Five Basic Steps for Efficient Space Organization within High Density Enclosures

White Paper 72

Revision 1

by Joe Kramer

> Executive summary

Organizing components and cables within high density enclosures need not be a stressful, time consuming chore. In fact, thanks to the flexibility of new enclosure designs, a standard for organizing enclosure space, including power and data cables can be easily implemented. This paper provides a five step roadmap for standardizing and optimizing organization within both low and high density enclosures, with special emphasis on how to plan for higher densities.

Contents

Click on a section to jump to it

APC 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

Introduction

Many data center professionals recognize the negative impact that crowded enclosures have on efficiency, on uptime, and on the overall look and feel of their data centers. The biggest concern is the impact on uptime, which is most often influenced by human error. A systems operator or employee who opens the door to a crowded enclosure can be easily confused by the mess of cables. Sometimes the wrong plug is pulled or jostled when attempting to move a piece of equipment.

Data center efficiency is driven by energy consumption, which is determined by the amount of power coming into the enclosure, and by the amount of cooling required to remove the heat generated by the equipment within the enclosure. The structure and organization of the enclosure plays a significant role in determining whether the heat removal process is efficient.

Furthermore, the overall condition of the enclosure reflects on the cleanliness and professionalism of the entire data center. Customers, investors, executives, and outside visitors all walk away with a negative impression when enclosure cabling resembles clogged knots of spaghetti (see **Figure 1**).

Figure 1

The crowded enclosure dilemma

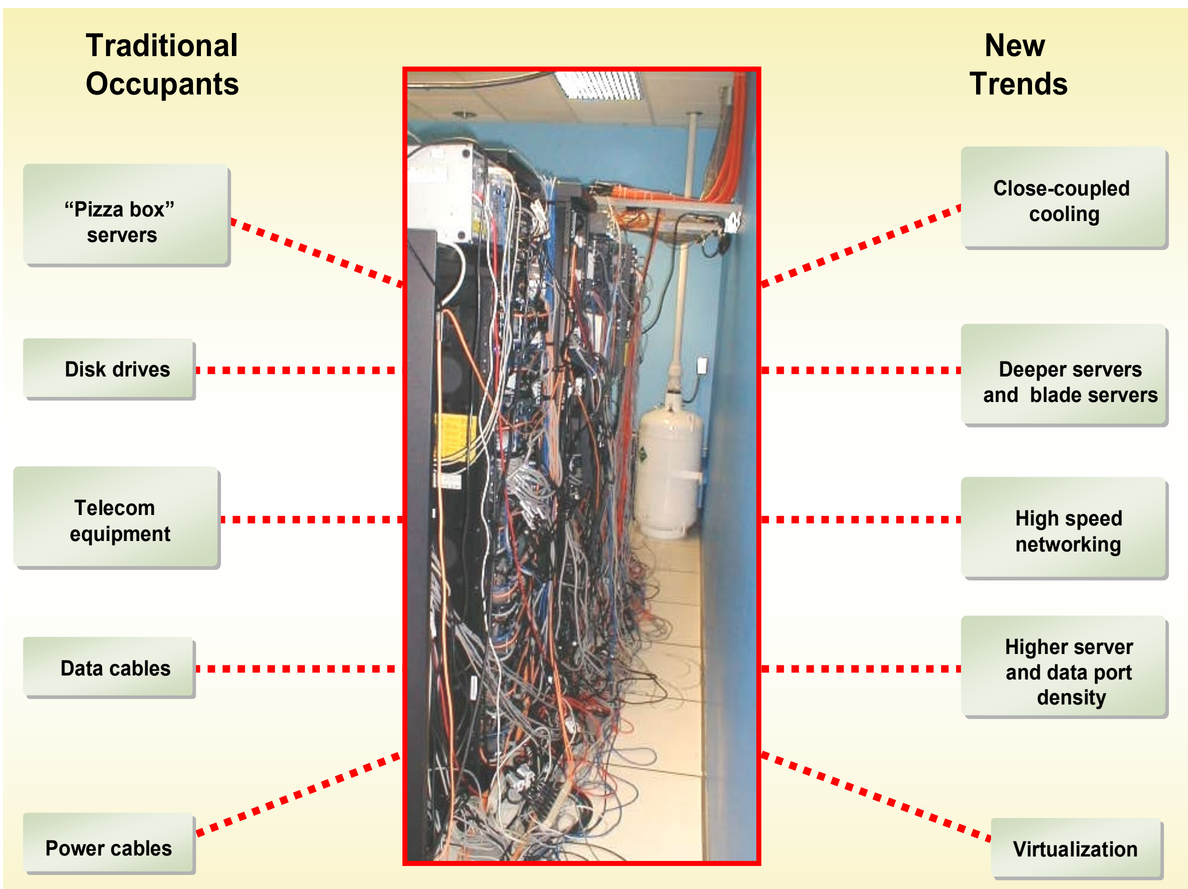

New trends such as virtualization, high density blade servers, and high efficiency add to the need for a comprehensive strategy for managing equipment and cabling inside of the enclosures. **Fortunately, new enclosure designs and sizes provide the flexibility and space to manage the new, more demanding data center environment.** This paper reviews five critical steps for facilitating the goals of improved uptime, efficiency, and physical appearance through improved internal enclosure organization.

