provide a good or quick enough return on investment
- Belief that it does not have much influence on other tenants or the entire facility

This "deploy it my way" strategy leads to a hodge-podge data center with low energy efficiency and space efficiency and potentially lower availability. **Ultimately this strategy leads to higher costs for colocation providers, who in turn, pass these costs on to their tenants.** Following an acceptable use policy (AUP) agreement made between colocation providers and tenants can improve the efficiency of using infrastructure resources and keep costs lower for both parties than it might be otherwise.

A considerable amount of colocation providers and tenants face low energy efficiency and space efficiency challenges. This paper investigates the causes of stranded capacity in colocation data centers and best practices are described that can solve these issues.

How should tenants find a balance between adopting best practices and maintaining flexibility in colocation data centers? An effective AUP is a good way to do this. This paper describes the attributes of a policy which mandates some rules on how to deploy IT equipment in colocation environments in a way that efficiently uses power, cooling, and space resources.

Causes of stranded capacity in colocation data centers

Stranded capacity is capacity that cannot be used by IT loads due to the design or configuration of the system. This lost capacity involves a lack of one or more of the following:

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

Any IT device in a data center requires sufficient capacity of all five above. Yet these elements are almost never available in an exact balance of capacity to match a specific maximum IT load. Invariably, there are data centers with rack space but without available cooling, or data centers with available power but with no available rack space. Stranded capacity occurs when one type of capacity cannot be used because one or more other types have been used to its maximum capacity. Stranded capacity is undesirable and can seriously limit the performance of a data center. 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 power capacity, but no available breaker positions due to the proliferation of low density racks

- 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
- There's a willingness to increase rack power densities, but there's not enough power or cooling distribution capacity available

Depending on the situation and the architecture of the power and cooling system, it might be impossible to eliminate any existing 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, freeing up 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 new blade servers when it is only operating at 200 kW of total load.

IT space in a colocation data center is typically leased in units of cages, racks, or in some cases, subdivided rack space into ½, ¼, or even single U spaces to accommodate smaller computing needs. Both colocation providers and tenants are facing the challenges to balance power, cooling and space capacities for servers in colocation data centers. Most colocation data centers suffer from one or more of the following causes of stranded capacity:

- Difficult to plan the entire leased space in advance
- Proliferation of low density racks
- Design limitation of cooling capacity
- Mixing of hot and cold air streams

Difficult to plan the entire leased space in advance

In many cases, colocation data center space is normally leased over time to various tenants with different space requirements, resulting in inefficient use of the data center space. Over time, as the leased spaces are carved out, the remaining space between these fragments represents stranded space that is difficult or impossible to lease to new tenants despite excess power and cooling capacity. This situation **can strand bulk power and cooling capacity,** but can be unlocked by increasing the average rack density.

Proliferation of low-density racks

Low-density racks are generally provisioned with a single-phase 16 to 20 Amp circuit depending on the service voltage. Despite the low power consumption (i.e. less than 2kW), each rack power distribution feed occupies a branch circuit position in a power distribution unit (PDU). The proliferation of such racks can quickly consume all available breaker positions **stranding the remaining power capacity of the transformer within the PDU**. Increasing the average rack density can help to recover stranded PDU capacity.

Design limitation of cooling capacity

The vast majority of colocation data centers are designed for close to 4 kW per rack. Compaction of IT equipment has lead to higher rack power densities. In addition, virtualization will often lead to much higher densities as virtualized servers tend to be grouped together. Also according to DCD 2012 Census, the average rack power density is 4.05 kW with 58% of racks under 5 kW/rack, 28% are 5 kW to 10 kW/rack and 14% of racks are over 10 kW/rack[1]. Cooling and or power distribution that is unable to support these higher density racks **may lead to stranded rack space**. Adaptable cooling and power distribution solutions allow increased rack densities that can recover this stranded rack space.

