Life Is On | Schneider Electric

Specifying Data Center IT Pod Architectures

White Paper 260

Revision 0

by Robert Bunger and Patrick Donovan

Executive summary

The desire to deploy IT at large scale efficiently and quickly has forced change in the way physical infrastructure is deployed and managed in the white space. Fully integrated racks complete with IT that roll into place, hard floor data halls, and air containment are just a few of the trends. Designing and deploying IT using standardized blocks of racks (or pods) facilitates these trends. This paper explains how to specify the physical infrastructure for an IT pod and describes optimum configurations based on available power feeds, physical space, and targeted average rack power densities.

Introduction

The growth of large and hyper scale data centers since 2010 has changed the way the IT space is designed and operated. In the past, provisioning for power, cooling, and space based on unknown or quickly changing IT requirements has led to stranded capacity and the ineffective use of the data center facility. In the competitive colocation data center market, for example, building out 1.2 MW data halls without comprehending variability in client rack densities can cause unpredictable cash-flow and underutilization of power and cooling.

In small and medium data centers, the rack or IT enclosure is typically the standard increment of compute that designers use to estimate the requirements for the facility. Knowing average and peak rack density along with the number of racks not only allows for efficient facility sizing, but also determines branch circuit ratings and airflow needed for each rack. The rack standard has also created a simple framework or unit of organization for deploying IT and operating the white space. Electrical faults can be contained and traced to a specific rack. IT moves, adds, and changes are made more efficient by understanding a rack's location and which physical IT equipment and VMs are contained within it, as well as knowing its power and cooling resource dependencies.

For larger data centers there is a demand to deploy IT more quickly and in larger increments. This has led to deploying infrastructure at the pod level. However, there is no similar industry standardization and "black box" method to deploy groups of racks and supporting larger infrastructure into the white space. Many data center operators have developed their own internal standards, but outside of hot-cold aisle configurations, the industry has not adopted a common architecture.

White Paper 160, *Specification of Modular Data Center Architecture* outlines the modules or deployment increments in data centers.

Data Center facility, comprised of
  IT Rooms, comprised of
    IT Pods, comprised of
      IT racks comprised of
        IT devices

In this context, an IT pod is defined as a group of IT racks either in a row or (more typically) a pair of rows, that share some common infrastructure elements like PDU, network router, containment system, air handlers, security, etc. Occasionally the term IT pod is used to refer to what we call in this paper an IT Room; that is not the use in this paper.