> **> Do it yourself?**
>
> Some data centers depend upon internal staff to properly cable their enclosures. Others ask a vendor to perform cable installation and set up. Pulling cables for structured cabling projects is laborious work and the vendor likely has a high degree of skill in routing cables inside of the rack. The core competency of the user staff more often lies in the ability to maintain the rack, post installation, on an ongoing basis.

Step #1: Plan for higher density

In most cases, high density can imply two distinct enclosure configuration scenarios. The first consists of an enclosure populated with 1U / 2U servers. The second consists of an enclosure with blade servers. In the case of the enclosure populated with 1U / 2U servers, the data center personnel will have to contend with many cables. Such an enclosure can be populated with up to 40 servers. Some current generation servers have as many as 10 network data ports per server if configured to maximum density. Most of these individual servers require two network cables and two power cords. Blade servers, because of their considerable input power requirements, will require a large rack-mounted power distribution unit (PDU). Also, the rear of the enclosure must be kept clear since the blade servers have many hot-swap modules that require rear access.

Before deploying blade servers or prior to launching a virtualization initiative, the selection of an appropriate enclosure environment needs to be assessed. When planning, the first question to consider is whether any of the existing data center enclosures are suitable candidates for hosting higher densities.

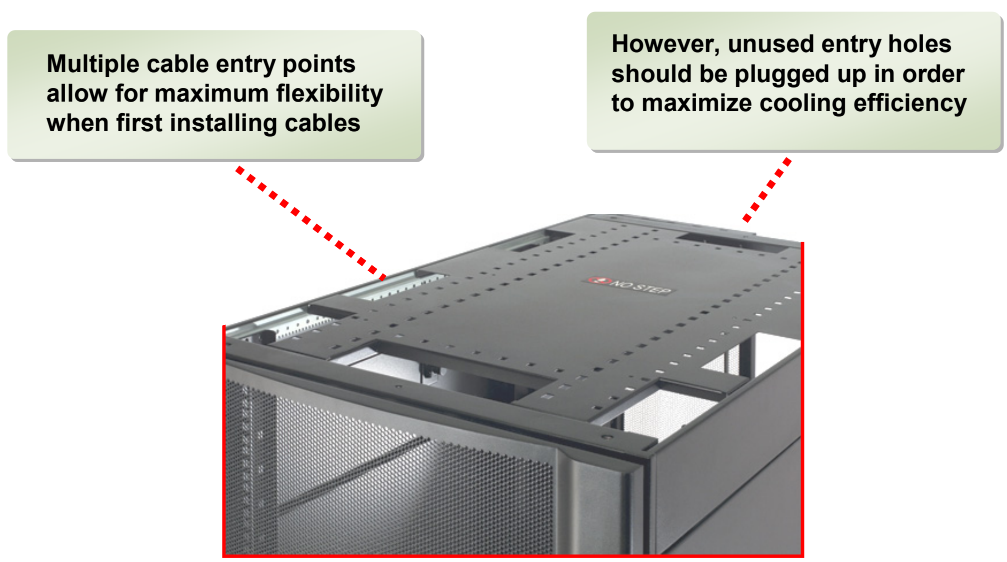

Figure 2

Available enclosure cable entry points can be critical to high density deployments

Higher density enclosures should offer the following accommodations:

- Vented front and back doors with at least 60% ventilated surface area
- At least two cable cutouts aligned with the rear channel of the enclosure to accommodate incoming, vertically managed cables (see **Figure 2**)
- At least 35 square inches (22,580 sq mm) of total area for cutouts in a standard rack and at least 70 square inches (45,161 sq mm) in a high density networking application
- The ability to accommodate a row-based cooling solution (these solutions should be considered if the enclosure will be expected to support a density of 10 kW per enclosure or more)

- The front of the enclosure should be capable of being sealed in such a way as to close off any return air paths (i.e., capable of supporting blanking panels)
- The ability to connect to other nearby enclosures
- Side panels should be provided to eliminate mixing of air between cabinets
- The ability to mount multiple vertical rack PDUs in the back of the enclosure

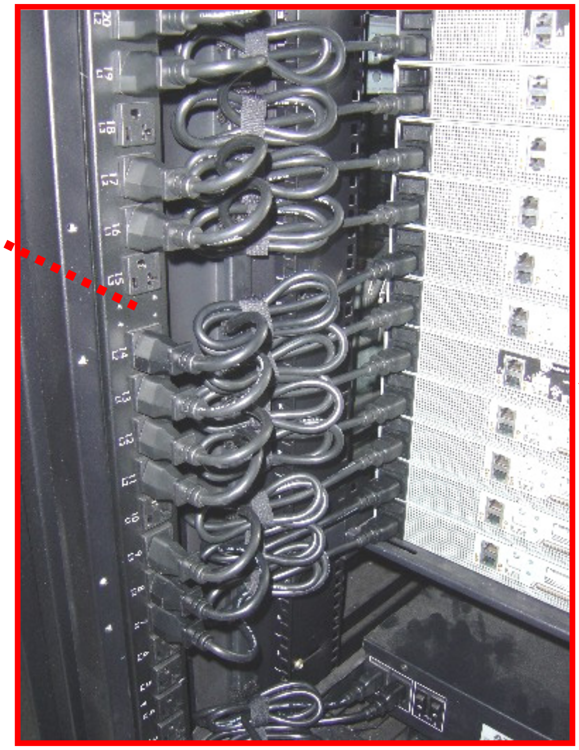

Figure 3

Vertical mount rack PDU for power cable management

Previous generation enclosures typically had vented doors but often did not offer advanced cable management options, airflow management, and accommodation for power distribution. New generation enclosures have the ability to adapt to higher density power and cooling solutions.