Mixing of hot and cold air streams

Stranded cooling distribution capacity is caused by one or more the following, which allow hot and cold air streams to mix:

- Lack of hot or cold aisle containment that prevents mixing of hot return air and cold supply air. Air mixing is exacerbated by the lack of a cold aisle / hot aisle arrangement due to the different rack configurations among tenants. This leads to increased server inlet temperature and degraded cooling capacity.
- Blanking panels are not used to close off open U-spaces leading to cold air leakage.
- Cable cutouts in raised floor tiles and gaps between racks can cause significant unwanted cold air leakage.
- Under-floor cabling acts as an “air dam” limiting air distribution at the front of IT racks causing hot spots while forcing air to short circuit back to the cooling unit.

More than 50% air leakage is routinely observed in data centers, which means that the cold air bypasses the IT equipment intakes and flows directly back to the cooling units[2]. This leads to **stranded cooling distribution capacity**. Containment solutions can recoup stranded cooling distribution capacity and allow for increased rack power densities.

Best practices promoted by an effective acceptable use policy

For some large colocation providers, tenants must sign an agreement that compels them to comply with the AUP when deploying and managing their IT. But smaller colocation providers often lack this kind of policy as they want to attract smaller customers who may treat this kind of policy as constraining. Some tenants are allowed to deploy their IT in a “deploy it your way” format with little to no design standards. As a result, these deployments do not conform to cold/hot aisle configuration which may prevent the use of containment strategies, all but guaranteeing stranded cooling distribution capacity. This type of environment makes it very difficult for the colocation provider to ensure that one tenant’s poor deployment practices do not cause hot spots for themselves or their neighboring tenants. These same problematic situations may occur with other resources such as power, space, and security.

In reviewing the four main causes of stranded power, cooling, and space capacity, we see that increasing the average rack density recovers the stranded resource in the first three. In essence, rack power density is the “shock absorber” between the constraint and the stranded resource. Air containment recovers cooling distribution capacity in the fourth case but is also an essential element to having a well-run and efficient data center. Best practices and solutions are available today that can avoid stranding valuable resources, but are compromised by poor operational practices and data center design[3].

---

[1] DCD Industry Census 2011: Forecasting Energy Demand. http://www.datacenterdynamics.com/research/energy-demand-2011-12

[2] http://www.upsite.com/media/whitepapers/reducing%20bypass%20airflow%20is%20essential%20for%20eliminating%20hotspots.pdf (accessed 4/9/2012)

[3] Data center design (not the topic of this paper) includes aspects such as ceiling height, provisions for bulk power and cooling expansion, 415 volt power distribution (North America only), cooling economizer modes, floor loading, data center physical infrastructure management, etc.

An effective AUP addresses these poor operational practices and provides mutual benefits for both colocation providers and tenants. This section will explain the following best practices and solutions for colocation providers and tenants to deploy IT in colocation data centers. These best practices should be specified in the AUP.

- Use overhead cabling
- Use wide racks to accommodate high density cable management
- Prevent mixing of hot and cold air streams
- Use supplemental cooling components for high density racks

Use overhead cabling

In data centers with raised floor-based cooling it is common to see large drafty holes in floor tiles to allow pass-through of under-floor cabling. These cut-outs might occur for nearly every rack and all PDUs in the data center and result in a significant amount of mixing between cold air and hot return air. Placing power and structured cabling overhead eliminates air mixing due to these floor tile cut-outs, eliminates air blockages, and also improves cable management. White paper 159, "*How Overhead Cabling Saves Energy in Data Centers*" provides further information on this topic.

Use wide racks to accommodate high density cable management

As discussed at the beginning of this section, rack power density provides the "shock absorber" between the present-day physical space constraint and an unknown future maximum power and cooling requirement. Increasing the power density over time allows a tenant to use the same leased space for a much longer period of time. However, for this strategy to work, the leased space must first account for the extra space required by wide racks (750-800mm) compared to standard width racks (600mm). Even if the initial IT deployment uses standard 600mm wide racks, eventually higher rack densities will required the use of 750mm wide racks to accommodate the extra cabling.