Figure 1
Standard IT enclosures and an example of scaled-out pod deployment

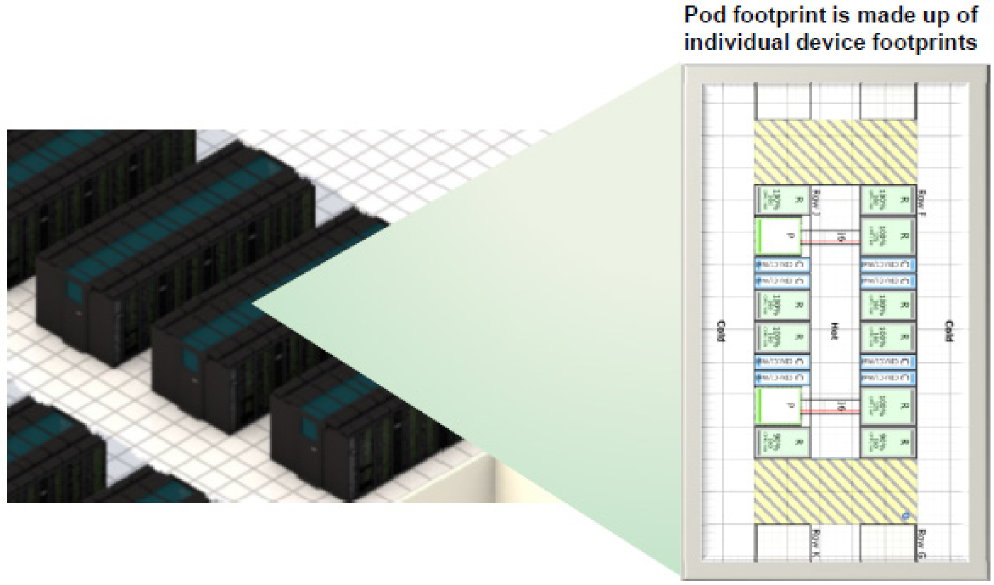

When groups of racks and the commonly available power and cooling resources are examined, optimal deployment sizes and architectures emerge. This paper explains how to specify the physical infrastructure for an IT pod and describes optimum configurations based on available power feeds, physical space, and targeted average rack power densities.

Advantages of pod level deployments and organization

Just as smaller data centers deploy and operate on a rack by rack basis, larger data centers tend to deploy IT using groups of racks at a time, or perhaps a room full at a time. Obviously this is because these deployment sizes match their business need for IT capacity. Racks that share common resources (e.g. a PDU, air containment system, a network switch) are deployed together as a pod. An IT pod can also serve as a logical grouping of business applications or designated to serve a specific client or line of business. An IT pod becomes a useful unit of organization that lies between the room and an individual rack. Knowing what a given pod's shared resources and dependencies are can help with capacity management and planning adds, moves, and changes. And it is at the pod level vs. individual racks where varying infrastructure technologies, architectures, and operational procedures within a site is more easily accomplished.

Deploying and organizing the white space by pod, supported by a dedicated PDU and power feed, makes it easier to vary power and cooling architecture within the same IT hall or room to meet specific business needs. Electrical redundancy can be varied pod by pod depending on the business criticality. Highly critical applications requiring dual feeds can be grouped together separate from less critical loads. Management practices (emergency operating procedures, methods of procedure, etc.) can then be varied pod by pod based on the criticality. Having a clear delineation at the pod level vs. rack by rack or server by server level is simpler, clearer, and therefore reduces risk of confusion and error. Trying to vary these things at a rack by rack level would be more complicated and add more distribution infrastructure overhead.

Organizing and grouping at the pod level also makes it much easier to vary the technologies being used to support the IT. For example, with a pod-based architecture, you could simultaneously power Open Compute Project (OCP) server racks (480/277V), as well as, traditional server racks (208/120V) in the same room or hall.

Another example could be a need to add a high density pod that might require row-based coolers or a rear door heat exchanger that is not needed or being used elsewhere in the room. Or perhaps future server gear might require liquid cooling and therefore, a pod designed to deliver water could be added to a room that had been using an air cooled architecture. This flexibility makes pod-based organization beneficial particularly in uncertain environments where pay-as-you-grow is valued.

Just as organizations standardize and limit the number of server and rack configurations to speed and simplify procurement, deployment, and operations, IT pod designs should be similarly constrained. Starting from a blank sheet of paper each time or allowing for a large multitude of pod sizes and configurations will certainly slow their implementation and complicate maintenance. Organizations should standardize as much as possible the size and overall architecture of the pod design. This standardization allows integrators to prepare product in advance and stock the pods as standard configurations to further reduce lead times.

To further accelerate deployment of IT and to make it easier to roll fully-configured racks in and out, a free-standing pod frame is recommended. A pod frame serves as a mounting point (instead of the rack itself) for air containment systems and services including network, power, and even in some instances, cooling ducts and piping. These can be built out as needed. Services can be brought to the pod without construction and it reduces the need for a raised floor or having to mount anything to the ceiling. **Figure 2** shows an example of a free-standing pod frame system.

Figure 2
Example of a data center pod frame system

To design, standardize, and deploy a pod-based architecture, the question becomes, how big should the pod be? What power capacity? How long should the rows of racks be, and so on. The next section will help answer those questions.

3 main drivers determining pod architecture

Optimizing the pod architecture requires examining the physical space, electrical feed(s), rack density, cooling requirements, and looking for overlaps in cost and flexibility "sweet spots." Network capacity or bandwidth is also a consideration, but this will not be covered in this paper since it tends to have the most flexibility and is highly varliable.