Step #2: Calculate enclosure power requirement

The user must select one of two methods to determine the maximum power required per enclosure:

1. Estimate the power usage of the equipment inside the enclosure. This method is commonly used when very high power loads are used such as enterprise class servers, blade servers, or high speed networking. The nameplate rating on this equipment may be somewhat misleading. Schneider Electric offers a tool to help estimate power usage which can help generate a more accurate reading (go to this link). Note that this IT equipment typically requires special purpose power dis-

> **Use of colored cables**
>
> For the most part, all power cables within the enclosure are colored black. Data cables come in a wide variety of colors. At the present time, no standards have been developed that link cable colors to cable functionality. However, use of colored cables can simplify the process of managing the equipment inside of the enclosure. For example, the color orange could be used for network traffic, the color grey for the management system, the color green for redundant network cable, and the color red for cables supplying critical systems.

tribution.

2. Assume a maximum power level of the enclosure based on an estimate of total data center utilization. For example, if the data center delivers 1 MW of power to the IT load, and the IT load consists of 100 racks, then a maximum power level could be estimated to be around 10 kW for the vast majority of the racks. This is easier to estimate and implement. This method is commonly used when a heterogeneous computing environment is expected and the specific IT equipment will be difficult to predict. The user manages the environment by restricting installation of more servers into enclosures approaching the maximum power level. Go to the APC TradeOff Tool™ Power Sizing Calculator for guidance on sizing data center power requirements.

Figure 4

Example of standard power cable and data cable layout

Related resource
White Paper 29
Rack Powering Options for High Density

Regardless of approach selected, the estimated power requirement will dictate the particular input power cord and plug configurations needed for the enclosure (see White Paper 29, *Rack Powering Options for High Density* for pros and cons of different input power cords).

All PDUs should have the ability to meter the input current at the branch circuit breakers. This allows the user to determine whether the circuit is approaching the maximum capacity or whether imminent danger of a circuit breaker tripping exists. Capacity management software, such as APC's Capacity Manager, can provide rack-level views which enable tracking of IT assets and available space and power.

Power distribution within the enclosure

When selecting a rack PDU, the user can either choose a single larger unit or multiple smaller units. If opting for a single large PDU (of 10 kW, for example) per enclosure, the overall number of PDUs to be installed will be lower. However, each of the input cables will be large and the rack PDUs may require internal circuit breakers which will increase their size. If opting for multiple smaller PDUs per enclosure, more space in the rear of the enclosure will be consumed that could otherwise be used for vertical cable management. Rack PDUs below 5.7 kW are a good choice for low to medium density enclosures.

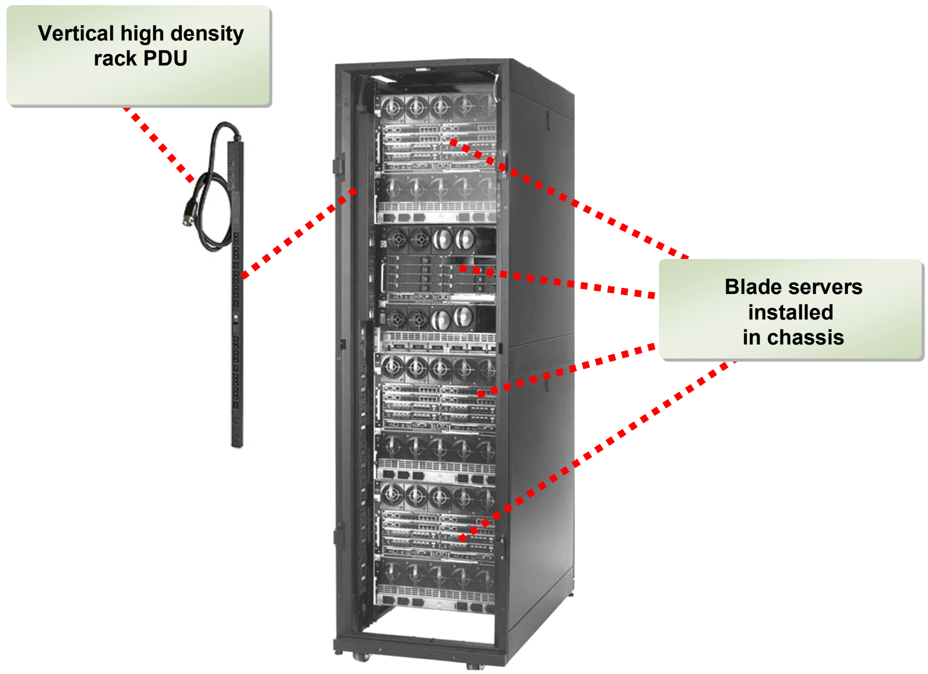

Figure 5

Enclosure with four blade chassis and high power density rack PDU

If users choose to implement redundant power supplies and 2N power systems, management of space inside of the enclosure becomes even more important. A 2N power system, for example, drives the requirement for two vertical rack PDUs and for twice the number of power cords going to each server.