Prevent mixing of hot and cold air streams

An effective AUP will specify IT rack deployments in a hot/cold aisle configuration. This is a foundational practice and it is well documented that as rack density increases, more airflow is required to cool IT equipment which leads to increased mixing of hot and cold air. When unconstrained, some of this hot air will recirculate back to the IT equipment's intake causing hot spots.

In existing data centers, the separation of hot and cold air streams is managed through hot or cold aisle containment solutions implemented at the row or pod level, and by using air management devices at the individual rack level. Hot aisle containment can save 43%[4] in annual cooling system energy cost compared to cold aisle containment, corresponding to a 15% reduction in annualized PUE. In addition to aisle containment, there are other air management practices including blanking panels, ducted racks, and other air distribution devices. In caged environments, hot exhaust air from a cage should be constrained to that cage to prevent infiltration into adjacent cages. Regardless of what type of containment is used, the AUP should always specify the use of blanking panels and raised floor grommets to seal any unwanted gaps and floor tile cut-outs to prevent air mixing. This practice is critical to preventing hot spots and increasing cooling system efficiency. White Paper 153, *Implementing Hot and Cold Air Containment in Existing Data Centers*, provides logical steps for identifying the best containment solution.

[4] http://www.apc.com/whitepaper/?wp=135 (accessed 4/7/2013)

Figure 1

Example of an air containment solution installed to prevent mixing of hot and cold air streams

Hot aisle containment　　　　Resulting air distribution

Use supplemental cooling components for high density racks

Data centers with raised floor-based cooling tend to have dynamic pressure variations depending on the amount of under-floor cabling, depth of floor, and the overall raised floor geometry. These pressure variations may cause some perforated tiles to supply more or less airflow than others thereby leading to unpredictable cooling. An enclosure with fully or partially-ducted air management devices can overcome these pressure variations and provide cooling distribution capacity on demand, especially for high density racks or the racks for which power density is greater than the average design value.

Figure 2a and **2b** show examples of rack-mounted, fully-ducted devices which provide the functions just described. To provide high availability, the devices of **Figure 2a** and **2b** are typically provided with N+1 fans and dual-cords. In addition, fan speeds can be controlled to optimize system performance. **Figure 2c** shows an example of a floor-mounted device which can adjust its fan speed to supply a suitable airflow rate depending on rack exhaust air temperature or under-floor pressure.

Figure 2a (left)

Rack-mounted fully ducted air return unit

Figure 2b (middle)

Rack-mounted fully ducted air supply unit

Figure 2c (right)

Floor-mounted local ducted air supply unit

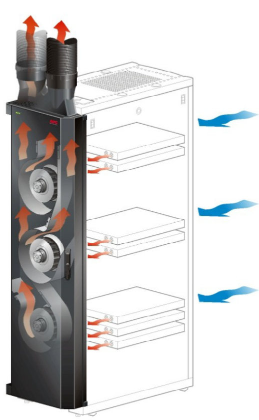

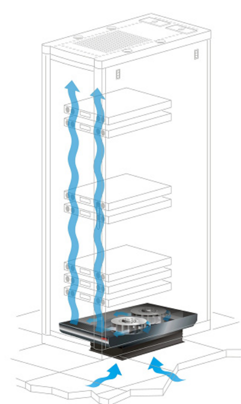

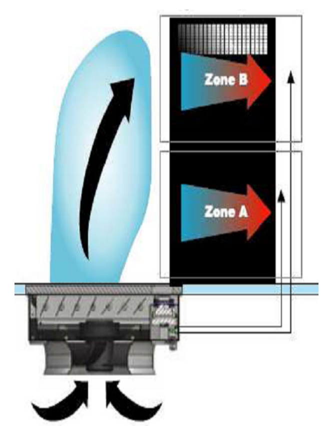

Other recommended best practices

There are other best practices that can be used to the benefit of both colocation providers and their tenants. While not typically found in AUPs, these bests practices are worth considering.