This paper identifies 3 main drivers for determining pod architecture:

1. Choice of electrical feed (power) to the pod
2. Physical space available for a pod (i.e. number of racks)
3. Average rack density required

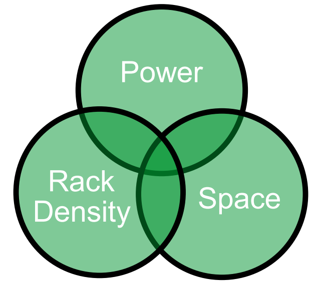

Electrical feed

In typical traditional data centers, bulk power is brought to PDUs or RPPs placed against the wall in the IT space, or strategically located throughout the floor for larger rooms. The power zones (floor area covered by a given power feed) of these PDUs are determined as racks are brought into the IT space and final distribution feeds or whips are run to the racks. Sometimes it is based on a predetermined plan or simply run from the closest PDU. As the data center fills with IT and upper limits on breaker panel space or transformer loading is reached, it is not uncommon to “borrow” breaker space from a PDU that is not physically close to the new rack. This makes operations more complicated and increases the risk of human error during maintenance or unscheduled downtime. Final distribution via busway has alleviated some of this operational risk, and power zones are pre-determined by the physical location of the busway.

Pods have dedicated electrical feeds. Just as a rack typically has dedicated A and B feeds from the immediate upstream PDU, a pod will have the same. The power available to the pod is determined by the voltage and sub-feed distribution amperage. **Table 1** shows the commonly available sub-feed distribution current worldwide.

Table 1
Available power for commonly available voltages and breaker sizes

Voltage	Breaker size (amperes)	Available power (kW)*
400/230 (IEC)	400	275
	250	170
480/277 (ANSI)	400	260
	250	165
208/120 (ANSI)	400	150
	250	70

* Available power numbers are rounded down for simplicity. In this example, 480V and 208V (ANSI voltages) are de-rated to 80% per United States national electric code (NEC).
** Amperage over 400A is available (e.g. 600, 800A) and might be appropriate in the unique requirement of high average rack densities

For the purposes of simplicity, two basic categories of pods – low power and high power – emerge and can be specified as shown in **Table 2** with commonly available voltages and breaker sizes.

Table 2
Typical data center voltages and breaker sizes yield 2 pod types; a low and high power pod

	Available Power (kW)	Voltage	Breaker size (amperes)
Low Power pod	~ 150*	400	250
		480	250
		208	400
High Power pod	~ 250	400	400
		480	400

* 70kW is ignored in this simplification but can have application in small data centers

Narrowing choices in this way can help simplify and clarify what would otherwise be a very complex menu of design alternatives. Standardizing on a smaller number of options can accelerate design and deployment while reducing error, and simplify operations.

An additional electrical consideration for determining pod size is the number of available breaker positions. This can become a constraint for low density pods where there is a higher number of racks for a given power feed.

For the typical IT equipment in data centers, we assume all of the power provided to the pod is turned into heat. That heat must be removed by the cooling system either using air, water, or refrigerant. **Table 3** shows the required volumetric flow of air and water for the two power sizes of pods. These numbers can speed sizing during the design of the facility cooling system and help with more accurate planning during phased buildout of an IT room.

Table 3
Showing the required volumetric flow of air and water for both the low and high power pods

Category	150kW	250kW
Airflow*	8,500 – 11,400 l/s	14,000 – 19,000 l/s
	18,000 – 24,000 CFM	30,000 – 40,000 CFM
Water flow**	7 L/s	11.7 L/s
	111 GPM	185 GPM

* assumes 11C (20F) to 15C (26F) deltaT across IT equipment
** assumes 5.6C (10F) deltaT and 100% sensible cooling and 0% glycol

Space requirements

A pod's footprint is its total width by its length. The width of a pod is determined by the depth of IT racks plus acceptable aisle widths that meet local codes. Length can theoretically be as short as one rack and as long as the length of the entire room, but there are obvious limitations on both ends. Rows that are too long can violate local egress / exit requirements, and short rows are an inefficient use of space. **Figure 3** below offers a simple depiction of space use efficiency starting with a 2 x 20 rack pod to 2 x 2.