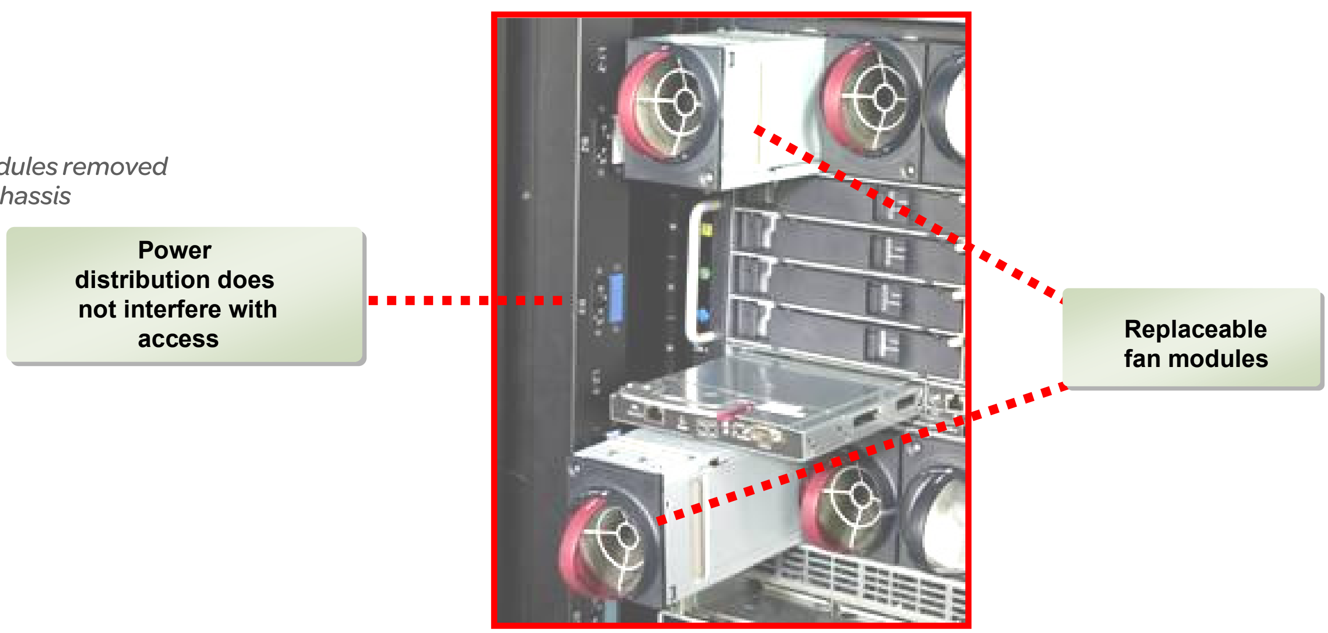

Figure 6

Blade fan modules removed from rear of chassis

Blade servers

Blade servers frequently require high density rack PDUs with high power capacities but fewer receptacles. An enclosure with four of the latest generation blade chassis can consume 15-20 kW but will only require 12 receptacles.

Blade chassis require clearance for addition and removal of components both in the front and in the rear of the chassis. Fan modules, power supplies, and management modules are frequently installed and replaced into the back of the chassis with service clearance needed from the rear of the enclosure (see **Figure 6**). This requirement means that the enclosure and rack PDU chosen for the application must have clearance for these components.

> **Single and dual power feeds**
>
> Compared to a single feed enclosure, dual feed enclosures require twice as many rack PDUs. In the case of the dual power feed, the two rack PDUs would mount on one side of the rear of the enclosure and the other side would be utilized for cable management. Best practice is for the PDUs to provide enough power and receptacles to power all the equipment in the enclosure. If a 2N redundancy requirement exists, two rack PDUs would be required. If more than two are required, best practice is to go to a deeper enclosure with more rack PDU mounting positions and to maintain all rack PDUs on one side of the enclosure.

Virtualization

Compared to a traditional data center environment, a virtualized server environment will require higher power densities and will be characterized by dynamic loads. When servers are virtualized, their power draw will cycle as applications are added and subtracted to and from each server. In the past, enclosure power draw only changed when new servers were physically added to the enclosure, but in this dynamic environment, increases in power consumption can occur without any physical changes.

Since virtualization maximizes the effectiveness and efficiency of every server, power densities will go up because every server can be operating near its peak performance. As a result, the ability to monitor the individual server power has taken on added importance.

Step #3: Select proper enclosure size

The most popular enclosures used today are 24 inches (600 mm) wide, 42 inches (1070 mm) deep, and 42U tall. This enclosure size has been preferred because the width is the same as standard data center floor tile and its depth is optimized for hot aisle / cold aisle enclosure layouts in a data center. The external dimensions of the enclosure facilitate the organization of enclosures on the data center floor.