- Manage capacity when deploying IT
- Use scalable and efficient power architecture
- Use a close-coupled cooling architecture

Manage capacity when deploying IT

In data centers with well defined power and cooling distribution, tenants are able to use data center infrastructure management (DCIM) software in order to minimize stranded power, cooling, and space capacity. This is especially beneficial with the advent of server virtualization and dynamic power variation caused by power management technologies in servers and communication equipment. For more information on dynamic power variations see White Paper 43, *Dynamic Power Variations in Data Centers and Network Rooms*.

Automated gathering of capacity data is the first basic step to analyzing supply and demand capacity, followed by modeling proposed changes. Tenants can analyze capacity both in historical and hypothetical situations before deploying IT, including fault conditions, and planned growth. For more information about power and cooling capacity management, refer to White Paper 150, *Power and Cooling Capacity Management for Data Centers*.

Some DCIM software tools available on the market today enable colocation providers to monitor, analyze, and report on capacity status and trends at the cage, zone, room, or facility level. If the provider is using this type of DCIM suite and as long as the tenant has access to these DCIM management tools (or information) that is relevant to their own cage or room, investment in DCIM can be avoided. This is another point to consider when evaluating potential colocation partners. If the provider does not use DCIM management tools, however, then it is recommended that the tenant buy and implement DCIM tools at that site.

Use a scalable and efficient power architecture

An ideal power system in colocation data centers allows for scalable bulk power and power distribution. In data centers with only one power path, this architecture is flexible in that a second UPS path can be added for increased reliability. Scalable UPS solutions also allow bulk power to increase seamlessly over time. **Figure 3a** shows an example of a modular, scalable UPS.

A scalable, modular PDU offers many of the same benefits as a modular UPS. PDUs of this type can be used to increase power density without having to add new wiring. A power circuit from a modular PDU is a flexible cable that is plugged into the front of the PDU on site to meet the requirements of each specific tenant. Ideally, tenants would use a 3-phase rack PDU (i.e. power strip) and an associated circuit capable of supporting a peak rack power density between 10 and 15 kW. This practice allows the rack to scale in density without having to change the branch circuit frequently. The rack PDU would provide a mix of C13 and C19 receptacles which would accommodate a large percentage of IT and networking equipment. This accommodates rack density increases over the course of the data center lifetime. **Figure 3b** shows an example of a modular power distribution unit with replaceable circuit breaker modules. For more information on scalable and modular power distribution, refer to White Paper 129, *Comparing Data Center Power Distribution Architectures*.

Figure 3a (left)

Modular UPS unit

Figure 3b (right)

Modular power distribution unit

Use a close-coupled cooling architecture

Where containment solutions are not possible, an alternative option is to use close-coupled cooling, including row and rack-based that can address these problems. Row-based or rack-based cooling is an air distribution approach in which the air conditioners are dedicated to a specific row of racks or a single rack. This is in contrast with room-based cooling where perimeter cooling units are "dedicated" to the entire room. Row-based cooling may be installed above IT racks, adjacent to IT racks, or in combination.

Figure 4a and **4b** show examples of row-based cooling units which provide the above functions. To provide high availability, the devices of **Figure 4a** and **4b** are typically provided with N+1 fans and dual-cords. In addition, fan speed may be controlled to optimize system performance. Close-coupled cooling architectures can be used together with traditional raised floor architectures to increase cooling capacity and / or rack density.

Figure 4c shows an example of a rack-based cooling unit which is directly mounted to the IT rack. All of the rated cooling capacity is dedicated to the IT rack which allows for the highest power density and targeted redundancy (i.e. N+1 or 2N redundancy for a single rack).