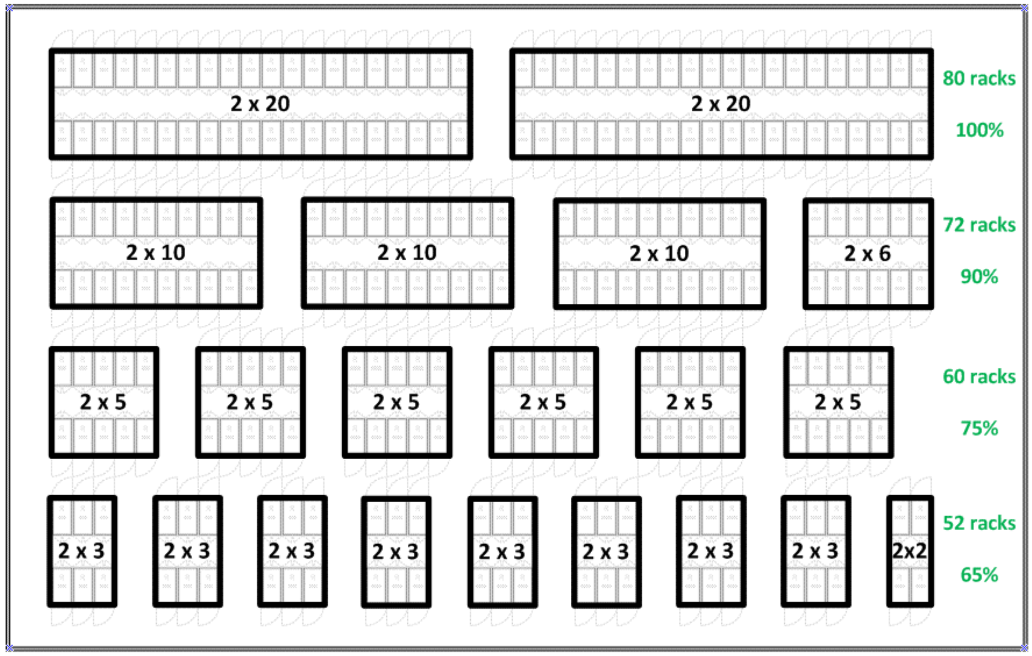

Figure 3
The larger the pod size, the more racks can be placed into a given space (percentages represent the ratio of racks to the base case of 80 racks)

Typical maximum continuous row lengths are between 20 and 25 racks. Longer rows are possible, but typically are only designed and used in hyper scale data centers.

To make the best use of space, design the longest practical pod for the room, balancing power and rack density. For example, a large colocation provider might choose a 250kW 40 rack pod (2 x 20) because it achieves an average rack density of 6kW per rack, which matches their business requirements.

The degree of modularity (i.e., “chunk” or deployment size) required will also be a consideration for determining the pod size. For example, if business conditions are uncertain, a smaller 150 kW pod might be deployed, to minimize capital outlay and reduce exposure to business risk.

Additional considerations

Laying out pods in the data hall is never as easy as placing blocks on a blank sheet of paper. Room shape, building columns, cooling architecture & ducting, pre-existing fire suppression, people access, security requirements, etc. are all important considerations when designing the IT space in a data center. For example, using row-based CRACs, modular PDUs, or pod-based UPSs will obviously impact the size of the pod and must be accounted for.

Rack density

Average rack density within a pod is a simple calculation of the available kW of power, divided by the number of racks. **Table 4** shows corresponding pod sizes for 3 density ranges.