Table 1

Enclosure size recommendations per density range

Density range	Application	Recommended enclosure dimension
2 to 5 kW per enclosure	1U/2U servers, mixed environment	24 in (600 mm) by 42 in (1070 mm)
5 to 8 kW per enclosure	1U/2U servers, mixed environment	24 in (600mm) by 42 in (1070 mm)
8 to 15 kW per enclosure	Virtualized 2U/4U servers, Blade servers	30 in (750 mm) by 42 in (1070 mm) or 24 in (600 mm) by 48 in (1200 mm)
16 kW per enclosure and above	Blade servers	30 in (750 mm) by 42 in (1070 mm) or 24 in (600 mm) by 48 in (1200 mm)
2 to 16 kW and above	High density networking	30 in (750 mm) by 42 in (1070 mm)

A properly designed enclosure of this size is very effective in protecting and supporting IT equipment up to a medium density level. However, in higher density enclosures (housing 5 kW and above) this enclosure size is no longer optimal as higher density PDUs and increased cable loads will cramp the inside of the enclosure (see **Table 1**).

Standard racks can be used for densities up to 8 kW. For higher density situations, an enclosure either wider than the 24 inch (600 mm) or deeper than the 42 inch (1070 mm) is chosen to provide the space needed for organizing additional data and power cables inside the enclosure. Wider enclosures have been a typical choice for networking applications and are now considered a logical choice for higher density server applications. Overall, the wider enclosure provides the most flexibility for equipment and cable organization. Deeper enclosures become an option when the uniqueness of the floor layout dictates a deeper rather than a wider enclosure or when more than two rack PDUs are required.

A common width for a wider enclosure is 30 inches; the 30 inch (750 mm) wide enclosure provides 6 inches (152 mm) of extra space on each side of the enclosure which allows for placement of vertical power strips and easier management of cables. In a more traditional 24 inch (600 mm) wide enclosure, the deeper sizes of the servers often create mechanical interference to proper airflow. The crowded space also limits serviceability and creates difficulties when routing cables. The wider configuration addresses these issues, by allowing sufficient clearances for servicing and by creating a less congested path for hot exhaust air to flow out.

Deeper enclosures (48 inch / 1200 mm) as opposed to the more traditional 40-42 inch (1070 mm) enclosures offer benefits similar to that of the wider enclosures. They deliver more space to the rear of the enclosure allowing for multiple rack-mount PDUs and easier cable management. A data center that has standardized on a 24 inch (600 mm) width enclosure can consider a 48 inch (1200 mm) deep enclosure when more internal space is needed.

The drawback of using a larger enclosure is the requirement for more data center floor space. However, an efficient interior design allows the data center manager to use all enclosure U spaces, maximizing the effectiveness of the vertical footprint of the enclosure. Larger enclosures are also considered for virtualization projects or for pure blade server installations when space constraints may be less of an issue.

Step #4: Implement smart cable management

When planning out the flow path of cables within the enclosure, a first step is to determine whether the cables will enter the enclosure from the bottom or top. If entering from the top, the location of enclosure roof cutouts and their proximity to the vertical cable channels need to be considered. If entering from the bottom, consideration must be given to any obstructions in the base (such as large equipment mounted in the bottom) that can interfere with the cable entry path.

Routing the cables into the bottom of the enclosure means that the cables will most likely be routed under a raised floor. The user loses visibility to the termination of the opposite end of the cable and this makes it more difficult to execute moves or changes later on. If cables come into the rack from the top, the overhead cable management makes it easier to trace cables and manage moves.

The most effective method for managing cables in high density environments is to implement patch panels or switches dedicated to cabling for a particular row of enclosures. These are typically located at the top or bottom of the enclosure. These patch panels or small switches will be terminated back to the core switch or router feeding its section of the data center. The core switch is typically located in another enclosure. This approach is effective because it separates the cabling inside the enclosure from the rest of the data center cable load.