Figure 4a (left)

Floor mounted modular row-oriented cooling unit

Figure 4b (middle)

Overhead modular row-oriented cooling unit

Figure 4c (right)

Modular rack-oriented cooling unit mounted with a rack

While these individual cooling units can be used within existing rows of racks, they can also be part of a pod as shown in **Figure 1**. A pod is typically formed by two rows of racks (e.g. 5 racks per row) with door locks and side panels and includes modular PDUs, UPSs, and cooling units to support the IT devices within the pod. A scalable modular pod can solve capacity problems in existing colocation data centers that were not designed for high density deployments. By isolating both hot and cold air streams, a high-density pod, at a minimum, neutralizes the thermal impact that high-density IT racks would otherwise have on traditional low-density data centers. For more information on this topic, see White Paper 134, *Deploying High-Density Pods in a Low-Density Data Center*.

Attributes of an effective acceptable use policy (AUP)

An effective AUP helps both tenants and providers to use one or more of the above best practices to reduce stranded capacity with increased availability and decreased operational costs. AUPs like this are becoming more common among colocation providers. Prospective data center tenants should evaluate a colocation provider's AUP as an important selection criterion and not base their decision solely on initial cost. **Table 1** provides examples of policies that should be included in effective AUPs.

Table 1

Example policies for AUPs

Attribute	Description
Power usage	**Tenants must comply with the following facility power usage rules when deploying their IT:** • Electrical power draw shall be less than 80% of the circuit breaker capacity. Violations of this policy may result in a special allocation fee per offending rack per month until the usage is corrected.
Cooling configuration	**Tenants must comply with the following cooling configuration rules when deploying their IT:** • IT racks shall be configured in a hot aisle / cold aisle arrangement. • All tenant equipment greater than 6 kW/rack[5] shall use containment and / or chimney racks; supplemental cooling components may also be used to increase cooling distribution capacity. • All tenants shall seal any open sections of racks using blanking panels to prevent cold air leakage. • Cable cutouts and gaps between racks shall be sealed to prevent cold air leakage. • Raised floor tiles shall not be left open except during maintenance work. • Doors to be used in aisle containment systems shall remain closed at all times unless entering or exiting the aisles and the doors must never be propped open.
Cabling	**Tenants must adhere to industry standards for cable management and the following rules:** • Power and structured cabling shall be supported overhead and not below the raised floor.
Rack and server	**Tenants must comply with the following rules when deploying their IT:** • Wide racks shall be used to accommodate high density cable management.

Conclusion

With many different tenants in the same data center, the likelihood of stranded capacity is higher for data centers when basic power, cooling, and space-related best practices are not enforced. Both the provider and the tenant can benefit from a reduction in stranded power, cooling, and space capacities. An effective AUP can help establish a balance between good practices and flexibility by deploying IT racks using a standardized approach that comprehends capacity management, air management, smart power distribution, and pods. AUPs like this are becoming more common among colocation providers. Prospective data center tenants should evaluate a colocation provider's AUP as an important selection criterion and not focus solely on initial cost.

About the authors

Paul Lin is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He holds a Bachelor's degree in Mechanical Engineering from Jilin University where he majored in Heating, Refrigeration and Air Conditioning. He also holds a Master's degree in Thermodynamic Engineering from Jilin University. Before joining Schneider Electric, Paul worked as the R&D Project Leader in LG Electronics for several years. He is now designated as a "Data Center Certified Associate", an internationally recognized validation of the knowledge and skills required of a data center professional.

Victor Avelar is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operation research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environments. Victor holds a bachelor's degree in mechanical engineering from Rensselaer Polytechnic Institute and MBA from Babson College. He is a member of AFCOM and the American Society for Quality.

[5] White Paper 46: http://www.apc.com/whitepaper/?wp=46 (accessed 4/7/2013)

Resources

How Overhead Cabling Saves Energy in Data Centers
White Paper 159

Cooling Strategies for Ultra-High Density Racks and Blade Servers
White Paper 46

Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency
White Paper 135

Implementing Hot and Cold Air Containment in Existing Data Centers
White Paper 153

Power and Cooling Capacity Management for Data Centers
White Paper 150

Comparing Data Center Power Distribution Architectures
White Paper 129

Deploying High-Density Pods in a Low-Density Data Center
White Paper 134

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

© 2013 Schneider Electric. All rights reserved.