Table 4
Illustrating the different average rack densities for different pod sizes given a particular power rating

	Available Power (kW)	Avg Rack Density (kw)	# of racks (Pod size)
Low Power pod	~ 150*	6	24 (2 x 12)
		12	12 (2 x 6)
		20	6 (2 x 3)
High Power pod	~ 250	6	40 (2 x 20)
		12	20 (2 x 10)
		20	12 (2 x 6)

It's important to note that this is <u>average</u> rack density. Pods with well-engineered air containment can have a large mix of densities. A 20kW rack can operate next to a 2kW rack very easily, as long as the overall pod power limit is not exceeded. The maximum rack density is determined by its corresponding branch circuit(s).

Key takeaway

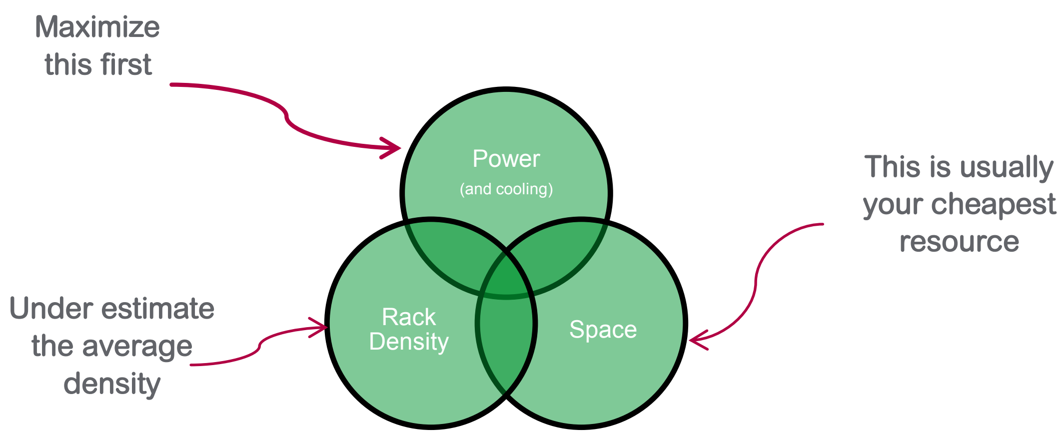

Balancing infrastructure capacity with space has always been a challenge due to rapidly changing IT equipment and business requirements. White Paper 155, *Calculating Space and Power Density Requirements for Data Centers*, discusses this issue in detail. In most cases, designers and operators should <u>under-estimate</u> expected rack density. In this case, if the installed racks operate at a higher density, you "use up" your more expensive resource – power and cooling, before the floorspace. In other words, **it is much costlier to deploy IT *below* the data center design density, than to deploy *above* the design density**. Again, this is true because the cost of space per unit of IT is almost always less than the cost of power and cooling per unit of IT. This tends to be difficult to accomplish since floor space is the only visual representation we have of a data center being "full."

Pod specification

As explained above, the choice of power feed, available space, and average rack power density are the three main drivers for determining the overall pod size and architecture. However, to more fully specify an IT pod for a detailed design concept, there are other pod-level attributes to consider and account for. Using a standardized specification for an IT pod can assist in comparing and contrasting alternatives and can help standardize pods across multiple data center designs and sites. **Table 5** lists the key attributes of a data center IT pod.

Table 5
IT Pod Specification table

Category	Item	Definition
Power	Pod power capability	Based on feeder circuit, panelboard, or pod UPS
	Input voltage & circuit(s)	208/400/415/480 & required feeder circuit to the pod
	Number of feeds to rack and redundancy	1N or 2N
	UPS location & type	Location of UPS: upstream, pod, rack, none
	Transformer	If a pod based/ PDU transformer exists
	Final dist voltage & type	Voltage to the rack
Rack / space requirement	Total racks and row length	Total racks in pod
	Rack type / size	600mm, 800mm, etc
	Inner aisle width	Width of the hot or cold aisle
	Clearance	Required clearance in front of racks and end of pod
	Average rack density	Based on total pod power and # racks
	Peak rack density	Based on max circuit to racks
Network	Architecture	e.g., leaf and spine
Services support structure	Method to deliver power, cooling, and network to the racks within pod	e.g., rack-attached, free-standing pod frame, ceiling-mounted, underfloor, etc
Cooling	Architecture	Air cooled; water cooled: in-row, overhead, rear door; DX; liquid
	Air / water flow	Volumetric flowrate required to remove pod heat
	Containment type	Hot aisle, cold aisle, rack, none
Environment & security	Physical security	Locked doors, cage, proximity sensor, room
	Monitoring	Pod monitoring via IP network / BMS, etc