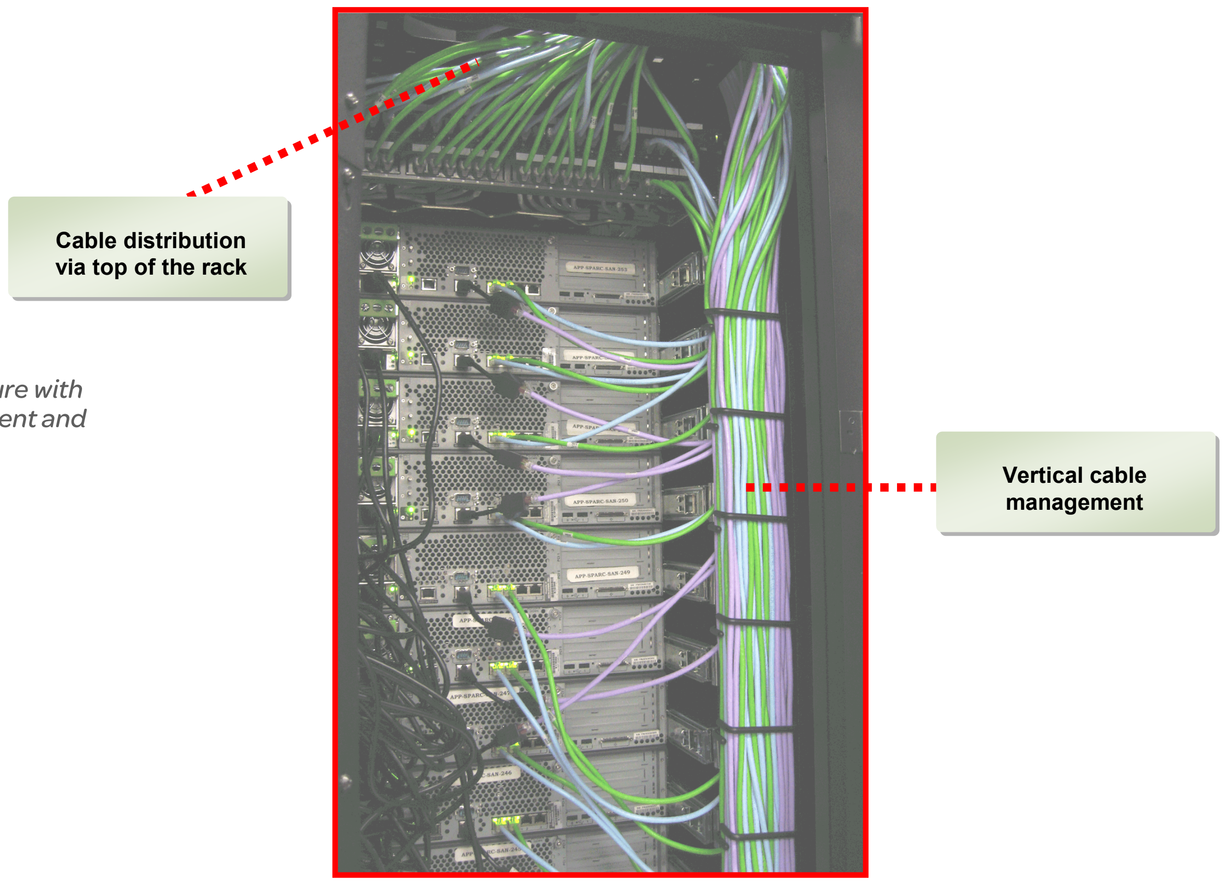

Figure 7

A well organized enclosure with vertical cable management and PDUs

Patch cords of various lengths will be connected between the patch panel and the IT equipment. The cords should be managed horizontally from the panel and / or the devices to the side of the enclosure and then managed vertically on the side of the enclosure (see

Figure 7). Since most IT components are designed with network cable terminations on the rear, the cable routing resides in the rear of the enclosure.

Users will often bundle data cables on the right rear side of the enclosure and utilize the left rear side of the enclosure for distribution of power cables. Enclosure management is facilitated when all server enclosures embrace a similar standard. It is important to calculate the total amount of vertically managed cable to ensure enough space in the enclosure.

Cable management accessories also need to account for space constraints and proper cable bend radius (so as to not damage the inside of the cables). This is especially important if thicker CAT6 or CAT6A cables are being utilized. See Application Note 137, *NetShelter SX Cable Management Guidelines* for cable sizing tables regarding cabling accessories (http://www.apcc.com/prod_docs/results.cfm?DocType=App%20Note&Query_Type=10).

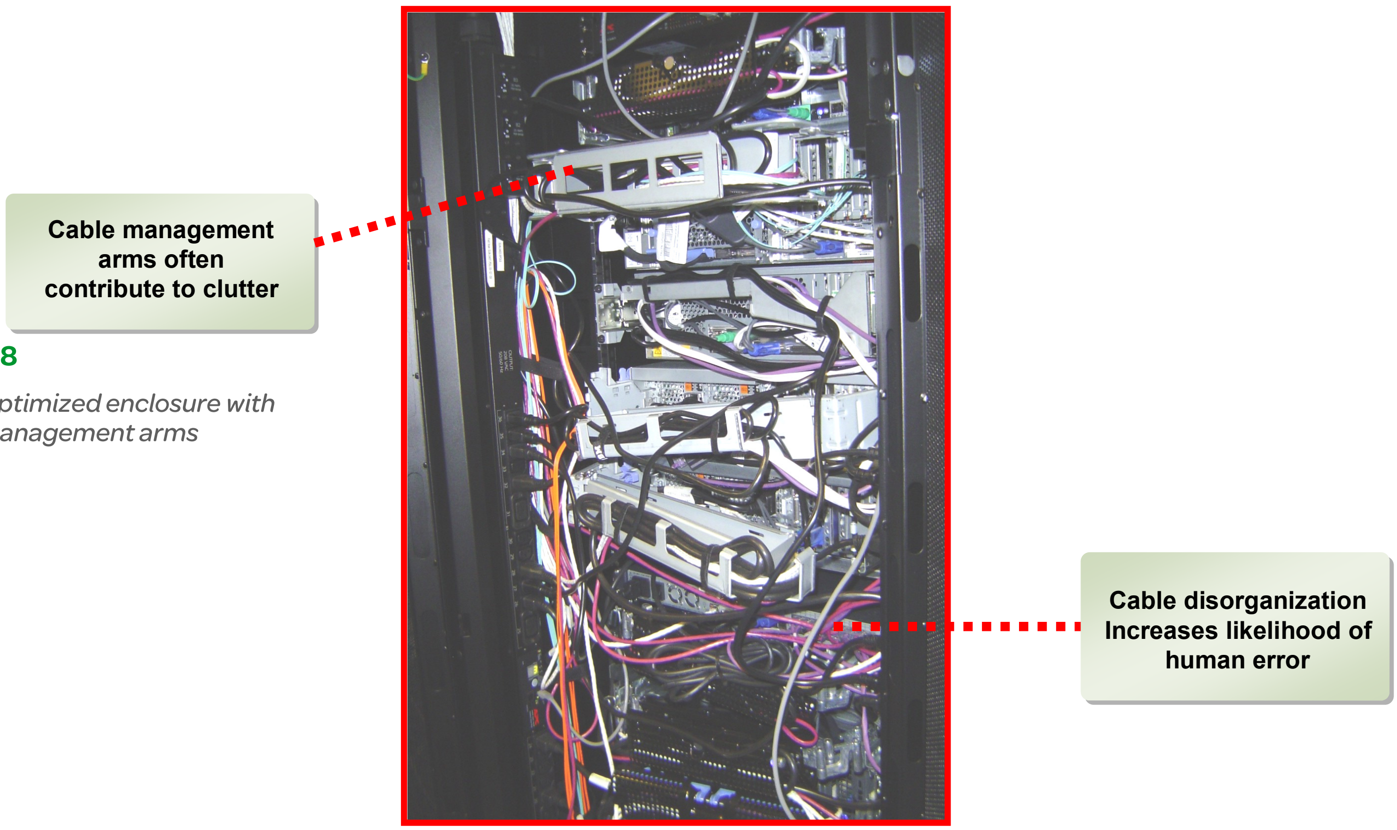

Figure 8

A sub-optimized enclosure with cable management arms

If fiber cables are deployed in the enclosure, the user may need to install fiber spools in the rear of the enclosure to spool the slack on each of the fiber runs. Given the requirement for spools and vertical cable channels, enclosures with significant amounts of fiber and copper cable will need to be wider or deeper than normal enclosures.

If cable lengths are too long, cable management arms can be used to control cable slack at the rear of each server. However, cable management arms have some drawbacks. They can interfere with rack PDUs in the rear of the enclosure. They can also interfere with the exit air path of the servers by creating an "air dam" around the exhaust area of each server (see **Figure 8**).

Networking enclosures

Related resource
White Paper 50
Cooling Solutions for Rack Equipment with Side-to-Side Airflow

Networking enclosures often house one or two switches that distribute cables to each of the enclosures housing the servers. These networking enclosures are wider than normal and have unique needs for cable management because of their much higher cable bundle quantities. The switches within these enclosures have components that are removable from the front such as fan trays. Therefore cables cannot be allowed to block the front of the switch. In addition, switches are frequently cooled from side to side rather than from front to back, so the enclosure must accommodate a method to route cold air to the side of the switch. See White Paper 50, *Cooling Solutions for Rack Equipment with Side to Side Airflow* for more information.