Recommended practices

Table 6 summarizes a list of recommended practices related to data center IT pods.

Table 6
Recommended practices related to data center IT pods

Best Practice	Description
Pre-specify & standardize pod specs	Standardize design and base it on commonly available voltages, breaker sizes, and space limitations; 150 kW and 250 kW pod sizes emerge as good starting points for design.
Use air containment	Separating hot and cold air streams increases efficiency of cooling plant, prevents hot spots, and offers flexibility in density (i.e. low and high density racks can be next to each other)
Maximize power & cooling resources first	Remember that space is usually your cheapest resource. Its costlier to deploy IT below the design density as you end up stranding the more expensive power and cooling resources
Pre-install higher capacity circuits to racks	Helps eliminate downtime and “hot work” in the event density needs to increase in a given rack; this also provides more flexibility for future growth without a significant cost adder for that capability.
Use free-standing pod frame	Pod frames provide containment and allow for mounting services (network, power & cooling) making it easy to add and remove racks, as well as avoid construction of ceiling mounted supports for services, and facilitating the use of hard floors.
Avoid borrowing circuits from adjacent pods or PDUs	Not maintaining an organized policy of each pod having a dedicated PDU and power feed can lead to confusion and error amongst operators that can lead directly to downtime.

Conclusion

Using a pod based approach in the IT space provides advantages in design, build, and operation, especially for larger data centers that tend to deploy fully configured racks in groups at a time. When commonly available voltages and breaker sizes are considered, optimum pod configurations emerge that make planning and design easier. Standardizing pod designs and limiting the number of configurations can help make pod-level deployments simpler and faster. Organizing IT racks into pods makes it easier to vary power and cooling redundancies and architectures based on the specific business needs within an IT room or hall. A free-standing pod frame system can speed deployments by reducing construction time and simplify maintenance by making it easier to roll racks in and out since containment and services are mounted to the frame instead of the rack. Using the pod specification table in this paper can help project teams clearly specify and repeat designs for current and future projects.

About the authors

Robert Bunger the Director of Data Center Industry Alliances for Schneider Electric's IT Division, is based in West Kingston, RI, USA. In his 18 years at Schneider Electric, Robert has held management positions in customer service, technical sales, product management, and business development. His work has taken him to Europe for two years and Beijing China for three, where he supported the data center solutions business growth in those regions.

Prior to joining Schneider Electric (American Power Conversion before being acquired by Schneider Electric), Robert was a commissioned officer in the U.S. Navy and served eight years in the submarine force.

Robert holds a Bachelor of Science degree in Computer Science from the US Naval Academy and a Master's of Science in Electrical Engineering from Rensselaer Polytechnic Institute (RPI).
Does not like long walks on the beach or quite dinners…

Patrick Donovan is a Senior Research Analyst for the Data Center Science Center at Schneider Electric. He has over 22 years of experience developing and supporting critical power and cooling systems for Schneider Electric's IT Business unit including several award-winning power protection, efficiency and availability solutions. An author of numerous white papers, industry articles, and technology assessments, Patrick's research on data center physical infrastructure technologies and markets offers guidance and advice on best practices for planning, designing, and operation of data center facilities.

Resources

Specification of Modular Data Center Architecture
White Paper 160

Calculating Space and Power Density Requirements for Data Centers
White Paper 155

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