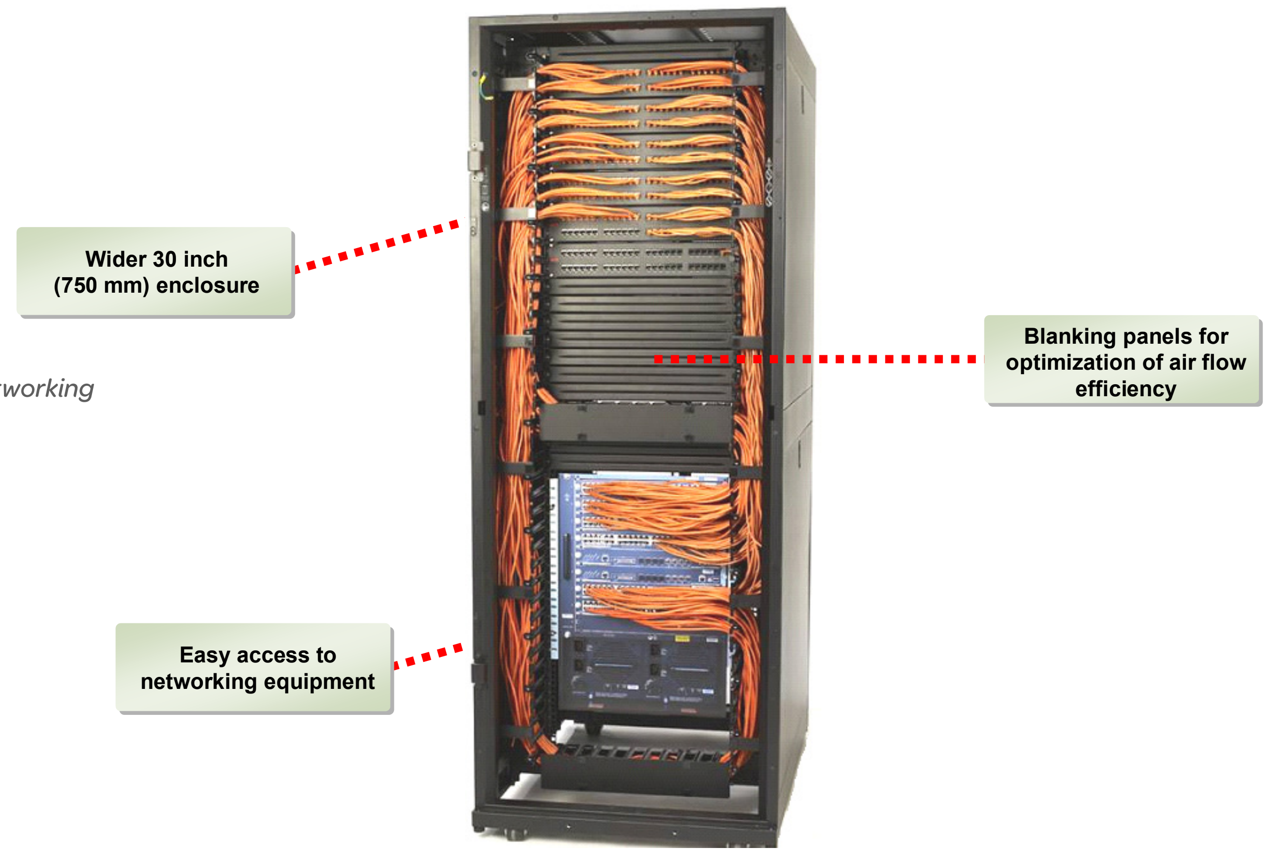

Figure 9
A high density networking application

Step #5: organize for efficient cooling

Most servers mounted in enclosures draw cool air in from the front and exhaust hot air out from the rear of the enclosure. This drives a requirement for the enclosure to have front and rear doors that are perforated. Enclosure doors designed with perforation percentages above 60% of the door's surface area are sufficient for higher density environments.

The enclosure front openings must also be sealed in such a way as to prevent hot exhaust air from being sucked in. A simple solution is to install airflow management blanking panels to cover all unused U spaces (see **Figure 9**). Airflow management blanking panels are tool-less and quick and easy to install. In addition, many enclosures have cutouts or other features to route cables from the front to the rear of the enclosure.

If these air management features are unused, they can introduce access paths for hot air to enter and circulate inside the enclosure. These cutouts must be closed with panels or grommets to optimize for high density air flow patterns.

Figure 10

Hot aisle containment version of close-coupled™ cooling

Some networking equipment components (e.g., routers and switches) circulate cold air and expel hot air in a in a side-to-side pattern. This can present some unique challenges which must be addressed with special equipment. Fan assist modules that pull cold air from the front and route it to the side or special duct kits are examples of equipment that can perform this function.

Enclosures that house power densities of up to 5 KW can be cooled both by utilizing perimeter computer room air conditioning (CRAC) units and by utilizing a hot aisle containment system (see **Figure 10**). A hot aisle containment system can prove effective for lower density enclosures because low density zones actually require the most amount of hot air containment. The further an enclosure is away from an air conditioning unit, the more hot air escapes and mixes with the cold air which has a negative impact on overall data center efficiency. See White Paper 130, *The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers* for more information.

Related resource
White Paper 130
The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers

A common method for higher density cooling (5 kW and above) is to deploy "close-coupled™" cooling. Close coupling implies the placement of the cooling source as close as possible to the servers in the enclosure. This technique can be deployed on either the row level or the individual enclosure level. In fact, hot aisle containment zones inclusive of door and ceiling can be created to house multiple rows of enclosures for optimal control of air temperatures and hot air removal. Go to the APC TradeOff Tool™ Data Center InRow™ Containment Selector for guidance on close-coupled cooling options.

Conclusion

The exercise of fitting more equipment, power, and cooling into data center enclosures can be greatly simplified by following 5 basic steps. Planning and organizing equipment and cable distribution within the enclosure can pay significant dividends in both increased efficiency (lower electrical bills), the ability to safely accommodate higher density (higher uptime and higher productivity per square foot), and for the overall improved look and feel of the data center.

Involving the physical infrastructure / enclosure manufacturer in the process can be helpful, particularly if power, cooling, and cable sizing expertise can be brought to the table.

About the author

Joe Kramer is General Manager of the Racks and Power Distribution Product Segment of Schneider Electric's Critical Power and Cooling Services Business Unit. Joe's 15 years of datacenter infrastructure industry experience includes 11 years of Schneider Electric rack and power distribution product design. Prior to joining Schneider Electric, Joe worked for Toshiba Corporation and for Systems Enhancement Corp and holds a Bachelor's degree in electrical engineering from the University of Missouri.

Resources

Click on icon to link to resource

Browse all white papers
whitepapers.apc.com

Rack Powering Options for High Density
White Paper 29

Cooling Solutions for Rack Equipment with Side-to-Side Airflow
White Paper 50

The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers
White Paper 130

Browse all TradeOff Tools™
tools.apc.com

Power Sizing Calculator
TradeOff Tool 1

Data Center InRow™ Containment Selector
TradeOff Tool 10

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